



Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Federal Department of the Environment,
Transport, Energy and Communications DETEC

Swiss Federal Office of Energy
Energy Research and Cleantech

Final report from 17 October 2025

QuickOpt

Hochdynamische Lastoptimierung von Elektrofahrzeugflotten fortgeschrittener Grösse





Publisher:

Swiss Federal Office of Energy SFOE
Energy Research and Cleantech
CH-3003 Berne
www.energy-research.ch

Subsidy recipients:

VEIU Tech GmbH
8703 Erlenbach
veiu-tech.com

Primeo Netz AG
4142 Münchenstein
www.primeo-energie.ch

Quickpac AG
8406 Winterthur
quickpac.ch

Autors:

Philippe, Degen, VEIU Tech GmbH, degen@veiu-tech.com
Giacomo, Pareschi, VEIU Tech GmbH, pareschi@veiu-tech.com
Yannick, Huber, Primeo Netz AG, y.huber@primeo-energie.ch

SFOE project coordinators:

Karin Söderström, karin.soederstroem@bfe.admin.ch
Luca Castiglioni, luca.castiglioni@bfe.admin.ch
Men Wirz, men.wirz@bfe.admin.ch

SFOE contract number: SI/502690-01

The authors bear the entire responsibility for the content of this report and for the conclusions drawn therefrom.



Summary

The QuickOpt pilot project aimed to demonstrate the potential of highly dynamic charge optimisation for large EV fleets to enable load flexibility and provide relief to the grid and the broader energy system. Two large EV fleets with distinct usage patterns were equipped with VEIU Tech's novel charge control solution, which optimises charging based on vehicle and operational data to increase the attainable flexibility. Multiple charging strategies and use cases were tested, targeting different optimisation goals (grid relief or energy services) in distinct setups, while maintaining full, regular fleet operations. An extensive monitoring setup was employed to quantify the impact of highly dynamic charge optimisation on peak loads and other grid-relevant metrics.

The live deployment validated the charge control solution's effective and correct operation, while also highlighting different implementation hurdles that remain prevalent across the charging ecosystem. The live control cases showed that both the car and the charging station implement charging sequences in their own way, sometimes depending on the vehicle's status and sometimes changing with different software releases. This introduces a critical level of unreliability into the control system, which can be a decisive barrier for advanced use cases such as ancillary service provision. Furthermore, the setup demonstrated that access to high-quality energy metering data often requires substantial effort, even though the data is already collected within the internal systems of cars and charging stations. Due to missing standards and regulations, an extensive additional metering setup may be required, adding both complexity and cost to the solution.

The fleets' suitability to the different use cases depended not only on the overall vehicle availability (based on their schedules) but most importantly on the total energy throughput of their charging sessions. While the regional utility's service pool fleet showed limited use and low energy consumption, the investigated delivery fleet exhibited an energy throughput roughly ten times higher. Although this increases the operating costs, it also significantly enhances the savings potential that can be realised with a smart charging system. For this more intensively used delivery fleet, a large share of charging energy could be shifted successfully, improving grid-congestion metrics and enabling potential economic gains in the range of 10 to 32 CHF/EV/month. For the less utilised service pool fleet, only a few valuable control actions per day were feasible, generating minor metric improvements and a weak business case with savings potential of 1 to 7 CHF/EV/month.

A notable exception was the provision of negative ancillary services, which was successfully tested under softened operational tolerances and showed a revenue potential of 28 CHF/EV/month with the lightly used service pool fleet. Even though the revenue potential is significant, the operational setup carries substantial business risk and uncertainty. The interfaces for vehicle charge control remain vulnerable to manufacturer-imposed software implementations and pricing changes; the complex business structure, with multiple stakeholders entitled to a share of the added economic value, further reduces the effective revenue outlook for each involved party.

The project also included the development of a digital twin to evaluate the use cases in a more controlled simulation environment and to explore a broader range of setups. The numerical results generally confirmed the trends and findings observed in the live use cases, thus lending them more validity. The model also revealed the potential of use cases that could not be tested in the field project, such as charging optimisation in the presence of a large PV plant: this scenario, for example, yielded the highest economic potential for the intensively used delivery fleet, achieving 52 CHF/EV/month (while also limiting the peak load).

In general it can be stated that less stringent use cases – such as dynamic tariffs or PV self-consumption optimisation – show limited earnings potential, largely due to current compensation structures. More advanced and strict use cases – such as peak shaving and ancillary services – appear more remunerative but demand higher reliability in flexibility provision. The following table depicts an overview of the main investigated flexibility use cases and how they compare across different metrics that are critical for a successful and profitable smart charging implementation.



Use case	SoC / drive data need	Complexity	Criticality of delivery	Expected operator revenue	Expected benefit for the grid
Local peak shaving	Low	Low-Medium	Medium-High	High	Low-Medium
DSO grid peak shaving	Low	Medium	Medium-High	Medium	High
Ancillary services	High	High	High	High	High
Local PV consumption	Medium	Medium	Low	Medium-High	Medium
Dynamic tariff optimisation	Medium-High	Low	Low	Medium	Medium-High

In conclusion, the project showed that today's charge control systems cannot fully guarantee the discussed reliability due to hard-to-predict vehicle availability and insufficiently standardised control interfaces. Nevertheless, continued progress in standardisation, wider rollout of connected metering, grid digitalisation, and the development of more decentralised flexibility markets and products will very likely unlock the economic case for EV fleet flexibility.

Zusammenfassung

Das Pilotprojekt QuickOpt hatte zum Ziel, das Potenzial einer hochdynamischen Ladeoptimierung für grosse Elektrofahrzeugflotten aufzuzeigen, um Lastflexibilität zu ermöglichen und das Stromnetz sowie das gesamte Energiesystem zu entlasten. Zwei grosse Elektrofahrzeugflotten mit unterschiedlichen Nutzungsmustern wurden mit der neuartigen Ladesteuerungslösung von VEIU Tech ausgestattet, die das Laden auf Basis von Fahrzeug- und Betriebsdaten optimiert, um die erreichbare Flexibilität zu erhöhen. Es wurden mehrere Ladestrategien und Anwendungsfälle getestet, die auf unterschiedliche Optimierungsziele (Entlastung des Stromnetzes oder Energiedienstleistungen) in unterschiedlichen Konfigurationen abzielten, während der reguläre Flottenbetrieb vollständig aufrechterhalten wurde. Mit Hilfe eines umfangreichen Monitoring-Setups wurde der Einfluss der hochdynamischen Ladeoptimierung auf Spitzenlasten und andere netzrelevante Kennzahlen quantifiziert.

Der Live-Einsatz bestätigte die korrekte und wirksame Funktionsweise der Ladesteuerungslösung und machte gleichzeitig verschiedene Implementierungshürden deutlich, die im gesamten Ladeökosystem nach wie vor bestehen. Die Live-Regelfälle zeigten, dass sowohl das Auto als auch die Ladestation Ladesequenzen auf ihre eigene Weise implementieren, manchmal abhängig vom Status des Fahrzeugs und manchmal abhängig von verschiedenen Softwareversionen. Dies führt zu einer kritischen Unzuverlässigkeit des Steuerungssystems, die ein entscheidendes Hindernis für fortgeschrittene Anwendungsfälle wie die Bereitstellung von Systemdienstleistungen darstellen kann. Darüber hinaus hat der Aufbau gezeigt, dass der Zugriff auf hochwertige Energieverbrauchsdaten oft mit erheblichem Aufwand verbunden ist, obwohl die Daten bereits in den internen Systemen von Fahrzeugen und Ladestationen erfasst werden. Aufgrund fehlender Standards und Vorschriften kann eine umfangreiche zusätzliche Messausstattung erforderlich sein, was die Lösung sowohl komplexer als auch kostspieliger macht.

Die Eignung der Flotten für die verschiedenen Anwendungsfälle hing nicht nur von der Gesamtverfügbarkeit der Fahrzeuge (basierend auf ihren Fahrplänen) ab, sondern vor allem vom



Gesamtenergiedurchsatz ihrer Ladevorgänge. Während die Serviceflotte des regionalen Versorgungsunternehmens eine begrenzte Nutzung und einen geringen Energieverbrauch aufwies, war der Energiedurchsatz der untersuchten Lieferflotte etwa zehnmals höher. Dies erhöht zwar die Betriebskosten, erhöht aber auch das Einsparpotenzial, das mit einem intelligenten Ladesystem realisiert werden kann, erheblich. Bei dieser intensiver genutzten Lieferflotte konnte ein grosser Teil der Ladeenergie erfolgreich verschoben werden, wodurch sich die Netzkennzahlen verbesserten und potenzielle wirtschaftliche Gewinne in Höhe von 10 bis 32 CHF/EFz/Monat erzielt werden konnten. Bei der weniger genutzten Service-Pool-Flotte waren nur wenige wertvolle Steuerungsmassnahmen pro Tag möglich, was zu geringfügigen Verbesserungen der Kennzahlen und einem schwachen Business Case mit einem Einsparpotenzial von nur 1 bis 7 CHF/EFz/Monat führte.

Eine bemerkenswerte Ausnahme bildete die Bereitstellung negativer Regelleistung, die unter abgeschwächten betrieblichen Toleranzen erfolgreich getestet wurde und mit der wenig genutzten Service-Pool-Flotte ein Umsatzpotenzial von 28 CHF/EFz/Monat aufwies. Auch wenn das Umsatzpotenzial beträchtlich ist, birgt die betriebliche Umsetzung erhebliche Geschäftsrisiken und Unsicherheiten. Die Schnittstellen für die Fahrzeugladesteuerung sind weiterhin anfällig für vom Hersteller auferlegte Software- und Preisänderungen; die komplexe Geschäftsstruktur mit mehreren Beteiligten, welche alle Anspruch auf einen Anteil am zusätzlichen wirtschaftlichen Wert haben, verringert die effektiven Umsatzaussichten für jede beteiligte Partei zusätzlich.

Das Projekt umfasste auch die Entwicklung eines digitalen Zwillings, um die Anwendungsfälle in einer besser kontrollierten Simulationsumgebung zu bewerten und ein breiteres Spektrum an Konfigurationen zu untersuchen. Die numerischen Ergebnisse bestätigten im Allgemeinen die Trends und Erkenntnisse aus den Live-Anwendungsfällen und verstärkten somit deren Aussagekraft. Das Modell zeigte auch das Potenzial von Anwendungsfällen auf, die im Feld nicht getestet werden konnten, wie z. B. die Ladeoptimierung bei Vorhandensein einer grossen PV-Anlage: Dieses Szenario ergab beispielsweise das höchste wirtschaftliche Potenzial für die intensiv genutzte Lieferflotte und erreichte 52 CHF/EFz/Monat (bei gleichzeitiger Begrenzung der Spitzenlast).

Generell lässt sich feststellen, dass weniger strenge Anwendungsfälle – wie dynamische Tarife oder die Optimierung des Eigenverbrauchs von PV-Anlagen – vor allem aufgrund der aktuellen Vergütungsstrukturen ein eher begrenztes Ertragspotenzial aufweisen. Fortgeschrittenere und strengere Anwendungsfälle – wie Spitzenlastabdeckung und Systemdienstleistungen – scheinen lukrativer zu sein, erfordern jedoch eine höhere Zuverlässigkeit bei der Bereitstellung von Flexibilität. Die folgende Tabelle gibt einen Überblick über die wichtigsten untersuchten Flexibilitätsanwendungsfälle und vergleicht sie anhand verschiedener Kriterien, die für eine erfolgreiche und profitable Umsetzung von Smart Charging entscheidend sind.

Anwendungsfall	SoC / Fahrtdaten notwendig	Komplexität	Kritikalität der Leistungser- bringung	Erwarteter Be- treiberumsatz	Erwarteter Mehrwert für das Netz
Lokales Peak-Shaving	Tief	Tief-Mittel	Mittel-Hoch	Hoch	Tief-Mittel
VNB Netz-Peak-Shaving	Tief	Mittel	Mittel-Hoch	Mittel	Hoch
Systemdienstleistungen	Hoch	Hoch	Hoch	Hoch	Hoch
Lokaler PV Verbrauch	Mittel	Mittel	Tief	Mittel-Hoch	Mittel
Dynamische Tarife	Mittel-Hoch	Tief	Tief	Mittel	Mittel-Hoch



Zusammenfassend hat das Projekt gezeigt, dass die heutigen Lademanagementsysteme aufgrund der schwer vorhersehbaren Fahrzeugverfügbarkeit und unzureichend standardisierter Steuerungsschnittstellen die diskutierte Zuverlässigkeit nicht vollständig gewährleisten können. Dennoch werden weitere Fortschritte bei der Standardisierung, die breitere Einführung vernetzter Messsysteme, die Digitalisierung des Stromnetzes und die Entwicklung dezentraler Flexibilitätsmärkte und -produkte sehr wahrscheinlich die Wirtschaftlichkeit der Flexibilität von Elektrofahrzeugflotten erschliessen.

Résumé

Le projet pilote QuickOpt visait à démontrer le potentiel de l'optimisation de charge hautement dynamique pour les grandes flottes de véhicules électriques (VE), dans le but de permettre une flexibilité de charge et de soulager le réseau ainsi que l'ensemble du système énergétique. Deux grandes flottes de VE présentant des profils d'utilisation distincts ont été équipées de la solution innovante de pilotage de charge de VEIU Tech, optimisant la charge des véhicules afin d'accroître la flexibilité mobilisable. Plusieurs stratégies de recharge et cas d'usage ont été testés, ciblant différents objectifs d'optimisation (soulagement du réseau ou services énergétiques) dans des configurations variées, tout en maintenant le fonctionnement complet et normal des flottes. Un dispositif de suivi étendu a été mis en place pour quantifier l'impact de l'optimisation dynamique sur les puissances de pointe et d'autres indicateurs pertinents pour le réseau.

Le déploiement en conditions réelles a validé le fonctionnement efficace et correct de la solution de pilotage, tout en mettant en évidence diverses difficultés d'implémentation encore courantes dans l'écosystème de la recharge. Les essais en pilotage direct ont montré que tant le véhicule que la borne, exécutent les séquences de charge à leur manière, parfois en fonction de l'état du véhicule, et parfois différemment selon les versions logicielles. Cela introduit un niveau critique d'incertitude dans le système de contrôle, pouvant constituer un obstacle déterminant pour des cas d'usage avancés tels que la fourniture de services auxiliaires. De plus, le projet a montré que l'accès à des données de mesure énergétique de haute qualité nécessite souvent un effort considérable, bien que ces données soient déjà collectées au sein des systèmes internes des véhicules et des bornes. En raison de l'absence de normes et de réglementations adaptées, une instrumentation de mesure additionnelle importante peut être nécessaire, ajoutant complexité et coûts.

La pertinence des flottes pour les différents cas d'usage dépendait non seulement de la disponibilité générale des véhicules (selon les horaires), mais surtout du volume total d'énergie transférée lors des sessions de charge. Alors que la flotte de service de l'entreprise d'utilité régionale présentait une utilisation limitée et une faible consommation énergétique, la flotte de livraison étudiée affichait un flux énergétique environ dix fois supérieur. Bien que cela augmente les coûts d'exploitation, cela renforce également le potentiel d'économies réalisables grâce à un système de charge intelligent. Pour cette flotte de livraison plus sollicitée, une grande part de l'énergie de charge a pu être déplacée avec succès, améliorant les indicateurs liés à la congestion du réseau et permettant des gains économiques potentiels de l'ordre de 10 à 32 CHF/VE/mois. Pour la flotte de service moins utilisée, seules quelques actions de pilotage réellement utiles par jour étaient possibles, générant des améliorations limitées et un modèle économique faible, avec un potentiel d'économie de 1 à 7 CHF/VE/mois.

Une exception notable concerne la fourniture de services auxiliaires négatifs, testée avec des tolérances opérationnelles assouplies et qui a montré un potentiel de revenus de 28 CHF/VE/mois pour la flotte de service faiblement utilisée. Bien que ce potentiel soit significatif, la configuration opérationnelle comporte des risques commerciaux importants et une forte incertitude. Les interfaces de pilotage de charge restent vulnérables aux implémentations logicielles imposées par les constructeurs et aux évolutions tarifaires ; la structure commerciale complexe, avec plusieurs parties prenantes se partageant la valeur ajoutée, réduit encore la perspective de revenus nets pour chaque acteur.



Le projet incluait également le développement d'un jumeau numérique permettant d'évaluer les cas d'usage dans un environnement simulé plus maîtrisé et d'explorer un éventail plus large de configurations. Les résultats numériques ont généralement confirmé les tendances et conclusions observées lors des essais en conditions réelles, renforçant leur validité. Le modèle a aussi révélé le potentiel de cas d'usage non testés sur le terrain, tels que l'optimisation de la charge en présence d'une grande installation photovoltaïque (PV): ce scénario a, par exemple, généré le potentiel économique le plus élevé pour la flotte de livraison intensivement utilisée, atteignant 52 CHF/VE/mois (tout en limitant la puissance de pointe).

De manière générale, il apparaît que les cas d'usage moins exigeants – tels que les tarifs dynamiques ou l'optimisation de l'autoconsommation PV – présentent un potentiel de revenus limité, principalement en raison des mécanismes de compensation actuels. Les cas d'usage plus avancés et stricts – tels que le lissage des pics de puissance (peak shaving) et les services auxiliaires – semblent plus rémunérateurs mais nécessitent une fiabilité accrue dans la fourniture de flexibilité. Le tableau suivant donne un aperçu des principaux cas d'usage étudiés et les compare selon plusieurs indicateurs essentiels pour une mise en œuvre réussie et rentable de la recharge intelligente.

Cas d'usage	Besoin en données SoC / conduite	Complexité	Criticité de l'exécution	Revenus attendus pour l'opérateur	Bénéfice attendu pour le réseau
Lissage local des pics	Faible	Faible-Moyen	Moyen-Élevé	Élevé	Faible-Moyen
Lissage des pics réseau (DSO)	Faible	Moyen	Moyen-Élevé	Moyen	Élevé
Services auxiliaires	Élevé	Élevé	Élevé	Élevé	Élevé
Autoconsommation PV locale	Moyen	Moyen	Faible	Moyen-Élevé	Moyen
Optimisation par tarifs dynamiques	Moyen-Élevé	Faible	Faible	Moyen	Moyen-Élevé

En conclusion, le projet a montré que les systèmes de pilotage de charge actuels ne peuvent pas encore garantir pleinement le niveau de fiabilité requis, en raison d'une disponibilité des véhicules difficile à prévoir et d'interfaces de contrôle insuffisamment standardisées. Néanmoins, les avancées continues en matière de normalisation, le déploiement accru de dispositifs de mesure connectés, la numérisation du réseau, ainsi que le développement de marchés et de produits de flexibilité plus décentralisés devraient très probablement débloquent le modèle économique associé à la flexibilité des flottes de VE.



Contents

Summary	3
Zusammenfassung.....	4
Résumé.....	6
List of figures.....	10
List of abbreviations	14
1 Introduction.....	15
1.1 Context and motivation.....	15
1.2 Project objectives	16
2 Approach and method	18
2.1 Project framework.....	18
2.2 System deployment.....	19
2.2.1. IT infrastructure.....	19
2.2.2. VEIU Tech system implementation.....	21
2.2.3. Grid load monitoring equipment.....	26
2.2.4. Fleet booking systems	32
2.2.5. Ancillary services link.....	33
2.3 Live use cases.....	33
2.3.1. Local (building) peak shaving	34
2.3.2. DSO grid peak shaving	35
2.3.3. Ancillary services	35
2.3.4. Local PV consumption.....	36
2.3.5. Voltage stabilisation.....	37
2.3.6. Dynamic tariff optimisation.....	37
2.3.7. Peak shaving classification.....	38
2.4 Simulation model.....	39
2.4.1. Scope and goal.....	39
2.4.2. Optimisation problem.....	39
2.4.3. Fleets characterisation.....	41
3 Results and discussion	47
3.1 Operational implication and user experience	47
3.1.1. Fleet status and charge process awareness	47
3.1.2. User experience and operational implications	49
3.2 Live use cases.....	52
3.2.1. Local (building) peak shaving	52
3.2.2. DSO grid peak shaving.....	58



3.2.3.	Ancillary services	64
3.2.4.	Voltage stabilisation	79
3.2.5.	Local PV consumption	79
3.2.6.	Dynamic tariff optimisation	84
3.3	Economic evaluation of live use cases	88
3.3.1.	Cost references	88
3.3.2.	Cost comparison	88
3.3.3.	Fleet implications	90
3.3.4.	Use case-specific considerations	90
3.4	Digital twin simulations	91
3.4.1.	Fleet load peak shaving	93
3.4.2.	Local (building) peak shaving	95
3.4.3.	DSO grid peak shaving	103
3.4.4.	Local PV consumption	109
3.4.5.	Summary of digital twin results	115
4	Conclusions and outlook	117
4.1	Deployment findings	117
4.2	Use case findings	119
4.3	General conclusions	121
4.3.1.	Implementation assessment	121
4.3.2.	Flexibility evaluation	121
4.3.3.	Economic validation	122
4.3.4.	Use case comparison	122
4.4	DSO's take	123
4.5	Outlook	123
5	Appendix	125
5.1	Back-of-the-envelope use cases evaluation	125
5.2	Digital Twin	128
5.2.1.	Fleet load peak shaving	128
5.2.2.	Local (building) peak shaving	128
5.2.3.	Local PV consumption	133
5.3	Complete Primeo Energie's pool fleet user survey	134



List of figures

Figure 1. Schematic organisation of the project and roles of each partner.	18
Figure 2. Project partner fleets.	19
Figure 3. Overview of the data infrastructure employed in the project for both monitoring and operational functions.	20
Figure 4. Schematic overview of the main components of the charge control system by VEIU Tech and their interactions.	22
Figure 5. Charge controller hardware.	23
Figure 6. Early prototype of the charge control adapter developed for the project. This prototype includes the Charge Controller board, a low-amperage communication wire connected to the Control Pilot (CP) line of the charging cables, and the five primary high-amperage conductors for the 3-phase power connection.	25
Figure 7. Deployed plug adapter with charge controller at an EVSE in Münchenstein.	25
Figure 8. Arrangement of the different load metering devices at various grid levels at the two project sites.	27
Figure 9. Janitza UMG-605 used for parking load monitoring.	28
Figure 10. Smart meter monitoring in Münchenstein with EnergyX device in the background.	29
Figure 11. Power pulse monitoring of the building meter in Hägendorf.	30
Figure 12. Mobile grid measurement unit installed at the level 6 transformer in Münchenstein.	31
Figure 13. Example of trip bookings on two days in Münchenstein.	32
Figure 14. Percentage of days an EV from Primeo Energie's pool fleet undertook one or more journeys and percentage of days the vehicles spent the night at a parking location other than the managed depot.	42
Figure 15. Distributions of departing times, arrival times, and trip durations by type of journey for the PE pool fleet.	43
Figure 16. Distributions of SoC at departure and arrival, and distribution of typical trip discharge for the PE pool fleet.	44
Figure 17. Duration (x-axis) and SoC discharge (y-axis) for every journey of the PE fleet by type of trip.	44
Figure 18. Share of days with 0, 1 or 2 journeys among eVitos in QP delivery fleet.	45
Figure 19. Distributions of departing times, arrival times, and trip durations for the QP fleet.	45
Figure 20. Distributions of SoC at departure and arrival, and distribution of typical trip discharge for the QP fleet.	46
Figure 21. Joint distribution of trip discharge and duration for the QP fleet.	46
Figure 22. Comparison between original data and samples from fitted functions for key quantities.	46
Figure 23. Live fleet dashboard for the Primeo Energie pool fleet.	48
Figure 24. Primeo Energie pool fleet user survey question 5.	50
Figure 25. Primeo Energie pool fleet user survey question 12.	51
Figure 26. Primeo Energie pool fleet user survey question 15.	51
Figure 27. Building load and charging power of EVs under management on various Friday-Saturday nights.	53



Figure 28. Aggregated fleet SoC on various Friday-Saturday nights.	54
Figure 29. Daily building peak loads in Hägendorf; coloured points indicate days with charge controls.	54
Figure 30. Building peak load (y-axis) and total electricity consumption (x-axis) for each workday of the peak shaving use case in Hägendorf.	55
Figure 31. Current flow monitored at the level 5 grid feeder to the Hägendorf site on the five tested weekdays. Clear spikes on the hour are visible especially during the night.	58
Figure 32. Charging load of the EV fleet in Hägendorf on the five tested weekdays: interventions on the portion of controlled EV are visible at 21:00 and 00:00.	59
Figure 33. Breakdown of level 6 transformer load in Münchenstein by building/consumer for a typical week. Loads are stacked and the high red peaks are due to fast charging.	60
Figure 34. Example of charge control intervention on three EVs to balance the sudden load caused by the external fast charging stations.	61
Figure 35. For ever calendar week: mean and maximum transformer load (lines and left y-axis); count of 5-min intervals with load above a certain threshold (markers and right x-axis).	61
Figure 36. Test signal curve with timing margins for secondary control given by Swissgrid for prequalification.	66
Figure 37. Ramps followed by available pool car for negative secondary control. Upper graph: Ancillary power request; middle graph: load measurement data from garage meter; lower graph: pool fleet vehicle status.	67
Figure 38. Average "connected" power available for negative control (charging) in Münchenstein.	69
Figure 39. Number of received activation requests by direction and duration: on the left, across all time durations; on the right, zoomed in to durations shorter than two minutes.	70
Figure 40. Frequency of activation requests by direction and day of the week (left) or hour of the day (right).	70
Figure 41. Average active power by control direction and day of the week (left) or hour of the day (right) over the entire time period (i.e. including no activation times).	71
Figure 42. Average "connected" power available for negative control (charging) with the PE pool fleet in Münchenstein. In orange at 15-minute resolution; in blue, after taking the minimum availability within each 4-hour block.	73
Figure 43. Depending on the length of the continuous availability period (x-axis), the y-axis shows the percentage of periods with N cars available (colours) for negative control with the PE pool fleet.	73
Figure 44. Average prices for TRL daily auctions (2025) by direction and day of the week (left) or time block of the day (right).	74
Figure 45. Average net earnings of ancillary service providers for tertiary regulation energy (2025) by direction and day of the week (left) or time of the day (right).	75
Figure 46. Average amount of tertiary reserve capacity activated by direction and day of the week (left) or time of the day (right).	76
Figure 47. Hourly profiles of all factors determining the profitability of the PE pool fleet to provide TRL-/TRE-: kW of available negative power (blue area, left y-axis), % of activation time (red stars, left y-axis), capacity price (green line, right y-axis), and energy price (yellow line, right y-axis).	77
Figure 48. Gross earnings for the ASP managing the PE fleet (12 EVs), by time block and source of income (TRL-/TRE- in 2025).	77



Figure 49. Gross earnings from different use cases of ancillary service provision via dynamic charging of EV fleets.....	78
Figure 50. Example of static charge control (algorithm 1) on two EVs before and after the activation of the "PV charging window".....	80
Figure 51. Example of threshold-based charge control (Algorithm 2), where three EVs start and stop the charging process depending on the power output of the PV plant.	81
Figure 52. Lines and left y-axis: weekly own-consumption and autarky rates. Bars and right y-axis: weekly solar generation and direct own-consumption with EVs. Live charge management weeks are highlighted in green.	82
Figure 53. Lines and left y-axis: weekly own-consumption and autarky rates. Bars and right y-axis: weekly solar generation and direct own-consumption with EVs. Live charge management weeks are highlighted in green.	83
Figure 54. Typical daily profiles and spread of: 1. SwissIX sport market price, 2. VARIO grid dynamic tariff, 3. Quickpac fleet charging load, 4. Primeo Energie fleet charging load. Both fleets are modelled with 50 EVs and their maximum load shift potential is shown in green.....	85
Figure 55. Charge shifting for dynamic price optimisation for an example day of the live use case in Münchenstein with the relative savings per car per kWh.....	87
Figure 56. Financial breakdown of revenue for the different live use cases in the two project locations and the two project fleets.....	89
Figure 57. Example of output of the digital twin: for each EV we obtain SoC profile (orange lines) and net battery power (blue lines). The last red line indicates the sum of all EV charging loads. The grey shaded areas represent the charging stops. Only three individual EVs are shown for clarity.	91
Figure 58. Normalized peaks of QP fleet charging load (kW/EV) for different fleet sizes and charging strategies. Charging rate of a single EVSE is 11 kW and simulation window is one week.	93
Figure 59. Normalized peaks of PE fleet charging load (kW/EV) for different fleet sizes and charging strategies. Charging rates of the fleet vary between 6 and 11 kW; simulation window is one week.....	94
Figure 60. Average peak loads obtained for the Münchenstein site with 50 EVs from the PE fleet for different technological setups (x-axis) and algorithmic strategies (colour shades).....	97
Figure 61. Monthly savings from peak load reduction (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein site with 50 PE EVs. Different colours indicate distinct technological or control setups (only heuristic management is shown). Small points represent results from individual stochastic simulations, while the rhomboids the average per setup. The rhomboids with the same colour depict the same setup with varying battery sizes (see Figure 60 for reference).	98
Figure 62. Average peak loads obtained for the Hägendorf site with 50 EVs from the QP fleet for different technological setups (x-axis) and algorithmic strategies (colour shades).	99
Figure 63. Monthly savings from peak load reduction (y-axis) vs. monthly costs for load management solution (x-axis) for the Hägendorf site with 50 QP EVs. Different colours indicate distinct technological or control setups (only heuristic management is shown). Small points represent results from individual stochastic simulations, while the rhomboids the average per setup. The rhomboids with the same colour depict the same setup with varying battery sizes (see Figure 62 for reference).	100
Figure 64. Average peak loads obtained for the Münchenstein site with 50 EVs from the QP fleet for different technological setups (x-axis) and algorithmic strategies (colour shades).....	101
Figure 65. Monthly savings from peak load reduction (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein site with 50 QP EVs. Different colours indicate distinct technological or control setups (only heuristic management is shown). Small points repres-	



- ent results from individual stochastic simulations, while the rhomboids the average per setup. The rhomboids with the same colour depict the same setup with varying battery sizes (see Figure 64 for reference). 102
- Figure 66. Typical daily load profiles and spread of the L5 feeders used in the DSO grid assessment (top two plots) and of the two project fleets (each with 50 EVs, bottom two plots). For the fleets, the maximum load shift potential is shown in green. 104
- Figure 67. Peak loads resulting from the gradual introduction of QP-like fleets (50 EVs each) in the Högendorf L5 grid. Uncontrolled peaks are shown in light blue, while optimal charge management in violet. Only for the scenario with 5 fleets, the additional impact of batteries is explored: either alone (green) or combined with smart charging strategies (orange). 105
- Figure 68. Monthly costs for grid peak load (y-axis) vs. monthly costs for load management solution (x-axis) for the Högendorf L5 grid with a growing number of QP fleets (50 EVs each). Each arrows series represents the economic development under the same load management approach; all results are obtained with optimal control. 106
- Figure 69. Peak loads resulting from the gradual introduction of QP-like fleets (50 EVs each) in the Münchenstein L5 grid. Uncontrolled peaks are shown in light blue, while optimal charge management in violet. Only for the scenario with 5 fleets, the additional impact of batteries is explored: either alone (green) or combined with smart charging strategies (orange). 107
- Figure 70. Monthly costs for grid peak load (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein L5 grid with a growing number of QP fleets (50 EVs each). Each arrows series represents the economic development under the same load management approach; all results are obtained with optimal control. 108
- Figure 71. Breakdown of monthly electricity costs with different load management solutions at the Münchenstein site with 40 PE pool EVs and a 150 kW PV plant. Right y-axis and stars indicate the PV own-consumption rate achieved in each case. 111
- Figure 72. Net monthly costs for electricity procurement (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein site with 40 PE EVs and PV plants of different capacity (individual plots). Points with the same colour refer to the same setup with varying battery sizes (respectively 150, 300, and 450 kWh). 112
- Figure 73. Breakdown of monthly electricity costs with different load management solutions at the Högendorf site with 40 QP delivery EVs and a 150 kW PV plant. Right y-axis and stars indicate the PV own-consumption rate achieved in each case. 113
- Figure 74. Net monthly costs for electricity procurement (y-axis) vs. monthly costs for load management solution (x-axis) for the Högendorf site with 40 QP EVs and PV plants of different capacity (individual plots). Points with the same colour refer to the same setup with varying battery sizes (respectively 150, 300, and 450 kWh). 114
- Figure 75. Summary of net economic benefits for a few key illustrative use cases as obtained from the digital twin simulations. The different columns correspond to different use cases and representative setups; the different colours to differ algorithmic methods. 116



List of abbreviations

AC	Alternating Current
API	Application Programming Interface
ASP	Ancillary Service Provider
BESS	Battery Energy Storage System
CP	Control Pilot
DC	Direct Current
DSO	Distribution System Operator
EV	Electric Vehicle
EVSE	Electric Vehicle Supply Equipment
FC	Fast Charging
GLPK	GNU Linear Programming Kit
IEC	International Electronics Commission
IoT	Internet of Things
IT	Information Technology
LED	Light Emitting Diode
LEG	Lokale Elektrizitätsgemeinschaft
MILP	Mixed-Integer Linear Programming
MQTT	MQ Telemetry Transport
OBD	On-Board Diagnostics
PCB	Printed Circuit Board
PDF	Probability Distribution Function
PE	Primeo Energie (fleet)
PV	Photovoltaic
QP	Quickpac (fleet)
SFOE	Swiss Federal Office of Energy
SoC	State of Charge
THD	Total Harmonic Distortion
TSO	Transmission System Operator
VPP	Virtual Power Plant
VT	VEIU Tech GmbH
ZEV	Zusammenschluss zum Eigenverbrauch



1 Introduction

1.1 Context and motivation

The future development in the energy system is leading towards a move to renewable energy production and general electrification on the consumer side. This trend enables a higher share of renewable energy in the energy mix as well as an increase in energy efficiency and therefore a reduction in the net energy consumption.

However this transition brings different changes to the current energy system that will need counter-measures and/or adaptations in order to keep the system running reliably. In general it can be stated that a new time setup will be dominant that will raise the need to use energy flexibility in order to align the consumption with the production and the energy infrastructure capacity limits.

Next to seasonal adjustments due to the imbalance in energy availability and consumption throughout the seasons, short term flexibility will be needed to synchronize the short term fluctuations in energy production as well as in the consumption. Further flexibility will be needed to cope with short term accumulations of consumption and production (especially solar) in the distribution grids, since they will face a reasonable increase in usage due to aforementioned reasons. Flexibility can therefore shave specific usage peaks and hence increase the efficiency and the use of capacity for cost efficient operations and reduce the time pressure for capacity expansion in the grid.¹

This flexibility always comes with the ability to store energy from one moment of good availability to a moment of high need. This can either be in the form of thermal storage of a cooling / heating system or in the most classic way of directly storing electric energy in chemical batteries.

Here the upcoming electric mobility is playing a vital role, since electric cars come along naturally with batteries that store a significant amount of energy, that is needed in order to make larger distance trips as used nowadays in transportation. This enables to store electric energy and therefore shift charging times to match good energy and grid capacity availability.²

The optimisation is however multi-layered so that different grid levels as well as production objectives do not have to align time-wise. Today already some simple flexibility systems for electric vehicles, also called smart charging system, are in use, focussing on the most imminent 'grid' capacity limit, being the building connection to the grid. The other upcoming flexibility system is the self-consumption optimisation, that targets charging for high – local – energy availability to the system.

This optimisation is however often done in a simple and limited manner because either (1) there is no need for further optimisation, (2) the technical setups do not allow it in an efficient way or (3) the regulatory incentives are not given to run such a system. These optimisations often have a wide range of possible optimal control actions for a given time window that lead to an equivalent local outcome, whereas for other energy availability and/or grid capacity situations a further optimisation would indeed be beneficial.³

These more granular optimisations evidently need a more granular storage state and time constraint knowledge since a tighter optimisation schedule is aimed for. This is the point where most of the current equipped and used systems come to their limits. Especially the state-of-charge (SoC) of EV batteries and the effective readiness time are often not available and therefore not considered. Those parameters are however eminent for an elaborate optimisation since they effectively define the flexibility.

Those parameters are often not available due to missing technical standards and probably protectionist tendencies of charging infrastructure and car manufacturer. Further the time component in the time-to-market process which is in conflict with a standardisation process seems to play a vital role in this early

¹ IEA: Electricity Grids and Secure Energy Transitions, 11.2023

² Di Natale L, Funk L, Rüdisüli M, Svetozarevic B, Pareschi G, Heer P, Sansavini G. The Potential of Vehicle-to-Grid to Support the Energy Transition: A Case Study on Switzerland. *Energies*. 2021; 14(16):4812. <https://doi.org/10.3390/en14164812>

³ Group for Sustainability and Technology - ETH Zürich: Enabling Flexible Electric Vehicle Grid Integration – ErVIn, 01.2023



phase of EV market penetration. Also it seems that car manufacturer do not (yet) have the perspective of seeing the car as a component in the energy system, but more as they saw it traditionally - a pure transportation device. This tendency leads to setups with missing remote data and charge control access of cars as well as incompatibility with certain desired charge processes on the vehicle side. Another conceptual observation is that charging stations are often considered the flexibility providers, especially because they seem to have a better access to charge control. However the nature of it being a “smart fuse” shows its limits when missing actual state information of the underlying energy storage system – the car battery.⁴⁵

Looking at the time constraints that come along with the usage of electric vehicles that have a major impact on the available flexibility, it shows to be a difficult constraint or objective to grasp. On the one side private car owner do not effectively schedule trips but mostly do it in a spontaneous and undocumented way. On the other side commercial cars are used along the business needs that can vary significantly depending on the work profile. Nevertheless business cars and especially commercial fleets do show quite distinct usage patterns and often use an IT tool to organize the operations such that a structured integration of that information into a flexibility system seems handier and more fruitful. In addition to that the larger number of vehicle inside a fleet also adds to the storage capacity and therefore the flexibility amount that can be accessed and controlled for – also local – system optimisations.

1.2 Project objectives

To address the previously identified challenges, this project has been initiated to implement a comprehensive smart charging control system for commercial EV fleets and evaluate its effectiveness in providing flexibility services to the grid. The specific project objectives were as follows:

- Verification of a connected charge control system: This objective involved verifying the seamless operation of a connected charge control system that can dynamically adjust the charging power of each EV while accessing real-time data on each vehicle's state of charge. The system, developed by VEIU Tech GmbH, was evaluated based on its ability to operate continuously, autonomously, and without interruption across various fleet configurations with diverse EV and electric vehicles supply equipment (EVSE) brands.
- Implementation and testing of flexibility services: The project aimed to implement and test at least three distinct flexibility services. This included identifying and resolving any issues that arose during the deployment of each service. Once reliable operation was confirmed, a qualitative and quantitative assessment of the benefits achieved through each flexibility service was conducted.
- Assessment of optimal deployment conditions: By deploying the flexibility services across different fleet and grid setups, the project sought to determine the conditions where these services can provide the main benefit and can be most leveraged. The goal was to identify scenarios where these services are most effective in alleviating stress on the local grid.
- Development of a digital twin for expanded assessment: To complement the live tests, a digital twin had been set up to simulate fleet-grid-flexibility service configurations that were not feasible in the physical project due to technical or operational limitations. This model enabled a broader evaluation of the potential grid relief that smart charging services can offer. The physical deployed tests served as a validation basis to ensure the reliability of the digital simulations.
- Analysis of economic viability and recommendations: Finally, the project leveraged findings from both the live flexibility use cases and the digital simulations to formulate insights and recommendations on the economic aspects of EV-based flexibility services. For instance, the analysis explores whether current power pricing structures provide adequate incentives for load management sys-

⁴ EUROPEAN COMMISSION -Directorate-General for Mobility and Transport: Access to In-vehicle Data and Resources, 05.2017.

⁵ Michele Bertonecello, C.M., T.M, T.S.: Unlocking the full life-cycle value from connected-car data, 02.2021.



tems to optimise grid relief or whether dynamic tariff structures could enhance the potential for grid relief.

By accomplishing these objectives, the project aims to demonstrate the practical benefits of smart charging control systems in commercial EV fleets, effectively providing flexibility services to the grid and alleviating infrastructure stress. The insights gained will contribute to developing more efficient and sustainable energy management strategies for future fleet and grid integrations.



2 Approach and method

2.1 Project framework

The project partners cover the full value chain of the smart charging ecosystem, ranging from EV fleet operators to grid operators and ancillary services brokers. In addition, the live tests have taken place at two separate sites: Högendorf (SO) and Münchenstein (BL). The following figure shows the main roles of each partner in the project.




Figure 1. Schematic organisation of the project and roles of each partner.

Specifically, the responsibilities of each project partner were distributed in the following way:

- VEIU Tech GmbH assumed the role of project lead, overseeing all aspects of project management including planning, execution, monitoring, and reporting. Their responsibilities encompassed coordinating activities with the other partners, managing timelines, and ensuring that project objectives were met in alignment with the project scope and milestones. In addition, VEIU Tech was responsible for the development, deployment, and operation of the charge control system. This involved both software development, and hardware integration to ensure compatibility with the various fleet configurations partaking in the project. Moreover, VEIU Tech managed data collection, storage, and analytics, employing the set up data infrastructure both for live monitoring of the system response and for elaboration of the achieved project results. Finally, they were responsible for developing a digital twin of the project setup and running numerical simulations that generated broader insights on the benefits of EV fleet charge management for the grid.
- Primeo Netz AG (hereafter “Primeo Energie”) was responsible for higher grid-level operations and monitoring at both project sites, functioning as the interface between the project and the broader electrical grid infrastructure. They additionally managed and monitored the local grid in Münchenstein, where they allowed and implemented multiple *ad-hoc* live load measurements to track the project evolution and results. Primeo Energie also served as the EV fleet operator in Münchenstein, where they facilitated the implementation of VEIU Tech’s charging control system within their fleet, coordinating scheduling and operational requirements. Finally, Primeo Energie shared their role as pool aggregator for ancillary services and other valuable grid flexibility services.
- Quickpac AG served as the EV fleet operator in Högendorf, responsible for the day-to-day management and operation of the fleet. They collaborated closely with the project team to implement VEIU Tech’s charge control system, ensuring that vehicles were always available for charging con-



trol interventions as required. They provided operational insights, facilitated communication between the fleet users and the project team, and assisted in power and fleet data collection efforts.



Figure 2. Project partner fleets.

2.2 System deployment

2.2.1. IT infrastructure

Overview

In the following chart the IT setup of the QuickOpt project is depicted. The general concept is a modular structure where different 'services' connect to a central entity in order to exchange data for various purposes. The services include the installed vehicle devices, the metering data as well as web services that connect to online data points such as data APIs for example. The data exchange with the central server was chosen to use the MQTT publish-subscribe protocol, that brings along multiple advantages for this IoT-based setup found in this project.




Figure 3. Overview of the data infrastructure employed in the project for both monitoring and operational functions.

The main principle of the central MQTT server was to act as a so-called MQTT broker that coordinates the data exchange between the different clients that connect to it. The server is considered a broker because it does not principally store the data, but only forwards data to clients that subscribed to the respective data point. The server stored data points when they were specially flagged when published to the server in order that upcoming subscribers would get 'the most recent' version of that data point.

This MQTT architecture enabled to keep the data exchange modular and scalable as well as to integrate with state-of-the-art software applications that support MQTT data handling. Further it can be stated that this architecture also reduced the effort for the security of the setup since all communications are encrypted to the MQTT server and authentication is also handled by it. It also reduced the effort for node-to-node data exchange since no direct network connection had to be set up or maintained since the traffic was routed via the publicly available MQTT broker.

Data transfer for monitoring and charge control

The two most relevant data streams of the setup and also the basis of the VEIU Tech system were on the one hand the measurement data from the vehicles and from all the included installed energy/power sensors in the grid system of interest and on the other hand the actual control signals that operated the charge optimisation according to the output of the charge manager.

Since unfortunately only a minor set of measurement devices and/or setups implemented an option for data dispatch with a MQTT client, a data forwarder routine had to be implemented in order to translate the respective data to MQTT data points that were then forwarded to the broker.

On the vehicle side, the board computers sent out the specific vehicle data points directly to the MQTT broker while also subscribing the charge control data points that were fed by the charge manager soft-



ware. The effective charge control came from the control translation to the charge adapters via Bluetooth, but this is described in more detail in a following subsection.

Data aggregation for data analysis and modelling

In order to document all the tests performed in the project and to do the necessary numerical analyses, all the relevant data travelling across the MQTT broker were also being stored in a local database. This project relied on the open-source infrastructure provided by InfluxData, which includes InfluxDB - a time-series database that can be installed and run locally – and Telegraf – a plug-in server agent to collect and write metrics in the database. In the project, Telegraf was used to subscribe to all relevant topics published in the MQTT broker – either coming from the charge manager or the sensors – and parse the payload in a convenient format before being written in the Influx database.

To easily monitor the system during the live operation, multiple graphic dashboards had been set up with Grafana. These dashboard were used on one side with a monitoring function, e.g. to verify the flawless operation of the charge control system or to assess the correctness of the charging response to given charge management logics; on the other side, dashboards were employed in a descriptive manner to identify possible patterns in fleet and grid behaviours and thus build more informed charge management schemes. The primary source for most dashboard visualisations was the Influx database itself; when necessary, direct access to other external data sources (e.g. the Solar Manager API) were implemented.

For more intensive post-processing and data analytics, scientific toolkits of Python 3.8 had been employed. As with the Grafana dashboard, data extracted from either the Influx database or other data sources were combined to perform the necessary analyses.

Third-party data exchange

For the integration of third party services the implementation varied. On the one hand API data integrations that rely on data point requests were implemented with a separate data acquisition routine that translated packets to the defined MQTT data points. On the other hand some systems supported the MQTT protocol such as the energy service pool manager of the project partner Primeo Energie, which supports MQTT client interfaces for their asset integration. In this way special access credentials for specified data points were created and the availability and control signals were directly exchanged via the MQTT broker.

2.2.2. VEIU Tech system implementation

At the core of the project stands VEIU Tech's system for vehicle data based charge control. It enables a holistic approach to the charge optimisation challenge that includes data based decisions and inclusion of different vehicle types. This is crucial for a diverse fleet setting that should be operating without any interruption and maximising the charge flexibility it can provide as energy services.

This is effectively done with an on-board device receiving vehicle data from the car diagnosis bus, which is available for most cars compared to cloud interfaces that only a limited of manufacturers do provide. Then a charge controller rounds up the system by enabling an unrestricted charge current control by interfacing with the Type 2 charge plug standard.

In the diagram and sections below the different parts and their respective deployments are described in order to point out the key features and their respective contributions to the whole charge optimisation system.




Figure 4. Schematic overview of the main components of the charge control system by VEIU Tech and their interactions.

Vehicle on-board computer

The vehicle on-board computer is a device that acts both as data retrieving platform, as well as charge control delivering platform. Based on the on-board diagnostics (OBD) bus system it leverages the standardised interface that is found over a broad range of vehicles and brands as it is also bound to regulatory settings. Not only for powering the device but also the access to the vehicles data bus is provided via the plug interface. Although the effective communication and data point settings differs between brands and sometimes models and manufacturers do not support those interfaces as an official system, significant work is done on it in the open source community and solutions for working with those interfaces are often available or provided through 3rd party systems.

The interface and the varying documentation bring along the fact that not all metrics are available from all cars. The availability of data can vary by a large amount making some cars only having minimal provision where other cars deliver close to any detail of information. However it can be stated that the main vehicle data is basically always available. That is, state-of-charge, mileage, high voltage (HV) battery current and/or plug/charge status can always be found or derived from the available data. In this EV charge management setting it would especially be nice to retrieve some dedicated electrical power data from a car-meter such as effective power and energy taken from the grid. Some metrics go into that direction but mostly aim for measurements at the high-voltage battery and therefore omit the charger losses.

A challenge that arose with the OBD interface data fetching was some unexpected 12V battery drainage. Not all car models would show this behaviour, or at least not to the same extent. This is especially true for the OBD systems, where a gateway is installed and the data have to be requested. It showed to be true that with actively requesting data on the bus, the 12V battery would be drained. The assumption is clearly that some vehicle components are kept active in this way, such that they all continue to consume current (and power) from the 12V battery that is significant enough to drain the battery.

Also the OBD device is directly powered by the 12V battery and does consume some power, but in a range that does not impact a car's 12V battery that is regularly recharged.

Even though this behaviour is manageable, it seems that generally for short trips the 12V battery in EVs is not dimensioned in a large enough way, such that regular power draw with limited car usage can put it out of order. This is effectively hard to understand and accept since it is a setting where a electric



vehicle that is plugged to a charging station with a large amount of energy in its high-voltage battery can go out of service due to an empty 12V battery system. This is a behaviour that can be omitted and will probably be solved in one or the other way in future car models. (Some models do not show this problem.)

In order to tackle this behaviour the OBD device can just shut off when a low voltage is measured on the 12V battery voltage and therefore minimally stress the battery itself. It can then wake up when the car is on, which is the case when the 12V battery is charged and the voltage surpasses a certain charging threshold. In this way it can be made sure that no car goes out of order, however it does limit the time that the OBD device is available and connected to the internet.

Charge controller

The charge controller makes up the critical device for the actual control of the charge process. It is a device by VEIU Tech, that intercepts and adapts the communication line between the charging station and the EVSE in order to control the available current/power that is communicated to the car. The car then charges according to the power made available over this communication. By adapting the direct communication, a granular control of the current and therefore the charging power is possible compared to other interfaces (EVSE/EV API) that often only consider charge start and stop for charge control.

Hardware




Figure 5. Charge controller hardware.

The charge controller is effectively a piece of electronics hardware that is on the one side connected to either side of the control pilot (CP) line of the communication wire and on the other hand to power supply. In special versions it can also be battery powered for the times where the charge line is not powered (no charging).

The control is effectuated by a microcontroller that receives control signals via Bluetooth and then puts them into effect via respective electrical signal generator circuitry. Further some basic UI lights are available to show the user in which mode the charge controller is operating.

Interaction with on-board computer

The charge control signals that the charge controller uses to effectuate an actual adaptation of the charge current communicated to the car is forwarded by the on-board computers described above. The on-board device is the device that is connected to the internet / cloud and that fetches current car data in order to run effectuate a data based charge control. The on-board computer connects to the charge controller via Bluetooth in the moment that the charging cable is plugged into the car. In that moment



the charge controller starts to advertise its Bluetooth connection and the on-board computer respectively understand this event and scans for a respective charge controller device in the Bluetooth space. This connection is not only helpful for having the charge control on the already connected on-board computer, but also enables to make a vehicle based charge control. That is through this connection it is always clear what car is connected to which charge controller and therefore which charging station. This is especially helpful for setups where multiple cars charge at the same charging station.

Plug adapter option

For configurations where the wallbox is either inaccessible or accessing it would void the warranty, an external plug adapter is required. This adapter is designed to interface directly with the charging cable (male end of the cable with the adapter's female interface) and the EVSE (EVSE's female socket with the adapter's male interface).

The adapter is designed to fully accommodate the Charge Controller, allowing its communication with the CP line of the charging cable. Additionally, it enables an unobstructed pass-through for the five high-amperage conductors of a 3-phase charging cable, including neutral and ground, using standard 32 A-rated cables. To ensure proper grounding of the Charge Controller, it is also connected to the ground conductor on the EVSE side.

The adapter's male and female connectors comply with IEC 62196-2 Type 2 standards, which include "locking holes" for secure attachment to both the wallbox and the charging cable. The adapter's plug has been modified to ensure compatibility with the specific EVSE models encountered in this project.

The Charge Controller requires a power supply to operate and two options have been considered:

- In principle the live cables going through the adapter could be employed, but these are powered only when the EV is actually charging. Therefore a battery is needed to buffer energy in order to supply the controller device during charge inactivity times. Long times without charging are therefore stressing the needed battery capacity and raising the need for large exemplars. However, given the project's setup requirements, this level of modification was deemed unfitting.
- An external power supply can be fed to the adapter. This was the approach finally pursued, with a second low-amp cable exiting the EVSE from a secondary "door" and reaching the adapter. To ensure that the adapter can be removed or attached whenever is convenient, this secondary power supply cable interfaces with the adapter via an IEC 320 C7/C8 230V 2.5A connector.

Two LED guides have been integrated at the top of the adapter, positioned above the LEDs on the printed circuit board (PCB). These LEDs communicate the following statuses: (1) whether the Charge Controller is active – indicating that it is operational and ready to receive and execute charge control commands – and (2) whether the Charge Controller is connected to the plugged-in EV via Bluetooth.




Figure 6. Early prototype of the charge control adapter developed for the project. This prototype includes the Charge Controller board, a low-amperage communication wire connected to the Control Pilot (CP) line of the charging cables, and the five primary high-amperage conductors for the 3-phase power connection.

The adapter's plastic housing, along with the male and female connectors, have been produced using additive manufacturing processes, while the metal pins are standard components that have been acquired on the market. The housing design incorporates four internal supports to securely position and stabilize the Charge Controller.



Figure 7. Deployed plug adapter with charge controller at an EVSE in Münchenstein.

Charge manager cloud service

The VEIU Tech charge manager is the cloud service that runs the charge optimisation throughout a full fleet of cars. This centralized manager assembles all relevant data points for a certain charge optimisation use case and computes respective charge control schedules for the single cars. It is especially important to have a centralized unit for use cases where the actual action of single EVs have an impact



on the control strategy of other cars, such as for the peak shaving use case. Further it facilitates the modular scaling of the application with cars and/or fleets and an efficient fleet and use case configurability.

The charge manager is run as a cloud service that connects to the MQTT data broker and exchanges data with the respective devices via the broker platform.

Deployment timeline

Since both project partners would house different infrastructure setups with both different cars (manufacturer and models) and different charging station installations the deployments of the VEIU Tech system would need to account for the differences found. On the one hand the installation of the board computers would follow the same procedure with the difference that for the operation the ones used in the Primeo Energie fleet would need to consider the 12V limitations and configure a minimal operating condition. The Quickpac fleet would not face this issue because the cars would be operational many more hours and more regularly.

The biggest difference would however be the installation of the charge controllers since one charging station type would provide enough space to house the charge controller inside the charging station and the other charging station type not. In the second case the adapter version of the charge controller needed to be used which had an impact on the production and installation effort and time.

The following deployment timelines were realized during the course of the first half of the project:

Deployments	Quickpac: Hägendorf	Primeo Energie: Münchenstein
<i>Board computer</i>	January 2024	February 2024 – August 2024
<i>Charge controller</i>	February 2024	June 2024 – July 2024
<i>Operation tests</i>	March 2024	June 2024 – July 2024
<i>Live use cases start</i>	March-April 2024	July 2024

The delay for the Münchenstein site was twofold. On the one side the challenge with the 12V battery status and triggered alarms from data fetching when the car was locked delayed the proper rollout of board computers to the last vehicles of the fleet. On the other hand the adaptations and mechanical design iterations on the charge controller adapter also delayed the proper rollout of the complete VT charge management system for the Primeo Energie fleet.

Further some delay is also due to work focus on the Hägendorf site for when we realized that the time at that site was limited due to the organisational transition of the Quickpac company.

2.2.3. Grid load monitoring equipment

The following diagram provides an overview of the approaches adopted at different grid levels to monitor the load of the grid and other relevant power metrics. The main goal at lower grid levels (level 7) was to directly measure the impact of the charge management process and thus quantify its grid relief benefits. At higher grid levels (level 5) the impact of the charge control would not be visible with the available fleet sizes, but the collected data were still valuable for the purpose of the project. Beyond in numerical simulations, these data have been employed in a few live tests to check whether the charge management could in principle react to those specific grid conditions and provide some relief, ignoring the negligible magnitude of the actual impact.




Figure 8. Arrangement of the different load metering devices at various grid levels at the two project sites.

In the following we detail the devices and setups employed to monitor the grid at every level.

Charging station monitoring

In Hägendorf a dedicated EVSE monitoring was omitted since it was possible to install an upstream parking measurement device that would exclusively monitor the set of charging stations under control. In this way an exact measurement and EVSE mapping is possible. The installation of an additional meter inside each EVSE was thus considered disproportionate.

In Münchenstein however the parking metering described below would include different loads, such as other not-controlled EVSEs, so that a perfect disaggregation would not be straightforward. Therefore, a Solar Manager device was installed for specific EVSE monitoring. The Solar Manager can read out, record, and visualise the active power of individual charging stations every 10 seconds via the "Zaptec



portal", the installation specific manufacturer's management interface. As Zaptec does not offer an open interface to read this data, the Solar Manager is generally used as a gateway. In this way, there was no need to develop a cloud integration for the Zaptec portal which would have lead to a comparably large development effort.

With this installation, the behaviour of the charging infrastructure could be visualised via the Solar Manager portal with which it is possible to see how the pilot project is changing the charging behaviour and the respective power consumption. In order to eventually integrate with all other power measurement data, the Solar Manager API was used to integrate with the QuickOpt data analyser.

Parking load and power quality monitoring

On the intermediate level, that is on power lines that would gather a set of charging stations (first level accumulation of power flows), grid and power quality meters were installed. The goal was to install such devices on either locations in order to not only measure the basic power and energy metrics but also more complex metrics that give away the power quality of the grid at that point. This would help to also monitor possible other effects on the power infrastructure that would arise from the charging processes and to investigate whether intelligent control would be able to mitigate those effects.

The rolled-out devices were Janitza UMG-605 devices that are often used for grid quality measurements and monitoring. It brings along a lot of measurement metrics and gave the full flexibility to gain knowledge of the underlying grid infrastructure with good quality.

Next to the standard metrics of power (P), current (I), voltage (V) of every phase also advanced metrics were monitored such as the frequency (f), the cosPhi factors, the THD (total harmonic distortion) factors as well as the harmonics (up to 25th) of every phase. Also due to missing identifiable behaviours linked to the charge control these metrics were not all evaluated in detail.



Figure 9. Janitza UMG-605 used for parking load monitoring.



Grid access points monitoring

Münchenstein site: Smart meter readout via EnergyX devices

In order to visualise the influence of intelligent charge control on the buildings load in the Münchenstein site, the measurement data from the existing NES smart electricity meters was analysed and aggregated. The NES electricity meters installed for billing made additional measurements at the various consumers superfluous. Separate measurements would only have been possible at the existing electricity distribution points and at great expenses. Accordingly, the use of existing electricity meters was promoted and successfully implemented with EnergyX, a connected meter readout system. To make the devices available remotely, a virtual network was created using ZeroTier, which simplified the integration of the decentralised devices via cellular network as if the devices were part of a local network. The active power data was then queried remotely every 30-60 seconds via the local / on-device EnergyX API.

For data management, it was important to understand how the meter data was measured. For CT measurements, it was necessary to check when the measured values were correctly scaled (transformer factor). The pilot project showed that the values at the NES-MEP interface are normally not scaled. Accordingly, the conversion factor had to be assigned twice in the EnergyX gateway, for the web interface and the for the local API.

These power value read-outs helped to understand the grid situation on the site in real time. Recording the individual measuring points provided a detailed view which allowed charge control to be carried out in line with the behaviour of the individual measuring points. One example of this is PV charging optimisation, which used the real-time data from the PV system's electricity meter. The on-site DC fast chargers were monitored too and could also be considered in the live optimisation.



Figure 10. Smart meter monitoring in Münchenstein with EnergyX device in the background.

Hägendorf site: Smart meter readout via LED pulse monitoring

The electricity meter reading at the Quickpac site in Hägendorf showed to be limited from what the installed meter would provide. An initial test of data readout via the optical communication showed to



only deliver the main metrics also provided by the display interface, that is, total energy for different tariff times and the peak power value of the last couple of months. Even when regularly polling the energy value in order to deduce the current power (via the derivative), this would not lead to a satisfactory result since the update rate would be too low – especially during low power usage times – and therefore the resolution would not suffice.

The workaround found in order to still extract a reasonable live power value was to use the flashing power indicator light of the meter. This light has a much higher update rate compared to the meter readouts, since it triggers at every fraction of a kWh. This blink period is then a direct indicator of how long it took to consume a certain energy amount (in this case 4 Wh) and can be transformed in a power value with update rates in the orders of seconds.

The tracking of the flashing light pulse was implemented with a small self-designed electronics part with a photo diode connected to an IoT microcontroller board. That board would then forward the measured pulse lengths via MQTT over a Wi-Fi network.

This simple method showed to work reliably and delivered satisfactory power resolution for the site monitoring and application. It is also quite generically applicable to any meter type since most of the standard meters have a light pulse indicator.




Figure 11. Power pulse monitoring of the building meter in Hägendorf.

Level 6 grid monitoring

Münchenstein site: ad-hoc measurement unit

In addition to the aggregated energy meter measuring points, a mobile monitoring device was installed at transformer station 2 in Münchenstein. This setup accounted for smaller, unmeasured consumers and line losses to individual consumers. However, it did not measure the transformer's own power loss. The comparison of data from the transformer measurements and the aggregated meter readings showed a high level of alignment, confirming the accuracy of the measurement setup.

In this setup, current on the phase conductors was measured using plug-in Rogowski coils, which can be easily attached around existing cables for flexible installation. These Rogowski coils provided adequate accuracy to capture grid condition data, and their ability to measure cables without interruption has been a major advantage. At transformer station 2, the connection to the low-voltage main distribution board (LVMD) consisted of four separate conductors per phase. One conductor per phase was



monitored with a Rogowski coil, and the measurement was then scaled accordingly by the measuring device.

The measuring device was connected via a LAN interface to a cellular network gateway, allowing real-time integration of the measurement data into the project's MQTT broker. Voltage, current, active power, and reactive power were analysed for each phase line, enabling charging optimisation based on the current status of the transformer station 2 and aligned with grid requirements.



Figure 12. Mobile grid measurement unit installed at the level 6 transformer in Münchenstein.

Hägendorf site: no monitoring available

A second level 6 transformer line measurement was planned for the Hägendorf site with the local DSO. This would have been analogous in the sense that the level 5 connection was part of the project partner Primeo Energie's grid. The intention was to use the same measurement equipment as used in the Münchenstein site and have an analogous setup that would entail the measurement device and the measurement coils in the local DSO's transformer station.

Unfortunately, the local DSO was not willing to have an external device installed for this pilot project and was not interested to participate in the project. The communicated reasons included the lack of resources to participate in a research project and the unclear benefits for their local grid. Therefore, for the Hägendorf site only level 4/5 grid data could be considered, which most likely show similar usage patterns since the transformers attached to the level 5 grid would show comparable grid setups (village with small industries).

Level 4 grid monitoring

Load data for the highest grid levels was directly provided by Primeo Energie as part of the legally mandated grid compliance monitoring. More specifically, this project considered the loads of the outgoing connections from level 4 transformers – thus corresponding to the level 5 network – linked to the two sites of Hägendorf and Münchenstein. Voltage, current, and power were measured in real time, with voltage and current measurements recorded every minute and power measurements stored every 15



minutes. Although the system captures data continuously, it archives measurements at these intervals for efficient storage and communication.

Given the limited impact anticipated from the charge control on this network level and in order to minimise the necessary workload, no live transmission to the MQTT broker was implemented. Instead, we opted for a manual monthly exports of the grid measurements. This method was deemed sufficient to perform numerical simulations or to undertake “fictitious” live tests, where actual charge optimisations were carried out based on past, retrospective grid measurements.

2.2.4. Fleet booking systems

In order to run the charge optimisation that was envisioned in this project and quantify the flexibility available to run respective services, it was crucial to know the usage times of the respective vehicles. Especially the planned departure time was important because it frames the available time window that was available for the charge management and its available flexibility.

While with private customers this is a major challenge because of irregular usage patterns and the additional effort to adjust those times, the operation of business fleet cars follows more planned schedule. The usage is normally bound to working hours, project location and/or shifts that are registered in some work logging tool. This is a great enabler for the VEIU Tech data based approach that tries to minimise the estimation and work with defined and current data.




Figure 13. Example of trip bookings on two days in Münchenstein.

Hägendorf site

The logistic fleet partner Quickpac showed to have a very defined shift schedule. The varying part would be the drive time, but at plugging time the duration of the available charge time would be defined, up to the start of the shift, normally at 5am. Rarely, single cars would not be scheduled for the shift of the day, which would lead to a larger charge availability time. However the car assignments for the shift would normally only happen in the morning (dependent on the daily packet load) and therefore would not bring any significant benefit to the charge flexibility. The used tool for the car scheduling was an internal tool such that an integration would not have been straightforward. Considering this and the minimal benefit, it was chosen to just use a predefined (configurable) ready time as a vehicle operation constraint.

Münchenstein site

On the other Primeo Energie fleet two main usages were present. One was that of personally assigned cars that were used regularly by employees of a specific team and the other was that of a bookable pool



fleet for the larger employee set. Pool cars were booked in advance by a booking tool and could be used accordingly. Those two sets of cars therefore showed a very different usage pattern and also the usage data was accordingly differently available. For charge control, the personally assigned cars were available during working hours, when the respective employees were in the office. Their availability could therefore mostly be counted during the working day.

For the pool cars the booking system used was “autoSense” that comes with a fleet API over which the operators can access the booking data. This booking data was precisely defining the departure times so that it did very well anticipate the charge availability times. So for the project’s charge optimisation the pool booking data was extracted and uploaded to the MQTT broker where it was available for display, monitoring and especially for the charge manager. With that information it could be made sure that all cars were always fully charged for the booking and still kept the full charge flexibility that was available. The only missing information in that specific booking system was the trip’s distance. This would have added some flexibility since it would have lowered the minimum charge level at booking start and enlarged the flexibility even more. Nevertheless it was clear that this was less common and an additional step for the user that was not necessarily just known, such that just a general working value (e.g. 95% SoC) seemed like a satisfactory solution.

2.2.5. Ancillary services link

In order to run real-time charge control along ancillary services requests the whole VEIU Tech charge management platform had to be linked to an ancillary services management platform. For this purpose a link to the *Equalio* platform was created that would connect the charge management system to the management platform of Primeo Energie’s flexibility pooler. The *Equalio* platform is a flexibility management platform that connects many smaller flexibility assets and pools them to act as a ‘virtual power plant’ (VPP) which has the size to act on the ancillary services market of the TSO Swissgrid.

In order to set up a live system for testing and not having to go through the full (pre-)qualification process by Swissgrid, a prototype-asset could be set up that was only rated with a very small power of 10kW which basically did not impact Swissgrid’s ancillary controls. This way the activation and requests could be tested and analysed while no compliant reaction times and reporting needed to be set up.

The live connection to the Equalio platform encompassed 4 data points. Two of them were set by Equalio and the other two by the flexibility asset. On the asset side an availability and an activation power value was transferred and on the management platform side an activation and requested power value were exchanged. The implementation with a MQTT client that connects to the project’s MQTT broker was chosen since the MQTT publish-subscribe concept matches very well the event based nature of the ancillary services requests.

2.3 Live use cases

One of the main goals of this pilot project was to understand in which cases a data-aware charge optimisation of fleets would help a specific set of the grid to reduce utilisation in critical moments. The specific settings would vary a lot since different levels of the electricity grid show different topologies, capacities, usage patterns as well as balancing capabilities. In order to assort these settings, use cases were defined for which a specific optimisation goal would be targeted. In order to get a quantifiable comparison this would then be compared to the uncontrolled base-scenario.

Even though this uncontrolled base-scenario sounds like a good reference or common ground, it showed to be more volatile due to small organisational changes and special cases. Nevertheless, this still accounts for the “normal” case and would still make a point in the analysis.

With the project focus on grids the three main use cases were defined including:



- Local (building) peak shaving
- DSO grid peak shaving
- Ancillary services

In addition, further use cases that also support the general grid operations have been considered and evaluated:

- Local PV consumption
- Voltage stabilisation
- Dynamic tariff optimisation

In this section we introduce the use cases with their target goal and technical setups that were used. Further the different used algorithms are listed including their data needs and potential outputs.

2.3.1. Local (building) peak shaving

Local peak shaving can be seen as the standard use case. This is especially true since the regulations of the time of the project incentivised low consumption peaks over a peak power tariff for customers with large consumptions (> 100 MWh) and more simple due to limited grid connection capacities that enforced a limited consumption peak for building loads.

The usual solution to shave the peaks within charging infrastructures is to either limit the controlled total power within the charging infrastructure or to tap the buildings meter and control the charging infrastructure to reduce charging when the building power exceeds a certain limit. Both solution have the limitation that without considering the SoC of the cars, no prioritisation can be done. That said, the limit of the infrastructure has to be quite big in order to make sure every car can charge in the amount of time needed. In other words, the infrastructure itself needs a significant slack compared to the actual charging needs of the connected vehicles since theoretically it needs to supply the theoretical minimum charge for the highest need to every car.

Taking these considerations into account, we could state that charging solutions at the time of the project typically aimed for a static power limit. This means they targeted a maximum power capacity for an installation or building to prevent exceeding its limits. Such an approach helps to avoid the need for infrastructure upgrades or allows the use of installations with limited capacity. This challenge is analogous to issues faced at higher levels of the electrical grid but occurs on a local building scale.

In contrast, the VEIU Tech charge management solution offers a dynamic approach. It allows for the prioritisation of vehicles to ensure they are fully charged by their required readiness times while still limiting the total power consumption as much as possible. This is achievable because, with knowledge of each vehicle's SoC and its ready-time, the system can prioritise and schedule charging. It ensures that vehicles begin charging with the necessary power at the appropriate times.

This utilisation of detailed vehicle data leads to dynamic optimisation rather than a static approach. Instead of merely ensuring that power consumption does not exceed a fixed limit, it actively minimises total power peaks. This strategy can potentially reduce power tariff costs under current pricing models and even more so if tariffs consider factors like the average daily peak rather than just the maximum monthly demand. By dynamically minimising the maximum power usage each day, the system avoids uncontrolled scenarios that might stay under the fixed limit but still result in higher costs due to peak demands.

In addition to the standard reference load capping the following algorithm setups were considered for this use case:

Input Data	Building consumption power data (NO vehicle data)
------------	---



Algorithm 0	Load capping with fixed threshold
Output / Evaluation Data	Building consumption power data
Further optimisations	Vehicle data (SoC, ready-times) based control, see below

Input Data	Vehicle data (SoC, ready-times), building consumption power data
Algorithm 1	Static-timed charge delay (fixed charging times)
Algorithm 2.1	Dynamic load drop (based on high peak appearance in baseload)
Algorithm 2.2	Dynamic peak load reduction with fixed threshold (dynamic charging below fixed threshold)
Algorithm 2.3	Dynamic peak load reduction with adaptive threshold (dynamic charging below adapting threshold)
Output / Evaluation Data	Building consumption power data
Further optimisations	More advanced building load/threshold prediction

2.3.2. DSO grid peak shaving

What the local peak shaving is for building owners / operators, the DSO grid peak shaving is for DSO and utilities. The basic approach is to minimise the peak load on different grid components, mainly transformers. This is the case since the transformers in the grid mainly determine the line capacities that emerge from the transformers' outgoing lines. This is the case for the "local" level 7 grid (with level 6 transformers) and goes up the regional grid to level 3. The impact of a single EV fleet is however not always observable, especially in the level 4 transformers. Therefore, the optimisation is more about spatio-temporal avoidance of peaks in the grid, that is omit charging times during high grid utilisation times. The evaluation is hence more about the ability to avoid those times and understand whether simple regular usage patterns bring the full clarity or whether a live feedback is relevant for respective peak shaving.

Since it is also peak shaving the algorithms follow the previous case.

Input Data	Vehicle data (SoC, ready-times), grid meter data, proxy data,
Algorithm 1	Static-timed charge delay (fixed charging times)
Algorithm 2.1	Dynamic load drop (based on high load appearance above fixed threshold)
Algorithm 2.2	Dynamic load reduction with fixed threshold (with communicated reduction percentage)
Output / Evaluation Data	Temporal charge pattern, grid meter data, charge control action log
Further optimisations	More advanced grid congestion predictions

2.3.3. Ancillary services



Next to grid usage optimisation the balancing of the energy is a main stability optimisation needed for a reliable and functioning grid. This is handled by the different ancillary service products managed by the TSO Swissgrid. For this pilot project a participation in the ancillary services was conducted through the asset pooling manager Equalio that makes part of the project partner Primeo Energie.

Considering the constraints given for the different ancillary services products, the secondary control (automatic frequency restoration reserves or aFRR) with a focus on negative control was chosen as a primary use case. The challenge in this use case was to consider the availability and manage the fleet to provide a charge control throughout the inquired times.

The actual control comes from a requested power consumption that needs to be triggered for a limited amount of time. The maximum requested power is defined beforehand and needs to be available a certain amount of time (1h-1week), depending on the ancillary service product.

Input Data	Vehicle data (SoC, ready-times), ancillary service request signals
Algorithm 1	Negative aFRR/mFRR; charge on hold until requested, charge just before ready-time
Algorithm 2	Positive aFRR/mFRR; control charge power and reduce charge when requested (load drop)
Output / Evaluation Data	SRL request signals; charging infrastructure meters
Further optimisations	optimise ready-time's charge levels (SoC) for better availability

2.3.4. Local PV consumption

Instead of focusing on downstream grid consumption capacity, this use case examines upstream capacity congestion, particularly in the context of PV integration. The objective was to address a future scenario with high PV penetration in the grid and avoid potential congestion issues related to local PV infeed.

In the specific project's setup, the optimisation was not aimed at maximising the effective own-consumption, as the site load demands were substantially higher than what the local PV plant would generate. Instead, the focus was on a local or regional optimisation, with the goal of charging an EV fleet during peak PV generation periods and, where feasible, utilizing PV-generated power from the local grid. This was also done by just simulating an own-consumption scenario in which only the fleet loads were considered and the building loads were omitted.

Input Data	Vehicle data (SoC, ready-times), PV meter data, local load data
Algorithm 1	Static time-based charge: charge during fixed ours with expected high PV generation
Algorithm 2	Dynamic PV threshold charge: charge when PV production exceeds fixed threshold
Algorithm 3	Dynamic own-consumption charge: charge when local PV production exceeds local consumption (i.e. maximise own-consumption and reduce grid feed-in)
Output / Evaluation Data	Temporal charge pattern, PV meter data, charge control action log
Further optimisations	More advanced PV predictions for the day



2.3.5. Voltage stabilisation

Another indicator for the prevailing grid state and power quality is the measured voltage at the grid endpoint or the meter. Low voltages arise when the actual grid line coming from the transformer is carrying a significant current, which leads to a voltage drop over that line. The topology also adds up on this, making the effect stronger on grid points with longer connection lines. The other effect on the voltage is that of a voltage drop coming from the transformer that supplies a lower voltage when under higher load.

Since the measured voltage is not solely dependent on the endpoint's usage, a charge control can be visible on the voltage measurement, but normally has not a big impact on the voltage. It seems like the fluctuations are normally significant, such that small changes are difficult to directly assign to a charge control action.

In general this method is not generally approved since the voltage loss in grid system has a highly local dynamic that depends on the topology, the user profiles and the system / hardware that is deployed in the grid. Further control activities on the transformer infrastructure can also interfere since methods, such as compounding, which aims for voltage regulation at the transformer by adapting the set point of the transformer for different load situations. Obviously this control breaks the stated dependence from grid load to measured voltage and introduces independent voltage changes that mix up the effective control.

Input Data	Vehicle data (SoC, ready-times), grid voltage data
Algorithm 1	Charging load drop for low voltage periods
Output / Evaluation Data	Grid voltage data, charge control log
Further optimisations	Prediction of expected nominal / average voltage

2.3.6. Dynamic tariff optimisation

In order to make the needed flexibility more accessible for end customers and the incentives of flexibility better understandable, different utilities have started to work with dynamic tariffs. These tariffs give a direct incentive to move loads from one time to the other whilst getting remunerated for this shift via a reduced price. With the novelty of the tariffs being dynamic, so varying from day to day and also with a higher resolution, energy providers or DSOs are able to incorporate short term effects like (un-)availability of solar production or holidays in order to adapt the incentives to other production and consumption patterns. Depending on the tariff itself – being a energy tariff or a grid usage tariff – the tariff is based on the energy supply and demand or on the congestion in the DSO grid.

For charge control this is a convenient method for flexibility provision, since it gives a clear incentive structure beforehand and the charge schedule can be calculated based on the expected driving needs. The challenge of changing tariffs can arise with very dynamic tariffs (e.g. updating every hour) however this is not seen today and will not reflect the added value that can be provided through such tariffs.

While dynamic tariffs can give incentives for grid-friendly flexibility provision, the challenge is for the tariff makers to also consider local effects in the grid. A varying local tariff depending on the location in the grid could help to better reflect real effects in the grid. However this added complexity is hard to handle for DSOs and breaks the simplicity of dynamic tariffs in general.

Input Data	Vehicle data (SoC, ready-times), day-ahead tariff data
Algorithm	Charge in low-price times



Output / Evaluation Data	Tariff data, charge control action log
Further optimisations	Forecasting price structure before publishing and/or forecasting trips.

2.3.7. Peak shaving classification

In the context of grid relief and efficient capacity usage the general term is peak shaving. Peaks stress the infrastructure and bring need to install more capacity which makes the grid less (cost) efficient. However the dynamics around peaks is quite diverse and depending on the appearances a different approach is convenient to tackle the flattening task.

Generally there are two main dimensions that classify peak shaving setups. On the one hand the appearance pattern of the peaks specify the problem and on the other hand the ratio of the controllable power compared to the peak portion of the load power determines the way that charge control can ideally be employed. This leads to a 2x2 matrix of peak shaving scenarios with different characteristics and control approaches.

Peak Appearance Type	Minor EV / Controllable Power	Dominant EV / Controllable Power
Dynamic (Varying times)	Load dropping (Fixed threshold based)	Load balancing (Live reference based)
Static (Repeating times)	Peak avoidance	Load distribution

- Load dropping:
This case applies when the peaks appear in a repeating manner and are hard to predict. Having the controllable load or in this case the EVs being minor to the peak load, the strategy is to drop all load that is controllable based on a direct control signal. This can be directly from a meter data stream or an indirect DSO processed control signal. This DSO signal can be a dynamic tariff incentive or a direct control signal.
- Load balancing:
In this case the peaks still appear in a repeating and unpredictable manner. However here the controllable loads dominate and need to be adjusted to not overshoot a desired value while the base load shows a peak. The controllable loads can still be active, they are just regulated in order to not forward the base peak to the total load. In this case a direct control or a control signal with a power rating is needed for an effective control.
- Peak avoidance:
In the peak avoidance case the peak appears regularly at the same (or similar) times and the peak is relatively high compared to the controlled load. Therefore a static incentive to circumvent those times is enough to reduce the final peak. A time-of-use tariff or lock-times are tools to implement this.
- Load distribution:
When the base peak is recurring in a predictable manner and the controllable load dominate we speak of load distribution. Basically the load can be scheduled based on the availability



and the “expected” base peak power. A time-of-use tariff or a fixed published (maximum) load / capacity schedule can be used in order to implement such a control case.

In the context of this project with EV fleets all scenarios can happen on basically all grid levels. It is different asset classes and deployment setups that create one case or the other. An example for dynamically appearing peaks are fast charging stations on the lower grid level or industrial tools or machines on the local scale. More static peaks come from standard pattern that are either scheduled or common like (boiler) heating peaks or evening cooking peaks. The size of the controllable (and flexible) load is basically the relative size of the EV fleet(s) compared to the investigated power load domain.

2.4 Simulation model

2.4.1. Scope and goal

Next to the live use cases, another important part of the project is the development of a digital twin that can simulate the operation of the charge control system in setups other than the ones accessible during the project. This digital twin is disconnected from the physical system, meaning that it's not meant to react to live measured data or to feed any instruction back. Nevertheless, the data collected along the project serve as major inputs for the model simulations.

In the simulations, fictitious EV fleets are modelled with their essential features (battery capacity, charging rate, trip length, etc.) and operated for a full week. The charging processes are then optimised in order to maximise various grid-centred objective functions while fulfilling all important operation constraints. A key aspect of these numerical optimisations is the perfect foresight that allows the charging processes to be adjusted based on future developments. In the work we attempted to contain this distortion by accounting only for constraints and data inputs actually available also in the physical system, e.g. time of the next booked trip. In fact, one of the goals of the digital twin is the quantification of the added benefit of integrated charge control systems that have direct access to SoC and booking data.

The simulation work is split in the following steps:

- Setup of the optimisation problem, presented in detail in the next subsection;
- Fleets characterisation, where the necessary information for the simulation is derived from the physical project and adapted for the purpose of the model. The goal is to extract the typical distribution of key quantities (travelling times, discharge rates) in order to subsequently generate fictitious fleets of any size with mobility patterns that emulate the real counterparts. Section 2.4.3 presents the detail characterisation of the main fleets available in the project.
- Exploratory simulations: here we start from setups close to the ones realized in the live system and gradually move away to test situations with different combinations of 1) fleet characteristics, 2) grid conditions, and 3) flexibility use cases. The results from these simulations are presented in Section 3.4.

2.4.2. Optimisation problem

Most of the simulations tested with the digital twin are optimisation problems, meaning that the fleet charging processes are optimised to maximise or minimise a single objective function. In a few instances this is not the case: for example, as reference the standard “plug and charge” logic is typically employed, which does not constitute an optimisation problem (although it could be formatted as such).

The simulated time window is always 1 week, with a time resolution of 15 minutes, for a total T of 672 time steps. In the following mathematical formulation t denotes the t -th time step, so that:

$$t \in \{1, 2, \dots, T\}$$



The number N of simulated EVs varies depending on the specific simulated setup; each EV is denoted with i , therefore:

$$i \in \{1, 2, \dots, N\}$$

Each EV is characterised by:

- A battery capacity E_i in kWh
- A minimum charging power $P_{\min i}$ in kW
- A maximum charging power $P_{\max i}$ in kW, which accounts also for the capabilities of the associated EVSE.

During the week each EV drives a set of J trips, whose number is not necessarily the same for the entire fleet. Each trip is indicated with j , hence:

$$j \in \{1, 2, \dots, J_i\}$$

Setting up the optimisation problem involves defining the following key components:

- **Objective function:** It is the primary expression that the model seeks to either maximise or minimise. It represents the goal of the optimisation, such as minimising the peak load, minimising the power expenditure with dynamic tariffs, or maximising the time available for ancillary services. The exact expression depends on the specific flexibility use case and it will be provided for each simulated problem.
- **Decision variables:** These are the variables the optimisation model adjusts to achieve the objective. The main decision variables in this charging optimisation problems are:
 - $x_{i,t}$: the charging power at the battery level in kW for the i -th EV at the t -th time step, which is a continuous variable;
 - $y_{i,t} \in \{0, 1\}$: binary variable that indicates whether the i -th car is charging at time step t ;
 - $e_{i,t}$: continuous variable with the energy content in kWh of the i -th EV battery at time step t ;
 - $totpower_t$: the sum of all charging powers at time t , typically used when setting up a peak load minimisation problem;
- **Parameters:** they are fixed inputs that provide context and influence the model but are not optimised themselves. These mostly define the trip schedule of the fleet:
 - $\Delta e_{i,j}$: the energy consumed at the battery level (kWh) by the i -th EV during its j -th trip;
 - $tdep_{i,j}$: the time step when the j -th trip of the i -th car begins;
 - $tarr_{i,j}$: the time step when the j -th trip of the i -th car ends;
 - $p_{i,t} \in \{0, 1\}$: binary parameter constructed from $tdep_{i,j}$ and $tarr_{i,j}$ which indicates whether the i -th EV is plugged in at time step t ;
 - $E_{feas i,j}$: derived continuous parameter which represents the maximum feasible energy content of the battery after each charging session, assuming the EV would always charge as much as possible. It may differ from E_i in case of short stops. This parameters is useful for later defining a SoC threshold that an EV should reach at the end of each charging session (see the constraint "Minimum desired SoC before new trip").
- **Constraints:** these are conditions that must be satisfied for the solution to be feasible and define the limitations on the decision variables. Our charge optimisation problems is defined by:



- Energy balance of each EV battery i (since we are not interested in the exact discharging rate of batteries during trips, we lump together the whole discharge in the last time step of the journey):
 - While discharging: $e_{i,t+1} = e_{i,t} - \Delta e_{i,j}$ for each $t = tarr_{i,j}$
 - While plugged: $e_{i,t+1} = e_{i,t} + x_{i,t} \cdot 0.25$ if $p_{i,t} = 1$
 - In the remaining trip time steps: $e_{i,t+1} = e_{i,t}$ if $p_{i,t} = 0$ and $t \neq tarr_{i,j}$
 - Physical battery limits: $0 \leq e_{i,t} \leq E_i$ where E_i is the total energy capacity
 - Minimum desired SoC before new trip: $e_{i,t} \geq \min(E_{feas\ i,j}, E_i \cdot f)$ for each $t = tdep_{i,j}$ where f is a strictness input parameter ranging from 1 (max strictness, the final SoC should be the maximum possible) to 0 (min strictness, no charging is mandated). In traditional optimisations with perfect foresight f is set to 0, allowing EVs to profoundly adjust their SoC as long as it does not reach 0 (or any other agreed limit) during any trip. A higher f approaching 1 prevents this “backward” transfer of trip information from the future, forcing EVs to be conservative and charge more in the present as deemed necessary. For all presented simulation results f has been set to 0.9 or 0.95.
- Charging power constraints:
 - Charge only when plugged: $y_{i,t} \leq p_{i,t}$
 - Limit to maximum power when charging: $x_{i,t} \leq y_{i,t} \cdot P_{\max\ i}$
 - Limit to minimum power when charging: $x_{i,t} \geq y_{i,t} \cdot P_{\min\ i}$

With the exception of the binary variables, the remaining quantities, constraints, and objective functions are formulated linearly, allowing the optimisation problem to fall within the scope of Mixed-Integer Linear Programming (MILP).

MILP models are computationally efficient relative to higher-order convex problems because they leverage well-established linear optimisation techniques, while taking into accounting a few non-linear but very realistic discrete decisions. Many solvers have been highly optimised for MILP, allowing for faster and more reliable solutions on larger problems compared to higher-order convex formulations. MILP models are also easier to interpret and analyse.

The MILP models have been formulated in python and solved with the GNU Linear Programming Kit (GLPK).

2.4.3. Fleets characterisation

The objective of the following fleet characterisation is the derivation of the input parameters for the optimisation problems, namely the plug-in times $tdep_{i,j}$ and $tarr_{i,j}$ and the battery discharge values during trips $\Delta e_{i,j}$.

This project interacted with two EV fleets, one located in Münchenstein and managed by Primeo Energie and a second one located in Hägendorf and managed by Quickpac. However, when it comes to their characterisation, these two fleets were not homogeneous. The Primeo Energie fleet has two use cases: a pool fleet used by office employees for same-day project trips and grid interventions and a second smaller reserved fleet used by a few employees for work, commuting, and private purposes. The Quickpac fleet is instead composed by two EV models that are operated differently because of the difference in all-electric range.



The following analysis provides a characterisation of the two fleets most used in the live use cases: the Münchenstein-based pool fleet of Primeo Energie (“PE fleet”) and the eVito delivery fleet of Quickpac in Hägendorf (“QP fleet”).⁶

Primeo Energie – pool fleet

The first aspect to characterise a fleet is the frequency of its trips. Figure 14 displays for every day of the week the daily number of trips per EV and the share of journeys that returned to the depot the next day. If a driver did not plug the EV in between two consecutive trips, it would count as a single trip: in this way the simulation would account for this common and realistic behaviour.




Figure 14. Percentage of days an EV from Primeo Energie’s pool fleet undertook one or more journeys and percentage of days the vehicles spent the night at a parking location other than the managed depot.

The plot shows that this fleet is only used on working days and even then 20% of the EVs may not be used for an entire day. However, it can happen that an EV performs more than 1 trip per day, or that the last journey involves an overnight stay at a different location.

The next important quantities to be estimated are the departing (t_{dep}) and arrival (t_{arr}) times. These clearly depend on the number of trips per day and whether they occur overnight, therefore a differentiation is made. Figure 15 presents the probability distribution functions (PDF) for these times by kind of trip, as well as the distribution of their durations. All PDFs are normalized to have the same total area under their curves, ensuring they are directly comparable despite differences in absolute counts.

⁶ The fleet based in Hägendorf was composed of two EV models: the Renault Kangoo Z.E. and the Mercedes eVito. Given their different electric ranges, these two models were used in different ways during daily operations. The vehicles participating in the project were all eVito; therefore, the following fleet characterisation, as well as all live and simulation results, refer exclusively to this portion of the fleet and its specific utilisation pattern.




Figure 15. Distributions of departing times, arrival times, and trip durations by type of journey for the PE pool fleet.

All distributions exhibit predictable behaviour: early-day trips depart and arrive earlier than late-day trips. Overnight journeys typically depart in the late afternoon and conclude in the morning of the following day.

To derive t_{dep} and t_{arr} , we can either fit and sample their respective distributions directly or fit and sample the trip duration and then compute t_{arr} based on it. Our method selection aims to retain as much information as possible. Pearson correlation analysis reveals that:

- For same-day trips, t_{arr} is strongly correlated with t_{dep} , while duration is rather independent from t_{dep} and primarily correlates with t_{arr} .
- For overnight trips, t_{dep} and t_{arr} are uncorrelated, and duration is mildly correlated with both.

This led us to independently sample t_{dep} and duration for same-day trips, and t_{dep} and t_{arr} for overnight trips. By independently sampling the uncorrelated variables, we retain the unique information each provides without redundancy, maximising information retention by capturing all orthogonal aspects of the data.

This approach is also logical because, in same-day trips, duration has a more tangible significance, influenced by factors such as travel distance and work length. In contrast, for overnight trips, duration is less meaningful, while t_{arr} is more influenced by office hours. Notably, even though we sample t_{dep} and duration independently for same-day trips, the inherent correlation with t_{arr} is preserved through their additive relationship.

Importantly, when multiple trips occur, t_{dep} of the subsequent trip is significantly correlated with t_{arr} of the preceding journey. To account for this, we modelled a linear correlation between the two and added Gaussian noise.

Finally, the last piece of information necessary for the simulation is the energy consumption per trip $\Delta e_{i,j}$. While energy consumption is the actual useful input for the simulation, fleet characterisation is more naturally expressed in terms of the SoC. This is because EV usage is more likely driven by the relative information communicated to the driver – the SoC percentage – rather than the actual unknown energy left in the battery. In other words, when comparing the usage of different EV models, we expect them to have more similar distributions in terms of SoC than in energy consumption; hence, SoC is a more suitable variable for fleet characterisation. Of course, for the simulations, the SoC would then be translated back into energy consumption, and the more realistic, diverse distributions of $\Delta e_{i,j}$ would be retained.

Figure 16 presents the observed distributions of SoC_{dep} , SoC_{arr} , and ΔSoC for the pool fleet in Münchenstein. The curves show the low usage intensity of this fleet, with trip discharges in the order of 5–15% of the SoC. Most trips depart with a full battery and the SoC_{arr} distribution is thus analogous to the



ΔSoC one. Same-day journey display the same SoC behaviour regardless of the number of trips taken during the day. However, overnight trips present a distinctly different behaviour, with driving charge depletion stretching to 20%–30% of the SoC, which is still rather conservative.




Figure 16. Distributions of SoC at departure and arrival, and distribution of typical trip discharge for the PE pool fleet.

Based on the PDF visualisations, we can fit and sample two distinct distributions for same-day and overnight trips. However, it's important to investigate whether a cross-correlation exists between the change in ΔSoC and the trip duration. Figure 17 confirms this expectation: for same-day trips, the journey duration and discharge rate are highly correlated due to the distance driven during the trip. In contrast, for overnight trips, the duration is only loosely related to the distance driven, so the link with ΔSoC is weaker or lost.

Therefore, when modelling an overnight trip, ΔSoC will be directly sampled from a fitted log-normal distribution. For same-day trips, ΔSoC will be derived from the trip duration based on a linear correlation, with added Gaussian noise to account for variability.




Figure 17. Duration (x-axis) and SoC discharge (y-axis) for every journey of the PE fleet by type of trip.



Quickpac – delivery fleet

As before, the first step is the analysis of the trips frequency. For the Mercedes eVito models located in Högendorf we obtain the following figure.




Figure 18. Share of days with 0, 1 or 2 journeys among eVitos in QP delivery fleet.

Being deliveries by car the core business of Quickpac, we find a much more frequent and regular usage of the fleet. Most cars are used exactly once per working day. In the following analysis we consider only one-trip days, being the contribution of multiple-trip days negligible. The departure and arrival times confirm the higher degree of standardisation in the fleet employment:




Figure 19. Distributions of departing times, arrival times, and trip durations for the QP fleet.

These times represent periods when vehicles were accessed prior to loading and departure or after arrival and unloading. This approach was chosen to better approximate the actual availability window for the charging process. Consequently, the duration considered is not the trip duration itself but the total time during which the EV was realistically unavailable for charging control.

Similar to the methodology used for the Primeo Energie fleet, we could either fit and sample t_{arr} directly or fit and sample the trip duration and then derive t_{arr} from it. Based on Pearson correlation analysis and the more tangible physical significance of trip duration, we again preferred to independently fit and sample the duration and derive the t_{arr} through addition with t_{dep} .

Regarding the Δe derivation, since all cars belong the same EV model, we performed the analysis in terms of SoC.




Figure 20. Distributions of SoC at departure and arrival, and distribution of typical trip discharge for the QP fleet.

These distributions also display the maximisation of the fleet usage, with several trips that reach very low charge levels. Also in this case it's safe to assume that most vehicle depart with full SoC_{dep} , so that the PDF of SoC_{arr} matches the distribution of ΔSoC .

Interestingly, the correlation between duration and ΔSoC is rather weak and most of the SoC variance is instead explained by the day-to-day journey variability. This shows that other factors were taken into account when planning the delivery trips, such as distance and cargo capacity maximisation. The sampling of ΔSoC will thus be independent of any other factor.

In summary, to characterise QP delivery fleet, only three independent distributions are required: t_{arr} , trip duration and ΔSoC .

The distributions chosen to fit the observed data are:

- departure time t_{dep} : log-normal distribution;
- trip duration: log-normal distribution;
- discharge ΔSoC : normal distribution.




Figure 21. Joint distribution of trip discharge and duration for the QP fleet.

The following figure compares the trip data coming from the project with samples obtained from the fitted distributions. We include also the arrival time to show how the chosen indirect approach still allows to adequately derive t_{arr} .




Figure 22. Comparison between original data and samples from fitted functions for key quantities.



3 Results and discussion

3.1 Operational implication and user experience

3.1.1. Fleet status and charge process awareness

With the process of introducing the charge management system on both project fleets, the usual user handling of the cars was slightly disturbed. Different signals that a user would use to make sure the charge process is set up correctly were changed and therefore unavailable. This was especially a problem since the charge initialisation process with the mechanical insertion and locking of the charging plug is still error prone. Some users would watch the charge initiation feedback to check whether the car would be properly plugged. However this signal is not available if the car is halted through the charge management. This came with the fact that various users were actually not properly trained and knowledgeable about the handling of the charging stations and the plugging process such that further instructions had to be given.

Especially with the delivery fleet, the shift managers were used to inspect the EVs on proper charging with the charging lights. Since repeatedly some charging stations would have errors leading to a stop of the charging process the EVs would get checked daily to ensure operability for the drivers the next day.

Now with the advanced charge control this would not be the case any more. A charging halt for peak shaving for example would not be able to be differentiated from an erroneous charging station that stopped charging by mistake. Of course this is especially impactful for setups that are not reliable, but operations showed that charging infrastructure is still missing some reliability margin and therefore still needs some charge verification process.

In the other case found in the Münchenstein location, an authentication system was part of the standard charge initialisation process and users would rely on the effective charge start since the successful authentication would give a sometime incomprehensible feedback. Combined with an independent charge management system this would confuse multiple users.

So the finding was that a better overview on the charge process had to be given to the user and/or the fleet manager. When being able to overlook the fleet status, and especially the charge management status, the user and the fleet manager as well would be able to understand anomalies that are not based on the charge optimisation and distinguish them from intended charge control actions. Having a remote overview would provide an even better service, since a better overview could be held and more time to work on other task would be freed.

With this conclusion, an online fleet dashboard was created, that would list all controlled EVs and give information about the vehicle state (SoC, plugged status, connection status) as well as the state in the charging process. A dedicated 'free to charge' or a 'stop for peak shaving' could be transmitted to the user which helps to understand the underlying process and differentiate from other errors that occur with the cars and or charging infrastructure.



Dashboard Primeo Pool & E-Mobility Flotte

Fahrzeug ID	Modell	VT System Live	SoC	Eingesteckt	Am Laden	Ladeprozess	Nächste Buchung
301	VW ID.3	●	82%	Nein	Nein	Keine Laderegulung	19.03.2025 - 07:30 Uhr
302	VW ID.3	●	78%	Nein	Nein	Keine Laderegulung	Keine Buchung
303	Skoda Enyaq	●	57%	Nein	Nein	Keine Laderegulung	20.03.2025 - 12:30 Uhr
304	Skoda Enyaq	●	50%	Nein	Nein	Keine Laderegulung	19.03.2025 - 08:00 Uhr
305	VW ID.3	●	98%	Nein	Nein	Ladung aufgeschoben	18.03.2025 - 13:00 Uhr
306	VW ID.3	●	100%	Ja	Ja	Keine Laderegulung	20.03.2025 - 09:15 Uhr
307	VW e-Golf	●	100%	Ja	Ja	Ladung aktiv für nächste Buchung	19.03.2025 - 07:00 Uhr
308	VW e-Golf	●	81%	Ja	Nein	Ladung aufgeschoben für SDL	19.03.2025 - 09:00 Uhr
310	VW e-Up	●	86%	Ja	Nein	Ladung aufgeschoben für SDL	20.03.2025 - 13:45 Uhr
311	VW e-Up	●	83%	Ja	Nein	Ladung aufgeschoben für SDL	20.03.2025 - 13:30 Uhr
315	BMW i3	●	100%	Ja	Nein	Keine Laderegulung	19.03.2025 - 12:30 Uhr
316	BMW i3	●	86%	Ja	Nein	Ladung aufgeschoben für SDL	20.03.2025 - 16:45 Uhr
911	Skoda Enyaq	●	4%	Nein	Nein	Keine Laderegulung	Keine Buchung
912	BMW i3	●	61%	Nein	Nein	Keine Laderegulung	18.03.2025 - 08:06 Uhr
913	BMW i3	●	71%	Nein	Nein	Keine Laderegulung	Keine Buchung
914	Skoda Enyaq	●	89%	Nein	Nein	Keine Laderegulung	Keine Buchung

Neue Abfahrtszeit einstellen (100% SoC): 911 -- : -- : Zeit einstellen

Figure 23. Live fleet dashboard for the Primeo Energie pool fleet.

These findings show how an advanced charge management system needs to be integrated in the operational processes of a fleet. This is especially the case for the verification of functionality for optimised fleet setups with little margins and for this phase when the underlying infrastructure is still maturing on its reliability and user interaction models. With an online tool this impact of the charge management system can be managed in an effective way and even add further benefits like the functionality of dedicated ready times or charge overruling for specific vehicles. It gives them a tool to interact with the optimisation system and maintain the overview of the operational fleet.

- The deployment of an advanced charge control tool showed how a perfectly reliable base charging infrastructure is crucial in order to not raise confusion during the charge process of a fleet. This is not a given standard today.
- It is key to integrate the charge process dynamics into a tool used by the fleet operators in order to keep the overview level that is otherwise simply given through ASAP charging.



3.1.2. User experience and operational implications

In order to understand the operational implications and how the users get along with the implemented charge control system the fleet users or managers were interviewed and surveyed. While with Quickpac in Hügendorf the users were not much involved, since they would only take the ready cars in the morning and give them 'back' in the afternoon, the Primeo Energie users were more involved in the management of trips and car handling. At Quickpac's site most of the operational work with the cars would be performed by the shift managers.

Those shift managers were interviewed on the operational impact to their work with the charge management system in place. In order to get the Primeo Energie user's feedback, a user survey was put in place. Both feedbacks show some similar trends, however with another focus.

Quickpac Delivery Fleet

The main task of the shift managers is to manage people and cars for every shift so that the packets get delivered throughout the day. After cars get back in the evening they had to make sure all cars would get charged until the morning. With the high fluctuation of delivery personnel and the little experience with electric cars, the actual vehicle users would not be very attentive to the charge process and would simply place the car and plug it in at any free spot and leave the operational verification to the shift managers.

They, on the other hand, had to handle multiple difficulties with EV charging: limited charging power over charging station groups (line capacity), insufficient charging stations for all EVs (shifts of charge sessions) and regular outages due to the unreliable data network for the load management of the different charging stations. With these implications already, some un-capped charging stations were identified to be used to charge vehicles with the ASAP strategy and in order to charge multiple cars one after another. This shows the ground laying complexity that was present due to a modest charging infrastructure setup.

The charge control then added some complexity to the operational processes since the shift managers could no longer rely on the cars being charged continuously after being plugged in. Unpredictable charge interruptions would mess up their plan of potentially charge other cars after the earlier cars would finish and also they could not rely on the signal they would see from the EVSE load management system. In addition to that the charge interruption would basically not look different to faulty charging station that would stop the charge for other reasons. Since this happened sporadically the shift managers would still need to verify it, since an uncharged vehicle in the morning would be unacceptable and a loss of half a day's people's work.

This additional insecurity and double checks needed on the fleet management side was considered a big additional effort and was not seen as a negligible impact. While a charge status fleet dashboard was not in place at that site this would have helped a lot to understand the charge status of the fleet. However the Quickpac personnel was already working with a fleet management dashboard and a separate handling would still add complexity to their work. Therefore it can be stated that the operational impact was negative and a seamless integration into their tools would help a lot. Further the quality and reliability of the charging infrastructure has a big impact on the compatibility with a system running on top of it, such that this can even be considered a requirement.

Primeo Pool Fleet

The pool fleet users of Primeo Energie were using the fleet quite independently via a booking tool and with rare requests to the fleet managing garage. In general the usage was streamlined and the handling left to the user or employee after a short introduction by the garage when joining the user pool.

In order to get a feedback of operational functionality and the use of EVs and the implemented charge management system, a survey was compiled and distributed among all potential pool fleet users. A



significant number of more than 100 participants shared their views on their experience with the operational system and their driving behaviour.

While the majority of the users would only drive little in frequency and distance, most of them would drive trips with 100km distance less than once a month and probably to that point book vehicles spontaneously (< 1h advance) more than 50% of the time. This has on the one side to do with the fact that the pool is large enough and no shortage situation is feared, and on the other side with the short trips, so that the impact on a following booking is small (in time and charge), so that most people get comfortable doing this. Both facts limit the actual flexibility that arises from a fleet usage in this way.




Figure 24. Primeo Energie pool fleet user survey question 5.

Further focusing on the comfort of using the electric pool fleet with the advanced management system the vast majority would not experience a negative impact on the usage and handling of the EVs and charging stations. A minimal set of 4% accounted some inconveniences to be annoying while no user would report a critical event impacting its operational needs. A more general but still minor discontent was communicated, stating that the uncertainty of the found SoC would restrict them in the use of the fleet. This would on the one hand be due to errors in the infrastructure leading to charge interruptions as well as mishandling of other users forgetting to plug and therefore charge the cars after usage. This comes right after the felt discomfort of the limited range that still has a psychological effect on users even though the vast majority is not dependent on a higher range. It however shows a similar lack of reliability in the whole charging process that is dependent on the users and also on the maturity of the infrastructure components. Adding an additional management system on top still discourages users to confidently use electric vehicles as they would be able.




Figure 25. Primeo Energie pool fleet user survey question 12.

Looking at the general confidence of adapting the usage behaviour of EVs for energy and grid flexibility, a large share is in favour and see the usage of electric vehicles as a mean to help tackling upcoming grid challenges. Some 40% of users are sceptical and don't see it being a functioning solution or even feel there is no need at all. Still the vast majority of the users (94%) are happy with starting a session with a reduced SoC of between 80%-90% in order to enable some flexibility to at least reduce charging costs. While this flexibility enabling behaviour sees a lot of acceptance, only 30% of the users are happy to add an estimated trip distance to a booking in order to adjust the prepared charge amount. This has probably to do with a big uncertainty in advance of the trip and a general higher implied risk for the user for his driving trip.




Figure 26. Primeo Energie pool fleet user survey question 15.

The full selection of questions and answers can be found in the appendix at the end of the document.



- Most users were fine having an advanced charge management system optimising the charge process as long as an error would not drastically impact their work.
- Generally there is a significant complexity at handling EVs so that an additional tool or system is seen as burdensome when not integrated with the existing systems that are already in use.

3.2 Live use cases

3.2.1. Local (building) peak shaving

Hägendorf

As an initial use case, the project setup in Hägendorf was subjected to peak shaving charge control. One challenge this setup presented was the partial roll-out of the charge control system; only a portion of the chargers were included in the charge control, while the remaining cars and chargers were considered uncontrollable loads, akin to other electric appliances.

Due to the uncontrolled charging behaviour of the remaining vehicles, the consumption curve exhibited a gradual increase in the late afternoon, peaking in the late evening. This pattern resulted from cars arriving after their daily shifts - either full-day or evening shifts. With an average charging time of approximately five hours, this led to daily consumption peaks during late evening hours and a gradual decrease starting after midnight.

Achieving optimal peak shaving results would imply approaching a flat consumption profile. This suggests shifting the flexible loads either to earlier in the evening or towards the early morning. While afternoon times might seem naturally suitable, the absence of cars during those periods renders this infeasible. Therefore, the most evident solution is to shift the load as much as possible towards the end of the charging session.

Two algorithms were gradually implemented:

- Static approach (algorithm 1): This approach would delay vehicle charging as much as possible, based on the vehicles' SoC and their required readiness times. This strategy shifted the controllable electrical loads to the early morning hours, thereby avoiding peak demand periods for these vehicles.

In a second step this approach was enhanced by introducing a charge-enabled window in the early afternoon. This allowed the vehicles to charge an initial amount during periods of low total electricity consumption, reducing the amount of energy required later and enabling further postponement of the remaining charge.

- Dynamic approach (algorithm 2.2): This approach acted as a peak load “limiter”, which activated if total consumption exceeded a certain threshold. This algorithm prioritised vehicles with a critically low SoC that needed immediate charging to be fully charged by their required readiness times. This feature ensured that all vehicles would be fully charged by the expected times, even when operating under a tight charging schedule aimed at minimising the total peak load.

Since the algorithm sought to minimise peak load rather than impose strict charging restrictions, ensuring vehicles were ready on time was given higher priority and was enforced consistently. As a result, the targeted peak load influenced the algorithm's behaviour, determining when vehicles would temporarily stop charging and when the charging schedule might be overridden to meet the readiness requirements. Therefore, having an accurate prediction or estimation of the attainable minimum peak load was beneficial. The overall electricity consumption curve showed relatively consistent patterns, so this threshold value did not require significant adjustments.



In principle all and only the 12 managed EVs should have charged at the controlled 9 EVSEs, but the urgency of daily business operations altered sometimes this agreed allocation. The following results are based on the 9 EVs that most often were charged at the 9 controlled EVSEs.

Figure 27 displays the building loads (solid lines) and EV charging loads (dashed lines) on multiple Fridays (until Saturday morning). The coloured lines depict three exemplary Fridays when active charging control was performed: the system did not simply cap the charging load, but actively rescheduled the charging sessions to the optimal time-window, while ensuring that all EVs would be fully charged by 5 am in the morning. The March controls used the first algorithm, while the May one used the second. The charge control allowed to reduce the building peak load by approximately 40 kW, while controlling a charging power of 40-50 kW.




Figure 27. Building load and charging power of EVs under management on various Friday-Saturday nights.

Optimal rescheduling of the individual charging sessions was only possible thanks to the access to the vehicles' SoC. Figure 28 displays the average SoC of the 9 controlled EV over the course of the same Friday nights. On the two managed Fridays in March the average SoC remains lower for longer time, avoiding the building peak load at 22:00, and slowly starts to increase after midnight just to in time to make sure that all cars are fully charged at 5:00. The charge control on May 3 instead allows the connected EVs to start charging following their usual tendency until about 18:00, when their charging gets halted. This allows to entirely avoid the high-peak time around 22:00 and to restart the charging process later in the night, in time to ensure full SoC at 04:30.⁷

⁷ In mid-April, the scheduled time by which vehicles needed to be ready at the Hägendorf site was brought forward from 05:00 to 04:30.




Figure 28. Aggregated fleet SoC on various Friday-Saturday nights.

In the plot one can see the so-called flexibility parallelogram, which encompasses all the profiles that the SoC could take to achieve full charge by the time required. This directly translates to the flexibility space available to DSOs and other energy or grid stakeholders.

Figure 29 presents the daily peak loads observed in Högendorf over the entire peak-shaving campaign, organized by day of the week. For each workday, the data points are sorted by week of the year, spanning from calendar week 5 to calendar week 21. The two single tests on Fridays are also depicted with the same colours of the previous plots.




Figure 29. Daily building peak loads in Högendorf; coloured points indicate days with charge controls.

Mondays and Fridays typically entail less fleet operation, hence lower peak loads, while Tuesday proves to be the most demanding workday. On average, the charge control system managed to lower the peak on every working day, although at a first glance the reduction may seem rather limited.



However, the fleet charging behaviour on some test days was affected by a few important factors:

- Along the project, some charging stations not associated with it became defective; consequently, the charging stations under the project had to be employed for more frequent charging rotations, leaving less room for charge optimisations⁸.
- The additionally charged EVs were also not part of the project and their charging sessions could not be rescheduled even if time allowed.
- The calendar week 16 was arguably one of the most demanding for the whole EV fleet.

This last point can be better appreciated in Figure 30, which compares the daily peak load with the total energy consumption for the whole day. The colour scheme is the same of the previous plots and the green points of calendar week 16 display a larger than average total energy demand. However, compared to the typical peaks expected for the amount of energy consumption, the controlled days achieved significantly lower peak loads. The linear regression (based on non-controlled and non-setup days) allows to appreciate the average peak load reduction, which was in the order of 20-30 kW also in weeks where the average controlled power was forcibly reduced to 30 kW.




Figure 30. Building peak load (y-axis) and total electricity consumption (x-axis) for each workday of the peak shaving use case in Högendorf.

Münchenstein

The local peak shaving case in Münchenstein concentrated at reducing peaks in the consumption of the Primeo Energie office building. It included the building's offices, the garage and the heating. During our live case session we however found out that during a certain time range a high power application was present while it was not during the other time. The strong assumption is, that it is the heating system (industrial heat pump), which was out for the other time. (Due to maintenance or similar) However this created two distinct circumstances since when the heating was on, this was a massive consumer dominating the peak. In the other case the garage with the EV fleet was actually being a relevant share of the consumption. Therefore in the first case with the heating one can speak of a "load drop" scenario for peak shaving whereas in the second case it was a "load balancing".

⁸ The number of defective charging stations grew along the project because the foreseen relocation of the fleet made their replacement or fixing not particularly urgent.



“Load drop” case with heating peaks

Since the heating peaks were very dominant with around 80kW additional consumption which meant for peak shaving that basically any consumption would be dropped during that time. At the same time it meant that during all other times it was irrelevant on how much power to use, because it would anyway only be below the heating peak. This led to a case where the EVs would normally charge and only stop the charge if a heating peak would arise. For effective peak shaving an overlap of heating and current EV charging would be needed. However as discussed before since all the other times with inactive heating would be used to charge the EVs at full power, there were only limited overlaps, since the EVs with the pool driving behaviour would have finished its charge already. Regularly it would overlap and reduce the potential total peak, however it was not guaranteed that an effect would be shown especially in absolute peaks that happened irregularly. In general the behaviour was very comparable to the fast charger peaks on the L7 transformer level discussed later.

The following table shows the frequency of charge and power overlap and the effective reduction and potential peaks that were prevented. (Shaded are days with the weekly peaks, in orange the average value)

Week 1: Control week (no charge control)

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	13:35	11:50	09:30	07:30	08:00	
Peak	187 kW	184 kW	204 kW	215 kW	184 kW	194.8 kW
# halted EVs	No charge control (control week)					

Week 2: Live week

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	07:00	08:00	06:50	09:45	08:00	
Peak	182 kW	147 kW	189 kW	199 kW	179 kW	179.2 kW
# halted EVs	1	0	2	0	0	0.6
Pot. Peak	192 kW	147 kW	209 kW	199 kW	179 kW	185.2 kW

Looking at the numbers of the daily peaks between the weeks one can see that the main peak was lowered and that it could be reduced. But on the other hand also that the average peak was lower in the controlled week even of the potential peaks when considering the amount of stalled charging sessions. Nevertheless without charge control the main peak would have been higher in that week and so the charge management was able to hinder an increase in the main peak due to EV charging.

Generally it can be said, that the fleet setup in Münchenstein generally does not add a large load to the peaks since the charge sessions are quite distributed and the chance of overlapping with a peak is not very high. It does still make sure that the controlled vehicles do not worsen the highest peaks when they arrive. This is clearly the case since there is enough charge flexibility for those short-term peaks and only a negligible amount of charging sessions are prioritized due to operational needs.



“Load balancing” case without heating peaks

In this case where no heating system was dominating the building's load, the EV fleet charging had a comparable share to the overall load with the rest of the building. In this setting a charge control for peak shaving would be active in the way of a 'load balancing' that would balance the load around a target peak value and only enable charging in times of low consumption or just with a limited number of cars.

This control setup lead to the following daily peak number over active and control weeks:

Week 1: Live week

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	10:00	12:10	16:00	07:30	08:10	
Peak	71 kW	56 kW	51 kW	58 kW	52 kW	59.4 kW
# halted EVs	2	1	2	2	2	1.8
Pot. Peak	91 kW	66 kW	71 kW	78 kW	72 kW	77.4 kW

Week 2: Control week (no charge control)

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	12:15	14:00	07:20	07:25	08:05	
Peak	72 kW	59 kW	66 kW	58 kW	68 kW	64.6 kW
# halted EVs	No charge control (control week)					

Week 3: Live week

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	08:20	07:50	08:10	08:00	14:00	
Peak	65 kW	66 kW	58 kW	64 kW	56 kW	61.8 kW
# halted EVs	2	3	2	0	1	1.6
Pot. Peak	85 kW	96 kW	78 kW	64 kW	66 kW	77.8 kW

The results of the live weeks and the control week shows on the one hand that the average daily peak could be reduced and on the other hand that a relevant amount of cars were halted from charging during the respective peak times. While the absolute peak could only be reduced in one week it did not in the other. However during that time 2 EVs were on hold with the charge, which still had enlarged the peak by a relevant amount in an uncontrolled scenario. This can't directly be accredited, because cars would only charge for short periods and not controlling the charging would probably have left the charging of one or the other cars finished before the actual peak.

Nevertheless it can be stated that the average peak could be reduced by some 3-5 kW or potentially up to an average of 15kW. This is not a realistic measure for the reason mentioned above. However the 3-5 kW are also rather low since the 3 weeks do not give a significant data base to make a clear statement.

In the second live week the actual week's peak was potentially reduced by 30kW. It can be a lucky shot as on other days it had 1 or 0 cars during these times but it still shows that it can reduce the peak by a large amount. This vague and potentially low impact comes along with the pool fleets driving charac-



teristic of short rides during the day. With this behaviour comes along that only short charging sessions are necessary and they only overlap rarely.

- The effect of local peak shaving is highly dependent on the driving pattern and the associated energy revenue of the fleet. A large energy revenue increases the charging time and therefore the charge sessions' overlaps and therefore the controllable power.
- On the other hand the impact is highly dependent on the appearing base load peaks, where the regularity and time correlation has a significant influence on the reduction potential.
- In general the potential is considerable in high energy revenue fleets.
- The current goal given as the highest peak of the month seems questionable since the reaching of it gives little incentive to generally reduce power needs otherwise.

3.2.2. DSO grid peak shaving

Hägendorf

In a second step, the grid peak shaving case was planned for the Hägendorf site. Unfortunately, the local DSO did not allow the installation of a meter at the transformer that supplies the Quickpac depot. Therefore, at this site the peak shaving use case could not be carried out for level 7, but only for grid level 5. However, the latter was not a live control case since 1) no real-time measurement equipment is available at that grid level and 2) the impact of the charge control actions is negligible.

An analysis of historical transformer line usage revealed regular peaks occurring precisely on the hour, with noticeably stronger peaks on the hour during the night. The most significant peaks were at 21:00, midnight, 01:00, and 03:00. These peaks almost certainly result from ripple control activations of electric boilers, which draw power during these staggered activation periods. Ripple control is a method used by utilities to manage load distribution by sending signals over power lines to control high-demand appliances.




Figure 31. Current flow monitored at the level 5 grid feeder to the Hägendorf site on the five tested weekdays. Clear spikes on the hour are visible especially during the night.



Since these peaks typically lasted only 15 to 20 minutes, implementing a brief reduction in power or load for half an hour could effectively mitigate them. Because the power peaks consistently occurred at the same times, we could establish fixed off-times, allowing the charging control system to avoid charging during these periods. This strategy raises an interesting question: to what extent does the charging of EV fleets overlap with these congested times, and how effectively can we reduce those peaks through appropriate charge control?

Due to the lack of live information from the transformer, only static control during these pre-established off-times was possible. It was decided to hold the charge process of any controlled, charging, and not urgent EV for half an hour at 21:00 and 00:00. Being the learning limited by the lack of dynamic control, we opted to test this use case for a single week in May 2024. In that period, 9 EVs were still available for control most of the nights.

Figure 32 shows the charging load of all the EVSEs belonging to the project – used also by all the project EVs and a few non-project EVs – for the 5 working days of the selected week. Clear charge reductions lasting for half an hour are visible at 21:00 and 00:00. With 9 controllable EVs, load reductions in the order of 30–40 kW were achieved. The sharp current increases of Figure 31 amount to about 10 A or 0.3 MW for a 16 kV grid. This means that the observed peaks of the medium-voltage grid could have been balanced by a simple pool of 100 available connected EVs.




Figure 32. Charging load of the EV fleet in Hägendorf on the five tested weekdays: interventions on the portion of controlled EV are visible at 21:00 and 00:00.

It is important to note that these peaks are quite specific to the current scenario and are likely to change in the foreseeable future. With an excess of energy becoming available during daylight hours due to increased solar power generation, boilers may be triggered during those times instead. This shift would lead to different consumption patterns and consequently alter peak load times. While predicting how the grid's load profiles will evolve is beyond the scope of this project, we observed an alternative grid load profile at the Münchenstein site, which can inform more forward-looking considerations.

Münchenstein Level 7: Fast charging balancing

Also at this site, one of the primary use cases involved reducing the local peak load. A preliminary analysis of load distribution on the level 6 transformer feeding the Münchenstein site showed that the highest



peak loads were primarily driven by a couple of fast-charging stations on the premises. This finding is significant, but also intuitive: the DC power rating of most new EV models reaches up to 150–200 kW, while typical household power consumption averages between 0.5–1 kW, with occasional peaks of 10–20 kW. Consequently, a single fast-charging session is equivalent to the sudden, simultaneous demand of a ~10-unit apartment building at peak usage.




Figure 33. Breakdown of level 6 transformer load in Münchenstein by building/consumer for a typical week. Loads are stacked and the high red peaks are due to fast charging.

For this use case, we thus decided to set up a dynamic charge control system that would monitor the fast charging load and intervene on the slow garage EVSE in case of employment of the fast charging stations. This approach aligns with the different purposes and philosophies of the two types of charging sessions: slow charging is typically not urgent and done for “opportunity” reasons (i.e., it costs nothing while the car is parked regardless). The fast charging is instead an urgent immediate request that should be accommodated in almost any case, and measures can and should be taken to safeguard the connected grid.

Three weeks in August–September 2024 were allocated to run first tests. During this period, eight EVs were available for charge control operations; however, as the fleet analysis for Münchenstein indicates, these EVs have a lower usage intensity compared to those in Hägendorf. This reduced usage limits the flexibility to advance or delay charging, as the overall energy demand is lower. For example, if an EV is charged to 100% SoC by morning but remains unused, there is no further charging demand that can be shifted to provide relief to the grid. Consequently, no more than three EVs were observed simultaneously participating in this load management use case.

Figure 34 illustrates a representative day in which three EVs were charging in the garage when two external fast charging sessions took place. Since the SoC and booking schedules of the three EVs allowed flexibility, the system was able to pause and hold their charging processes until the fast charging sessions concluded.




Figure 34. Example of charge control intervention on three EVs to balance the sudden load caused by the external fast charging stations.

Given the limited availability of control actions at optimal times throughout the week, no substantial reduction in the absolute peak load was anticipated for the Münchenstein site. Therefore, Figure 35 presents multiple KPIs for each monitored week, including the count of 5-minute intervals where transformer load surpassed a specified power threshold. Seasonality effects are apparent from the plot, with a noticeable drop in load during the mid-summer weeks of July and August. However, a slight reduction is observed also in the targeted weeks (34–36) compared to the adjacent, uncontrolled weeks. While these reductions are modest – primarily due to the low utilisation of the fleet – they indicate that even this straightforward use case may offer meaningful relief to both the grid and transformer load.




Figure 35. For ever calendar week: mean and maximum transformer load (lines and left y-axis); count of 5-min intervals with load above a certain threshold (markers and right x-axis).

Later in the project another 3 weeks test was run with the full fleet being under control, which still did not show significant improvement. The general driving and charging behaviour of the Primeo Energie fleet,



which shows a spread pattern where the cars only charge short amounts of time during the day, shrinks the probability of having multiple cars charging at the same time of appearing peaks. Therefore the impact of a load drop case is only limited since next to sometimes not even being charging at those time, most often its only a small fraction of the fleet that is charging, if not only single cars. This diminishes heavily the added value of such a charge control case. However not considering it and running another charge optimisation case can lead to disadvantageous situations on the transformer level.

Below are two lists depicting the different occurring peaks and how the pool cars would be able to control their charging process accordingly. A mentionable effect was the exact same numbers of peaks and also reduced and night peaks in both weeks which seem to show a quite steady appearance patterns throughout the weeks.

Live week 1

Weekday	Monday	Tuesday	Wednes-day	Thursday	Friday	Average
Peak type	FC	FC	Other	FC	Other	
Peak reduced	Yes	No	Yes	No	No	
# halted EVs	1	0	1	0	0	0.4

Total FC peaks: 18

Total reduced peaks: 8

Total FC peaks in evening/night hours: 6 (Potential reduction through delivery fleet)

Live week 2

Weekday	Monday	Tuesday	Wednes-day	Thursday	Friday	Average
Peak type	Other	FC	FC	FC	FC	
Peak reduced	Yes	No	Yes	Yes	No	
# halted EVs	1	0	2	1	0	0.8

Total FC peaks: 18

Total reduced peaks: 8

Total FC peaks in evening/night hours: 6 (Potential reduction through delivery fleet)

Münchenstein Level 5: District peak avoidance

In contrary to the L5 transformer data from Hägendorf, where Quickpac's depot was located, in Münchenstein the load profiles would not make out clear peaks. In Hägendorf the peaks were dominated by the hourly ripple-control activations in the nighttime. In Münchenstein however the daily peaks would mostly aggregate in the morning hours. In the hours from 8 to midday the highest consumptions would be visible with a still varying peak time. In some cases the peak would also move to the early afternoon hours.



This load behaviour is still quite typical for this grid level. No clear repeating peak but with an aggregation around a part of the day. This goes into a static setup where we can speak of 'peak avoidance'.

With this setting in mind a charge control setup was configured to avoid morning charging due to the aggregated peaks in the morning. Since there were still some afternoon peaks a dynamic peak shaving was put in place to react on spontaneous peaks in the afternoon hours.

The transformer data we had was only historical so that we simulated the grid data for the charge control by feeding "quasi-live" data from a simulated meter with data from the according month in the previous year.

This combination of algorithms would be a special case and demonstrate the ability to work with two concepts while trying to reduce charging during peak times. The dynamic control part would add a significant complexity that would normally be solved with a longer 'avoidance time' with the drawback of longer charge halting times.

The tables below show the appearances of peaks during the test weeks and how the control could affect the daily peaks. (In shades there are the weekly peak days and in grey are peaks outside the basic peak-time, which is the morning.)

Week 1: Live week

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	-	12:00	13:00	10:35	10:15	
# halted EVs	-	2	1	3	1	1.75

Week 2: Control week (no charge control)

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	12:35	10:18	09:40	09:00	07:40	
# charg. EVs	0	1	0	0	0	0.2

Week 3: Live week

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	07:40	08:05	11:20	15:20	07:20	
# halted EVs	2	2	2	3	1	2.0

Week 4: Control week (no charge control)

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday	Average
Peak time	08:10	08:10	11:20	08:40	11:05	
# charg. EVs	0	0	0	0	1	0.2

- As on the local level the appearance pattern of external peaks are relevant for the impact, with the observation that higher level grids show more regular and static peak behaviour.



- Again a high energy revenue fleet is able to have a bigger impact on the reduction goal.
- The general correlation of the grid peak and the fleet charging times highly impacts the flexibility potential.
- The grid peaks seem to follow a very local behaviour and differ significantly in between sites.

3.2.3. Ancillary services

The ancillary services use case was only implemented and tested at the end of 2024 due to the project timeline and time dependencies with the project partner's asset pooler. As described in section 2 a solution was implemented to add the controlled fleet as a flexibility asset to the Equalio VPP (virtual power plant) or asset pool. The pooling happens in order to comply with the 5MW minimal power unit that the TSO Swissgrid auctions for the ancillary service provision and activation. That means that the availability and the activation is shared with a large asset pool that performs the same ancillary service.

While this concept is very beneficial as it adds the possibility for smaller assets to also participate in this service auction, the downside is that the whole pool with different assets and different working principles are auctioned in the same way and are therefore also activated the same way. As we describe later some different bidding strategies lead to different activation patterns (activation frequency).

Ancillary service product considerations for EV fleet control

In order to understand the most useful and matching application for the EV fleet flexibility the different products were investigated on time requirements, timing requirements and energy requirements.

Negative ASP

With negative ancillary services, energy is taken out of the electricity system which is equivalent to a charge power increase in the case of EV fleets. The most natural way to do this with EVs is to start charging from a state of holding the charge. The other option is to increase the charging power, while this signifies that the charging process needs to be running at the moment of ASP activation. This on the one hand reduces the power increase range (often by 50%) and the availability on the other since this means that only the charging times of the cars would be considered.

In the context of a fleet that has a high charging time portion of the plugged time the second approach could make sense. However most fleet will have a relatively low charging time portion of the plugged time and therefore the first approach would be beneficial by increasing the availability time as well as the provided power.

Positive ASP

The positive ancillary services entails the process of adding energy to the electricity system which means that in a non-V2X capable EV context this would mean a reduction of the charging power. This implies that the corresponding EVs will need to be charging to be available to the service. The V2X capability however adds the possibility of feeding power to the grid and supplying a positive ancillary service to the TSO without the need of being charging during the availability/activation times. Looking at the standard case the trade-off is in trading charging power for service power versus availability. The larger your charging power, the larger your supplied power service, however the shorter your availability. This service therefore is highly dependent on ones charging needs since it defines the time the system is available for the service. A fleet with high charging time portion of the plugged time is very beneficial for this application.



Power provision (PRL/SRL/TRL)

The power provision is generally seen as the availability provision. A certain amount of power is reserved for a certain amount of time to be activated for a varying amount of time. The difference between the categories (primary, secondary and tertiary) lays in the response speed that is expected. The actual service is auctioned by Swissgrid.

Being the main difference, primary reserve can only be offered on both directions (positive and negative) and is normally evened out in energy. Also a fast reaction time of 2 seconds is mandatory. Both constraints make it unfeasible for EV fleets since most charging and charge control setups do not comply with the timings and a bidirectional activation needs a constant medium charge level which can then be adjusted in both directions. On this second constraint especially with standard (AC) charging a lower charging limit of cars reduces the bandwidth of playable charge power which adds another barrier for such an implementation. For a implementation with DC chargers and V2X ready cars please consider the 'V2X Suisse' project that focuses on this application.

Next to the timing difference for secondary and tertiary control, currently also the availability slots that are auctioned by Swissgrid differ. While for tertiary control (TRL) 4h slots are available for power provision bidding, secondary (SRL) control uses a full week as bidding slot. This has a large impact on the suitability for EV fleets. EV fleets still mainly exist in order to be driving rides, especially during the day. Therefore the EVs are generally not available during that time and the effective power that can be provided is limited to the actual minimum available cars during the week. This value mostly goes towards 0 as the analysis shows further down. The 4h time slots for tertiary control however make it possible to subdivide the availability and bid or offer higher power availability for corresponding times with high EV availability.

Power/Energy activation (SRE/TRE)

The activation of the ancillary service power is considered as the "energy" that is auctioned and is done in a separate bidding auction. Every asset that is providing the power provision (availability) necessarily also takes part in the energy auction. However it can bid a high energy price and would therefore only be called in a later time once cheaper assets would have been called.

Swissgrid offers a separate auction way called "free bids" that enable to participate in the energy auctions without being active in the provision (availability). This auction is done short term (1h in advance) and can therefore be settled in a very spontaneous or flexible way. However those offers compete with the offers that the assets in the provision give out without making the provision revenue. Therefore the sole participation in the "free bids" lowers the potential revenue. Nevertheless it would enable a quite irregular availability of EV fleets to still participate in the ancillary service market with its flexibility and use a situational bidding strategy that meets the fleets providing possibilities at that point of time.

The challenge with the fleet-optimised "free bid" setup is that a large enough fleet pool needs to be assembled with a comparable usage such that the minimum power package size of 5MW is met. This has a big impact in the scaling of such a setup or product and therefore brings along high hurdles.

Timing considerations:

Reaction Timing

For the participation and the supply of ancillary services Swissgrid has specific timing constraints in place for each product category of the ancillary service products. While for primary control the reaction time lies below 2s the secondary control has a full activation time of 2 minutes and tertiary control even at 10 min. However it is expected that a certain fraction of the power is already delivered earlier following a ramp up slope given through a specific rate. The graphic below depicts the ramp margin that is accepted for the SRL qualification.




Figure 36. Test signal curve with timing margins for secondary control given by Swissgrid for prequalification.

While for secondary control the full ramp up lies at 200s (about 3.5 minutes) reaction time, the initial ramp must be available after 10s. The tertiary has a full ramp time of 10min with varying reaction times that are asset dependent and have a significant flexibility margin.

The found reaction times of the charge control system showed a clear inability to perform primary control, a very tight (in-)ability to perform secondary control and the ability to perform tertiary control. This is given by the reaction time of minimally 10s for any change of charge status while the by far biggest chunk would come from the internal charge status delay of the cars themselves. The actual latency from the issuing of a request on the MQTT broker down to the control signal of the car was below 1s (max ~500ms). However, the current state change was dependent on the cars actual status especially for the charge start transition. The initiation of the charging would in contrary to the other state changes need up to 60s which stands in contrast with the 10s needed for stopping the charge and changing the charge current (and therefore power) of the charge process.

This fact makes it incompatible to use a charge start for secondary ASP, since a initial power delivery is expected after 10s. Therefore only a charge current (/power) adaptation is legitimate for secondary control and this clearly limits the scope since a constant charge is needed in order to be able to adapt a running charge process.

Figure 37 shows an example of ancillary control action with a stepwise increase and deactivation of negative (secondary) control. One of the available cars is chosen to follow the requested ramp.



SDL activation requests



Figure 37. Ramps followed by available pool car for negative secondary control. Upper graph: Ancillary power request; middle graph: load measurement data from garage meter; lower graph: pool fleet vehicle status.

The measurement of the garage meters collect data points every 2-3 s. With this resolution following timings were measured:

Delay type	Delay
Charge start delay:	56 s
Charge reduction delay:	8 s
Charge increase delay:	3 s
Charge stop delay:	8 s

With the analysis of a vast set of control actions, also including other control pipeline, like the direct vehicle API or the EVSE control, following average values could be observed.

Delay type	Avg. Delay	Max. Delay
Charge start delay:	45 s	<60 s
Charge reduction delay:	8 s	<10 s
Charge increase delay:	8 s	<10 s
Charge stop delay:	8 s	<10 s



Reporting

The reporting for the accounting of the activations needs a specific and dedicated process that tracks and logs the differences of the expected power balance of the asset when activated for an ancillary service. While for bigger assets this is done with a dedicated high quality metering equipment in the case of a EV fleet this is not as easily implementable and quite impactful for the costs implied for the ancillary case. While a metering system for the whole fleet is possible, often this setup with no other impacting third party asset is found very seldom. As viewed in the picture above, the actual meter's data is an aggregation of multiple loads, in this case of uncontrolled vehicles. As seen at 09:47 another vehicle stops charging and soon later a third car starts to charge before the controlled car stops the charge process due to the stop signal by the ancillary pooler. So there is a power activity which is not from the active asset and needs to be filtered out for the reporting of the ancillary service activation. An additional logic for disaggregation of the single cars/assets is needed that can be very specific to the found setup. This adds a further hurdle to the implementation of a ancillary service provision with EV fleets and also additional costs.

Timings of the activations

While on the one hand the specific timing requirements on the activation slopes is crucial, the timings of the activations themselves are also very impactful on the case itself. An analysis that is described below in detail shows that secondary control comes along with potentially very short activations (<1min) while tertiary control has a minimal activation time of 15min. The below minute activation cycles are not very suitable for charge control since especially the charge start process plays in the same order of time. So for SRL it can be stated that charge interruptions seem to be more reasonable (SRL+) while for charge session start a longer activation time like with tertiary control is more favourable.

- While not incorporating V2X technology, the fleet usage pattern dictates whether negative or positive control can be offered in a reasonable way.
- The timing constraints asked for the secondary control do not match the possibility of the charging infrastructure (including the cars) by nuances. It feels the unmatched constraints are historically given and not relevant to the flexibility product.
- The punctual limited availability raises need for specific asset bidding and therefore a specific asset pooling in order to reach the TSO bidding size which makes it very hard to start with small numbers.

Ancillary services - Fleet activations analysis

Average available power

During the live operation of ancillary services in Münchenstein, we recorded the number of EVs whose charging processes were temporarily put on hold, making them available - if activated by Swissgrid via the Equalio platform - to resume charging and thus provide the expected negative control energy (SRL- or TRL-). For the purpose of the demonstration, the offered available power was consistently limited to 10 kW, corresponding approximately to the power offered by a single charging EV. However, we also observed instances in which two or more EVs could have been activated simultaneously, meaning greater capacity was technically available.

The average electric power that was potentially available for negative control over the course of a typical week is presented in Figure 38.




Figure 38. Average "connected" power available for negative control (charging) in Münchenstein.

The figure clearly illustrates a distinct daily pattern: as the pool EVs return from their daytime missions and services, they are parked in the garage and their charging processes are deliberately put on hold for as long as operationally feasible. As a result, most of the available power for negative control accumulates in the late afternoon and evening. This availability gradually declines overnight into the morning hours, as many EVs have a low SoC and have scheduled trips that necessitate the initiation of charging, overriding the secondary objective of providing flexibility for ancillary services.

This pattern does not hold on weekends: on Saturdays and Sundays, there are typically no scheduled trips (with rare exceptions), so EVs are kept on hold for longer times. They exit the pool of vehicles available for negative control only if they disconnect from the charge management system or when their 12V battery reaches a critical threshold. This explains the more gradual decline in power availability from Friday into the weekend, reaching the lowest point on Monday morning.

The plot also includes inter-week variability, represented by the 25th and 75th percentile lines. These percentile bands follow the average trend but highlight that, particularly in the morning hours of weekdays, there are non-negligible chances of having no EVs available for negative control at all.

Activation requests

The following analysis provides an overview of SRE activation requests received from the Equalio platform between April and September 2025. While these results are likely influenced by the pooler's specific bidding strategy and may not be broadly generalizable, they offer a first impression of how often the Primeo Energie (PE) pool fleet was actively triggered during the live ancillary services demonstration and help to derive some initial insights.

The analysis covers activation data for both SRE+ and SRE-, though only SRE- requests were integrated with the charge management system and resulted in actual control actions. SRE+ activations were recorded purely for statistical purposes.

The first figures show the histograms of activation requests, categorized by their direction and duration: the plot reveals that the fleet received more SRE- than SRE+ calls, but more importantly that most activations were very short, usually under 10 minutes and very often shorter than 1 minute. However, there were a few longer events, including some exceeding 1 hour, with the longest continuous activation lasting 3 hours.




Figure 39. Number of received activation requests by direction and duration: on the left, across all time durations; on the right, zoomed in to durations shorter than two minutes.

The next two plots illustrate how activation requests were distributed by day of the week and hour of the day. The weekday distribution appears mostly random, with one notable exception: Sundays recorded the highest number of activation calls. The hourly distribution is more revealing, showing a clear concentration of activations during daylight hours. This combination suggests a possible correlation, especially for SRL-, with imbalances caused by higher-than-forecasted PV generation. These mismatches are more impactful on the grid during low-load periods, such as Sundays. Furthermore, it is worth noting that until recently, many utilities and balance groups focused their PV forecasting models on working days, leading to less accurate predictions for Sundays and Mondays.




Figure 40. Frequency of activation requests by direction and day of the week (left) or hour of the day (right).

These distributions reflect only the frequency of activation calls, not their duration or requested power levels.

The next two figures address this by showing, for the same weekly and hourly breakdowns, the average requested power, including all periods with no activations. The low average power levels reinforces the insight that activation calls are generally short and infrequent.




Figure 41. Average active power by control direction and day of the week (left) or hour of the day (right) over the entire time period (i.e. including no activation times).

Given that the EV fleet provided 10 kW of both positive and negative control to the ancillary service provider (ASP), we can approximate the expected activation time share by dividing the average requested power by 10. This reveals that, on average, the fleet would be activated for negative control for about 1% of the time (0.6% in the case of positive control). During daylight hours, especially for SRE-, this likelihood increases to around 2.5–4%. While these figures are low, they are comparable to the national average for SRE activation frequency in Switzerland, which is roughly 4%.⁹ These observations also support the finding that the connected ASP follows a “premium” bidding strategy, targeting less frequent but higher-value activations.

The sharp spikes in average SRE+ power in the hourly plot suggest that the data is still influenced by isolated large events, such as occasional long-duration activations. Given the relatively short observation period (6 months) and the likely high-price bidding strategy of the ASP, the number of activation requests captured in the dataset is inherently limited. As a result, while the data provides useful indicative insights, it does not support broad or statistically robust conclusions.

Ancillary services - Economic assessment

While the live deployment of the ancillary services use case enabled the validation of the full technological stack required for its effective and reliable operation, the relatively short duration and fixed time window of the demonstration limit the ability to draw broad conclusions regarding its economic viability.

To address this, we conduct a broader assessment that remains grounded in the specific monitoring data from the project’s fleets but extends the analysis across a generalized time frame and longer duration.

The assessment follows these steps:

- Quantification of the electric power made available by the EV fleets for the provision of ancillary services.
- Characterization of average market prices for secondary and tertiary reserve capacity (SRL±, TRL±).
- Characterization of average revenues from balancing energy provision, i.e. from actual activations (SRE±, TRE±).

⁹ [Swissgrid \(2020\). TSO report on balancing.](#)



As discussed above, the most promising use case is the provision of negative control energy through the initiation of charging processes. In particular, we focus on the economic potential of negative tertiary control (TRL-) for two main reasons:

- The required availability window for TRL bids can be as short as 4 hours, provided the ASP participates in the daily TRL auctions. In contrast, SRL is at the time of the project awarded on a weekly basis, which the project's EV fleets could not consistently guarantee.
- The response times of electric vehicles are better aligned with TRL requirements. Moreover, the 15-minute activation duration of TRL is far more compatible with typical EV charging response than the shorter activation times (few seconds to a few minutes) seen in SRL activations.

The analysis below therefore concentrates on this specific use case, with a focus on the PE pool fleet.

- Activation requests for secondary positive or negative control were often very short, frequently below one minute.
- On average, Primeo Energie's pool fleet would have been activated for 1% of the total time it offered negative secondary control.
- Negative activation calls were more frequent on weekends and during daylight hours, suggesting a correlation with PV-induced imbalances.

Quantification of EVs availability for the provision of ancillary services

The figure above presented the average weekly profile of electric power made available by the PE pool fleet for the provision of negative control (TRL- or SRL-). The analysis of price signals for TRE- and TRL- presented below indicates no meaningful weekly pattern, but rather a more pronounced daily variation. Accordingly, we refine the whole analysis by transitioning from a weekly to a daily resolution.

To ensure consistency and relevance, we focus exclusively on working days, as they show more stable usage patterns and greater flexibility potential for negative control. Working days are defined here as the period from Monday 12:00 to Saturday 12:00.

In Figure 42 the orange line represents the average power available for negative control across working days. For a realistic economic evaluation, however, it is essential to consider the specific conditions of the ancillary services tender process. In particular, Swissgrid allows ASPs to bid in daily auctions for TRL- in 4-hour time blocks. This means that ASPs must guarantee access to the full bid capacity continuously for the entire 4-hour window. Assuming that the ASP does not aggregate multiple assets to meet this requirement (i.e. no asset pooling), the EV fleet can only offer the power it can reliably sustain for 4 consecutive hours. This constraint motivated the creation of the second blue line in the figure: it represents the minimum power the PE pool fleet could effectively provide in each 4-hour slot, averaged across all relevant days and weeks of the live operational campaign.




Figure 42. Average "connected" power available for negative control (charging) with the PE pool fleet in Münchenstein. In orange at 15-minute resolution; in blue, after taking the minimum availability within each 4-hour block.

As expected, the power made available by the fleet when applying a 4-hour continuous availability constraint is lower than when assessed at a 15-minute resolution. Nevertheless, the analysis shows that the PE pool fleet can consistently provide 20–30 kW of negative control between 16:00 and 04:00 on working days.

This insight prompted a broader question: how does the required duration of availability for ancillary services (SRL/TRL) impact the flexibility potential of this EV fleet? To explore this, we generalized the previous analysis by evaluating multiple availability durations - not just 15 minutes and 4 hours. For this broader assessment, weekends were also included.




Figure 43. Depending on the length of the continuous availability period (x-axis), the y-axis shows the percentage of periods with N cars available (colours) for negative control with the PE pool fleet.



The resulting bar plot demonstrates that the fleet can offer sufficient flexibility for negative control across a wide range of short-term availability windows, up to 4 hours. However, as the required duration increases to 12 hours or a full day, the average number of EVs available drops significantly. For an entire week, the fleet cannot reliably provide even the equivalent of a single EV under continuous availability constraints.

This result underscores why providing SRL⁻ is practically unfeasible without aggregating multiple assets, such as combining EV fleets with other types of flexibility resources. Conversely, the bar plot highlights that the current 4-hour availability window required for daily TRL[±] tenders is well-aligned with the inherent flexibility of EV fleets for negative control.

Characterization of average prices for capacity provision (TRL[±])

The only capacity tender that the EV fleet can currently participate in under existing regulations is the daily tender for TRL[±], where ASPs place bids for each 4-hour block. The following analysis assesses whether systematic daily or weekly pricing patterns exist in the TRL[±] market, as these would significantly influence the revenue potential of the EV fleet. The analysis is based on 2025 market data, during which daily TRL[±] tenders were held consistently. This contrasts with 2024, when TRL[±] was predominantly awarded through weekly tenders.

The results from Figure 44 indicate that TRL⁻ is considerably more profitable than TRL⁺. This disparity can largely be attributed to Switzerland's hydropower-dominated flexibility landscape: for many hydroelectric generators, offering TRL⁻ capacity entails committing to a minimum generation schedule, even during periods of unfavourable market conditions, leading to higher opportunity costs.




Figure 44. Average prices for TRL daily auctions (2025) by direction and day of the week (left) or time block of the day (right).

The left figure displays the average awarded TRL capacity prices by day of the week: no meaningful weekly pattern emerges, supporting the earlier conclusion that weekly temporal structure can be disregarded in revenue estimations. The right figure instead shows a pronounced daily pattern, particularly in TRL⁻ pricing. Two key factors explain this trend:

- High PV generation during midday hours typically leads to low spot market prices, increasing the opportunity cost of mandatory generation for TRL⁻ provisioning.
- Forecast errors in PV output remain substantial during peak solar hours (12:00–16:00), increasing the need for balancing capacity from the TSO.

These combined effects result in elevated TRL⁻ prices during the 12:00–16:00 time block, highlighting this window as particularly promising for consumption-based flexibility bids, such as with EVs.



Characterization of average revenues for balancing energy ($TRE_{\pm}$)

A comparable analysis can be performed for the price trends of balancing energy. However, the actual revenue for the ASP and the underlying flexible assets (the EV fleet) is not simply equal to the activation price paid by Swissgrid for regulation energy: in fact the ASP is also subject to a financial settlement with the balance group where the flexible assets are physically located. When an activation occurs, it is the balance group – not the ASP – that physically delivers or absorbs the activated energy via its connection to the flexible assets. This leads to the need for a commercial energy settlement between the ASP and the balance group. This settlement is carried out at the Swiss day-ahead market price (SwissIX) for the corresponding 15-minute period. Therefore, the net revenue for the ASP during a balancing energy activation is determined by the difference between the activation payment (based on the bid or clearing price of the specific TRE product) and the SwissIX market price at the same time.

Figure 45 analyses the weekly and hourly patterns of the ASP's net margin for balancing energy, as derived using the settlement mechanism described above. The trends for TRE_{-} and TRE_{+} resemble those observed for TRL_{-} and TRL_{+} , with TRE_{-} being more profitable, especially during periods of PV generation. This primarily results from the asymmetric forecast errors in PV generation and from the way the final price for $TRE_{\pm}$ is set: the MARI and TERRE platforms (international cooperation platforms for tertiary balancing energy) operate on a pay-as-cleared basis, meaning the price of the marginal activated bid determines the remuneration for all activated assets in a given interval. This mechanism amplifies the effect of high-cost marginal bids during imbalanced conditions.




Figure 45. Average net earnings of ancillary service providers for tertiary regulation energy (2025) by direction and day of the week (left) or time of the day (right).

Similar to TRL, no meaningful weekly pattern is observable in TRE prices. However, the hourly variation remains significant and will be considered in modelling the economic potential of EV fleets in providing balancing energy. One notable exception is a price spike for TRE_{-} on Sundays. Unfortunately, this effect is lost when we limit our economic analysis to working days only, following the power availability patterns of the EV fleet.

The final step needed for the economic assessment is an estimate of activation frequency. This is the most uncertain input, as it strongly depends on the bidding strategy of the ASP (whether targeting fewer, high-paying activations or more frequent, lower-priced ones). For this analysis, we assume an average behaviour, meaning we estimate the average activation frequency of all assets offering $TRE_{\pm}$ in Switzerland and apply that rate to the EV fleet.

Figure 46 shows the average activated MW for TRE_{+} and TRE_{-} by day of the week and by hour of the day. As in previous analyses, no significant weekly pattern is observed. However, clear daily patterns



emerge: TRE- activations peak during daylight hours, while TRE+ activations occur more often during periods of higher national load.




Figure 46. Average amount of tertiary reserve capacity activated by direction and day of the week (left) or time of the day (right).

To calculate the probability of activation in each time window, we divide the average activated MW by the total capacity procured via TRL (480 MW for TRL+, 508 MW for TRL-). It's important to note that this approach currently does not account for the additional competition from TRE free bids, which would reduce the activation likelihood for the EV fleet. As such, the resulting TRE revenues should be considered a best-case estimate for an average bidding strategy.

Economic assessment of TRL- with PE pool fleet

Figure 47 summarises the hourly trends of the four key inputs used for the economic assessment (all data refer to working days in 2025, averaged by 4-hour block):

- Power made available by the PE pool fleet for negative control (blue area, left y-axis); note that the fleet consists of 12 EVs.
- Activation frequency of TRE- (red stars, left y-axis).
- TRL- capacity price (green line, right y-axis).
- TRE- activation price (yellow line, right y-axis).

The figure illustrates how well the natural fleet behaviour aligns with the typical price and demand patterns of TRL- and TRE-. While most of the available power falls during nighttime – when ancillary service demand is low – the fleet still captures a significant portion of the high-value period between 12:00 and 20:00.




Figure 47. Hourly profiles of all factors determining the profitability of the PE pool fleet to provide TRL-/TRE-: kW of available negative power (blue area, left y-axis), % of activation time (red stars, left y-axis), capacity price (green line, right y-axis), and energy price (yellow line, right y-axis).

Finally, Figure 48 presents the revenue potential per time block, broken down into TRL- and TRE- earnings. Capacity provision ensures steady income throughout the day, totalling around 11 CHF/day. Energy activations contribute a notable boost in midday, adding approximately 7 CHF/day. Altogether, the fleet could generate a gross daily income of 18 CHF for the ASP on an average working day. Assuming 95% success in TRL auctions and completely excluding weekends, this results in an annual income of 4'400 CHF or about 370 CHF/EV.




Figure 48. Gross earnings for the ASP managing the PE fleet (12 EVs), by time block and source of income (TRL-/TRE- in 2025).



General economic assessment of ancillary services

The above analysis focused on the case where the PE pool fleet is optimised to offer TRL- and TRE-. However, other combinations of ancillary services and fleet types also merit evaluation. Figure 49 presents the total earnings per EV for various scenarios, using the same methodology as before (all cases are based on 2025 working days only):

- PE TRL-: This is the baseline case, summarizing the results from the previous section.
- PE TRL- (excluding top 1% TRE- prices): This scenario tests the impact of removing the most extreme TRE- price spikes, which were observed to be rare but highly skewed. The bar shows that even if the ASP misses the most profitable 1% of activations, earnings remain robust, indicating low dependency on outliers.
- PE TRL- (weekly auctions): Here, we assume TRL- is awarded via weekly tenders (which allocate ~458 MW), rather than daily ones (~50 MW). This reflects a setup where the ASP can pool together various flexible assets to meet the full availability requirement of one week. However, TRL- weekly auction prices are significantly lower than daily ones, which reduces profitability for the EV fleet.
- PE SRL-: If the ASP can aggregate multiple flexible assets, it may also participate in SRL- and SRE- markets. SRL- capacity payments are generally more attractive than TRL- (weekly tenders), boosting total income. However, the SRE- activation frequency is low, which limits the energy-related earnings.
- QP TRL+: This case considers positive control provision using the Quickpac (QP) fleet. Due to its intense driving schedule, the fleet charging load can serve as a stable load for positive “stop charging” control. Nevertheless, TRL+ power availability from the QP fleet is limited to 16:00–24:00, resulting in modest revenue. Although the higher TRE+ activation frequency during evening hours adds some value, the overall outcome confirms that EV fleets are poorly suited for providing positive control.




Figure 49. Gross earnings from different use cases of ancillary service provision via dynamic charging of EV fleets.

- EV fleets cannot provide continuous control power over extended time windows: the maximum reasonable duration is around four hours. This aligns well with the current daily tenders



for tertiary control and with the expected transition of secondary control to a similar structure.

- Without V2X capabilities, the typically low energy depletion of EV batteries makes participation in negative control (holding charge and starting upon request) clearly more practical and profitable than positive control (charging and stopping upon request).
- Primeo Energie's fleet based in Münchenstein could generate up to 370 CHF/EV a year for the managing pooler when offering negative tertiary control through daily tenders.
- Quickpac's delivery fleet based in Hägendorf had shorter idle periods, which limit its potential for negative control; the deeper battery depletion during operation does not sufficiently compensate by improving positive control opportunities, resulting in an overall low economic potential for ancillary service provision.

3.2.4. Voltage stabilisation

Hägendorf

In the voltage stability use case the charge control would use a similar algorithm compared to the live peak shaving algorithm. Instead of cutting down the charge power when the power consumption is overshooting a threshold value, the voltage control algorithm would reduce or even drop the load when a certain voltage would be undershot. Again as in the peak shaving case the threshold would have to be adapted and set, in order to have the EV fleet charge in the times of high enough voltage. Compared to the local peak shaving the effect of a control reaction was not well observable. However still small corrections were visible and showing that an effect could be attained. The benefit of this control method is that it does not only include local consumption effects but also intermediate grid effects. If not in the same weight, but still along a broader optimisation target.

When speaking about voltage control the concept of actually having the voltage being outside the norm voltage boundaries ($\pm 10\%$) and therefore corrected would not be viable in both locations. The voltage would always show very safe values that would not be in danger of exceeding any limits given by the regulations.

Due to the weak focus and small expectations we only ran the case in the Hägendorf locations. Further the transformer setup in Münchenstein was quite generous so that the effect on the voltage would be minimal.

- Next to the difficulty of accounting the charging actions taken in this use case, the general impact was not effectively visible and it remained unclear whether the actions would be helpful also outside the local setup.
- The assumption that a voltage stabilisation would act towards congestion relief inside the grid had to be challenged since the correlation seems not to be given.

3.2.5. Local PV consumption

Münchenstein: Static time-based charge (Algorithm 1)

The first tests conducted in Münchenstein involved basic static charge controls based on preset times of day. These tests aimed to assess EV charging response in a market structure where the low and high tariffs, as well as the feed-in tariff, reflect the peak in PV generation around midday. Both DSOs and utilities have incentives to encourage consumers to maximise own-consumption of PV-generated elec-



tricity during these hours. Here, we present results for a static approach that prioritises EV charging from 13:00 to 16:00. Following this, we examine results from a dynamic approach that adjusts charging based on real-time power output from the PV installation.

Figure 50 illustrates the mechanism of this static charge control and its impact on select EV metrics. In this setup, EVs returning from their trips are plugged in at 11:30 and 12:30, but their charging is paused until 13:00, i.e. the beginning of the "PV charging window." At this point, charging is activated, allowing current to flow to the EV battery, gradually increasing its SoC. Given the relatively high initial SoC of the two EVs – typical for the Münchenstein fleet – charging is completed before the window closes at 16:00. In cases where charging did not complete within this window, the process was suspended until the next scheduled booking, which triggered a final charging session to ensure the SoC reached 100% as needed.




Figure 50. Example of static charge control (algorithm 1) on two EVs before and after the activation of the "PV charging window".

The impact of this use case on relevant metrics is presented in the next section together with the results from the dynamic case.

Münchenstein: Dynamic PV threshold charge (Algorithm 2)

In this dynamic setup, a power output threshold from the PV plant was established to control the charging of selected EVs. This test took place in September, when PV generation had already decreased relative to summer levels, prompting the choice for a conservative threshold of 6 kW (for a system with a peak output of around 20 kW).

Figure 51 shows the simple logic of this approach: EVs returning from trips with non-critical SoC levels delay charging until the live-monitored PV output exceeds the 6 kW threshold. When this threshold is met, charging begins; if output drops below 6 kW, charging is paused again. This cycle continues until either the EV reaches full SoC or an upcoming booking requires the EV to charge, regardless of PV generation.

To minimise frequent on-off cycles, two buffer mechanisms were implemented: (1) PV output must consistently remain above or below the threshold for several consecutive minutes before initiating or



pausing charging; and (2) the threshold for starting charging is slightly above 6 kW, while the threshold to pause is slightly below 6 kW.




Figure 51. Example of threshold-based charge control (Algorithm 2), where three EVs start and stop the charging process depending on the power output of the PV plant.

Figure 52 presents the results for the static (Algorithm 1) and dynamic (Algorithm 2) solar use cases. To derive the results, we assumed the system consisted solely of the on-site PV plant and the metered loads from the EVs garage, hence excluding the remaining building's load. The number of EVs usually charging in the metered area was approximately 20, but during presented live weeks only 9 EVs were actively engaged in the charge control system.

This system would have a nominal PV generation capacity of 20 kW and a peak EV consumption exceeding 160 kW. However, due to the low usage intensity of the EV fleet, the PV plant generates 40%–70% of the energy consumed by the EVs. Therefore, this setup is considered appropriate for evaluating how effectively EVs can shift their loads to absorb a greater share of PV generation.

The figure shows:

- The weekly generation of PV energy (yellow bars, right y-axis): this can fluctuate significantly between weeks, which complicates the quantification of the benefits generated by the charge management system.
- The share of solar energy which is directly consumed by the EV fleet, hence “own-consumed” (blue bars, right y-axis): it doesn't contain all the energy consumed by the EVs, but only the amount that has been simultaneously generated and consumed within the same 5-minute interval.
- Own-consumption rate (red line, left y-axis): the ratio of own-consumed energy to the total PV generation.
- Autarky rate (orange line, left y-axis): the ratio of own-consumed energy to the total EV consumption.




Figure 52. Lines and left y-axis: weekly own-consumption and autarky rates. Bars and right y-axis: weekly solar generation and direct own-consumption with EVs. Live charge management weeks are highlighted in green.

The figure shows that the impact of seasonality is visible, as the decreasing solar generation in September and October naturally leads to higher own-consumption and lower autarky rates. In general, the impact of the charge control interventions is noticeable but not outstanding. The relatively small size of the controlled fleet means that random variations in fleet mobility and availability significantly affect the flexibility that the charge manager can access and exploit. Similarly, day-to-day variations in PV output also substantially influence how the system responds, especially in the dynamic case.

Nevertheless, weeks 37 and 39 exhibit a behaviour that deviates from the general trend, likely due to a combination of active charge control and lower-than-average PV generation. By shifting charging to times when verified PV generation exceeded 6 kW, the dynamic algorithm increased the own-consumption rate by 5–10%.

Münchenstein: Dynamic own-consumption charge (Algorithm 3)

In the summer of 2025 the last and most dynamic algorithm was tested (Algorithm 3). In this context the objective was exactly to maximise PV own-consumption, as it's typically the case in single households or ZEV setups. This is commonly pursued because the feed-in tariff paid by DSOs is usually lower than the price paid for electricity drawn from the grid. As mentioned above, the setup of Primeo Energie in Münchenstein is not ideal, since PV production is significantly smaller than the building load. Therefore, the same “fictitious” system described earlier was taken as a reference for the live use case: on one side, the entire PV generation; on the other, a consumption load consisting only of the EV charging demand in the garage. In this way, the PV and EV load profiles were comparable, and the charge management algorithm had sufficient “degrees of freedom” to shift charging loads and improve the own-consumption rate.

The control followed “Algorithm 3”, which compared the aggregated load of the garage meters with the PV generation and activated the charging processes only when PV production exceeded the garages’ load. Even when PV generation was insufficient, the charge manager ensured that vehicles were still charged in time for their next booking, disregarding PV availability if necessary.



The use case was run for four alternating weeks during the summer of 2025, with the “no management” weeks serving as a baseline to quantify the benefits of charge management. By 2025, the number of controlled EVs had increased to 12, but a few EVs not participating in the charge management system still contributed to the monitored garage load.

Figure 53 follows the same notations of the previous graph: the bar plots show the weekly energy produced by the PV plant and directly consumed by the EVs; the lines depict the resulting autarky and own-consumption rates.

While the own-consumption rate is heavily dependent on the total production, a net increase during the live weeks with charge control can be observed. The own-consumption is a metric that shows to be misleading especially for weeks with little production. Generally, an increase of own-consumed energy of around 80 kWh could be achieved with the controlled pool fleet.




Figure 53. Lines and left y-axis: weekly own-consumption and autarky rates. Bars and right y-axis: weekly solar generation and direct own-consumption with EVs. Live charge management weeks are highlighted in green.

As mentioned before this evaluation and therefore also these findings only consider a very specific time window of the year. Combined with the chosen PV and consumer setting this indicates a very specific extract of the PV own-consumption case. When looking at the total two buildings at the Primeo Energie site that were considered in this project, the general consumption is always higher than production, which indicates a perfect own-consumption case without the need for optimisation of the charging process. On the other hand if the production is much higher, than the consumption, the incentive is fully there to increase own, consumption. However in this case a shift of the charging times is not always needed, because at most of the time during the day the production is “high enough” in order to profit from the own produced PV energy.

When considering a shift of the charging times from evening or nighttime to solar times, this also increases the own-consumption. However this is then considered less dynamic charge optimisation but rather a static charge configuration to charge during the day. This is also achieved with static configurations e.g. based on time-of-use tariffs.

Therefore the “dynamic own-consumption” case does only stand out when there is a comparable amount of consumption (including the flexible assets) as well as production on a site. Furthermore the strong



seasonality makes it a use case with strongly varying conditions and therefore varying added value. The variations are documented in the digital-twin simulations.

When comparing the different use cases and algorithms it shows that the different ones share behaviour for some circumstances and angle. The “dynamic PV threshold” case is comparable to a “dynamic own-consumption” case with high PV production compared to the available (flexible) loads. In both cases as soon as there is significant PV generation the flexible loads are enabled. The “dynamic PV threshold” case can be seen as a “dynamic own-consumption” optimisation but at the district / municipality level rather than on the local level. There, it is assumed that the total PV production is anyway surpassing the loads in that district and therefore it compares to a high-production own-consumption case. At this broader level, a system would be required to access live PV generation data from local or regional producers, even those outside the organisation of the EV owners or operators. Deploying a data-sharing system for this purpose could provide valuable benefits for DSOs. Even more high level (e.g. regional), with a certain balancing effect, the PV generation pattern can be seen as recurring and a static optimisation can also render desired effects. Here an intermediate path can be chosen with dynamic tariffs, where day-ahead weather predictions are considered and a “quasi-static” load distribution can effectively show the desired effect.

- Solar optimisations are subject to high fluctuations between days and weeks. This makes an evaluation over single weeks difficult.
- A dedicated own-consumption algorithm is simpler also for accounting reason, however limits the potentially participating fleets to companies owning a PV plant.
- Depending on the fleet operational pattern a visible, albeit no major, improvement of typical PV-related metrics could be achieved: in Münchenstein, the dynamic charge control “Algorithm 3” increased own-consumption and autarky rates by 10%, shifting approximately 80 kWh/week from grid feed-in to direct own-consumption.

3.2.6. Dynamic tariff optimisation

As a further use case, dynamic tariff based optimisation was considered and investigated. Even though dynamic tariffs are now often used to do a sort of DSO peak shaving by emulating a congestion price, it is a more general process approach that this project still wanted to investigate. Especially because a direct economic benefit can be calculated from the provided flexibility and the underlying charging pattern.

Dynamic tariff characteristics

Next to the time-varying wholesale day-ahead market prices for energy (published by the ENTSO-E association) more and more DSOs are following flexible tariff options for their grid tariffs imposed to their customer. While the involved DSOs and energy suppliers of the project partners did not have a dynamic tariff active during the time of the project, the Swiss DSO Groupe E had already set up their VARIO tariff, that they would compute and publish in a daily cycle and a granularity of 15 minutes. It being a grid tariff, it reflects the main average load in a suburban – rural grid part and therefore comparable to the high-level setup used in this project. While in the lowest grid level the effect are very local and therefore also varying, clearer trends are found in higher levels that lead to an averaged view of the grid that condenses in a single DSO tariff.

In order to rate the potential for optimising costs with the available charge flexibilities, the price distributions and variances were compared in the following figure with the availability times and flexibility potential of the different fleets.




Figure 54. Typical daily profiles and spread of: 1. SwissIX sport market price, 2. VARIO grid dynamic tariff, 3. Quickpac fleet charging load, 4. Primeo Energie fleet charging load. Both fleets are modelled with 50 EVs and their maximum load shift potential is shown in green.

The main visible trend for both tariffs, is a cheap price or opportunity during midday due to PV production and high prices during morning and evening peaks. In both tariff structures a clear difference is observable throughout the seasons but also inside the season there is some variance that is mainly due to changing availability of renewable energies. This effect is larger for the energy prices while they seem to follow a generally flatter curve. The grid tariff seems to have stronger daily effects that actually do not vary much in between the days. With those tendencies in mind it can generally be stated that the highest benefits can be obtained by moving loads from morning or evening to midday or afternoon. Also the night shows reduced prices, however with a smaller difference in price.

Comparing these trends to the charge curves and availabilities of the fleets it shows that generally the Primeo Energie fleet already aligns rather well with the low-price times. And mainly a shift from morning to afternoon seems effective. The Quickpac fleet however seems to hit the high price evening peak and therefore a shift can have a significant effect on the charging price. The only downside there, is that only



the night-time opportunity can be used since the cars are not available for charging during the very favourable midday times.

Münchenstein with Groupe E's VARIO grid tariff

With this use case applied to the Primeo Energie fleet in Münchenstein two factors were dominant. On the one side a large portion of the fleet would normally charge throughout the day, including afternoon, as well as the little energy charged in average by the users for doing mainly only small trips. Both effects diminish significantly the resulting optimisation potential since the reference is already relatively well matching the price incentive. The table below shows the effective benefits that were able to be reached with live tariff optimised charging during a week with the 12 fleet EVs of Primeo Energie's pool fleet. Wednesday and Thursday stand out, which had only little fleet activity and most of the used vehicles would charge around noon .

Example Week: Savings for price-optimised charging of the 12 fleet EVs per day

Weekday	Monday	Tuesday	Wednesday	Thursday	Friday
Estimated flexibility savings	1.60 CHF	1.90 CHF	0.30 CHF	0.10 CHF	1.65 CHF
# of postponed charging sessions	4	4	3	2	8

In the graphic below one can see that the timing of the car's arrival is essential for the respective saving's potential. On vehicles arriving late morning the charge optimisation has a negligible impact since they would already charge in low-price time. The most impact can be seen from cars arriving early morning. This does happen at overnight bookings, however those cars are often still booked sometime in the morning and cannot wait until the low-price time after noon. Most of the cars go out on a trip in the morning and arrive before noon or in the early afternoon. In that case no big saving can be achieved since the charging time is already favourable. Only for afternoon arrivals it can get interesting but only for late arrivals and there are not so many of those with the Primeo Energie fleet. All these effects clearly diminish the overall savings potential.



- A data-driven advanced charging system can make use of dynamic tariffs very efficiently to manage EV fleet flexibility. This is especially true for varying tariffs so that the fleets can easily adapt to those updates.
- The investigated dynamic tariff showed a limited day-to-day variation, so that a static charge schedule would be able to have a comparable effect.
- The fleet schedule and the price spread highly dictate the possible gains that can be achieved through tariff optimisation; the tested project fleet in particular displayed limited potential.

3.3 Economic evaluation of live use cases

3.3.1. Cost references

Tariffs:

For the economic evaluation of the live use cases, some reference tariff values were needed. While there is quite a variance depending on the DSOs and setups, an average value still depicts a good baseline for the different tariff values that were needed for the evaluation. The following table lists the chosen values that are mainly an average over the DSO tariffs for large commercial consumers (C4) including the project partners. (Rounded to whole numbers.)

	Energy & grid working tariff	Solar infeed tariff	Low-voltage peak power tariff	Medium-voltage peak power tariff
Average tariff	23 Rp. / kWh	13 Rp. / kWh	9 CHF / kW / month	8 CHF / kW / month

Charge control:

Comparing different techniques including the VEIU Tech charge control system an average cost of *6.70 CHF per EV per month* could be estimated for remove vehicle data access and charge control. Even though not all manufacturers support the same level of interface and also the tariff system varies a lot between car manufacturers the average feels reasonable and also matches the costs of VEIU Tech's charge management system that works independently of manufacturers.

Flexibility pooler:

The revenue share taken by a flexibility pooler for the inclusion into a VPP was estimated at 25%. This is a reference that we took from an effective setup with a project partner.

3.3.2. Cost comparison

This section compares the calculated revenues and revenues splits for the specific live use cases that were executed live. The cases are only valid for the specific situations that included a specific time of the year when the live control was executed and the respective technical and operational setup. In tendency a good setup was chosen, so the assumption can be made that the revenues are slightly above the yearly average.

A more general calculation was made with the digital twin, where a more generic setup could be chosen and also seasonal effects could be integrated in the computation. For those results please look further down in the next sub-chapter.



In the following results we assumed that a charge flexibility provider would take 50% share of the resulting savings or revenue generated with the fleet flexibility. This would leave the other 50% to the fleet operator. This is an assumption that considers that the fleet operator and/or owner would still need to get a significant benefit in order to subscribe for such a scheme. This is especially the case since there is a non-negligible impact on the fleet operation and a usage adaptation need.




Figure 56. Financial breakdown of revenue for the different live use cases in the two project locations and the two project fleets.

The overview shows quite a distributed range of revenue splits per case and also per location / fleet. Also the effective impact of the charge control costs is well visible and how it challenges the different use cases. (*) The dynamic tariff and ancillary service in Hägendorf use cases were calculated based on the standard driving pattern observed in the live use cases and without live control.

The fact that the revenue share is held at 50% has the effect that not all cases show the same benefit for the fleet operator. This makes sense since the system operator is dependent on the cases' revenue and therefore a percentage split creates an incentive for the fleet operator too, in order to act in accordance to the optimising charge control. The resulting gain for fleet operators would in total be higher than the effective profit of the system operators since they need to subtract their own costs. However it is unclear whether the fleet operator saving values are relevant enough compared with the operating costs of a fleet vehicle that are estimated at around 700–1'000 CHF/month.¹⁰

¹⁰ [TCS Testberichte-Ratgeber](#)



- The cost setup with a fleet operator, a needed system operator and control access costs divide the reached flexibility revenue significantly.
- Generally, the revenues achieved in the live use cases were modest to small.
- Compared to the operational costs of fleet vehicles, the potential savings are relatively small such that it is not clear whether an adoption of such a system would be considered.

3.3.3. Fleet implications

A clear and impactful difference is visible between the two fleets of Quickpac and Primeo Energie. Especially in the grid cases of local and DSO peak shaving, as well as for the dynamic price optimisation the Primeo Energie's pool fleet in Münchenstein shows a much weaker case in terms of total savings. This is effectively the case since the vehicles are scarcely used and especially because they drive much shorter distances. This is due to the nature of the usage for commuting to a work site and not to do effective transport work. The small distances lead to a smaller energy turnover which leads to shorter charging sessions. This on the one hand reduces the effective energy costs involved as well as the possibility for charging sessions to overlap and create power peaks.

On the other side the scarce usage of the cars leaves a high idle time of the vehicles which is beneficiary to the ancillary services case. This is also visible in the better economic performance of the case compared to the other project fleet in Hägendorf. As discussed in the results of the ancillary services' case, some other downsides play in in that case with usage pattern, and technical limitations bringing limits to the effective potential.

The Quickpac delivery fleet in Hägendorf shows a much bigger energy turnover and denser charging patterns which lead to a bigger savings potential. The limiting factor there is the relatively short time availability of the cars that limit the flexibility time windows and therefore the potential savings.

- The possible flexibility revenues highly depend on the energy revenue used by the fleets.
- The alignment of the fleet charging process with flexibility needs are impactful.

3.3.4. Use case-specific considerations

Generally it can be stated that the peak shaving cases look more promising than other use cases. The dynamic tariff use case can be seen as a modified peak shaving case that due to its setup and price structure seem to not bring the same benefit to the end customer. But conditions can change in the future.

While in general the opportunity costs of not being aligned with e.g. energy tariffs (time-of-use or even dynamic in the future) is not considered. However in the DSO peak shaving case also the savings from a delayed grid enforcement is not considered. Only the tariff paid to higher level grid infrastructure is considered in the project calculations.

The ancillary services case would be the best performing case however the needed inclusion of a third party pooler diminishes the remaining potential revenue. Furthermore the effective active operation is questionable because of the unreliable availability and the small power size of a fleet, in order to go for an aligned optimal bidding strategy.

One way this could be relieved would be through a DER flexibility platform where smaller assets can directly bid their specific flexibility. However smaller compensations must be expected since the competition would also significantly increase.



- Flexibility products seem to not be aligned with the flexibility possibilities of EV fleets yet.
- Every case has its own speciality that hinder the effective integration of fleets.

3.4 Digital twin simulations

Figure 57 illustrates a representative output from the simulation platform. In this example, the objective function was standard parking peak load minimisation, meaning there were no other loads besides the simulated EV fleet. A one-week simulation involving 10 EVs was conducted, but for clarity, the plot only displays output data for the last three EVs (EV 7 – EV 9) .




Figure 57. Example of output of the digital twin: for each EV we obtain SoC profile (orange lines) and net battery power (blue lines). The last red line indicates the sum of all EV charging loads. The grey shaded areas represent the charging stops. Only three individual EVs are shown for clarity.



For each EV, time series data of the battery SoC and net charging/discharging power were generated. The final red plot presents the aggregated charging load of all 10 EVs, allowing us to observe how the optimiser aimed to minimise the overall peak load. The EVs were allowed to charge at rates between 4 kW and 10 kW, and the final observed peak load was 5.6 kW.

In this example, we can particularly appreciate the interplay between the charging demand and the "minimum desired SoC before a new trip" constraint, expressed as:

$$e_{i,t} \geq \min(E_{\text{feas } i,j}, E_i \cdot f) \text{ for each } t = t_{\text{dep } i,j}$$

Here, f was set to 0.9, and the role of this constraint is visualised in Figure 57 by the purple markers, which represent the energy content thresholds that each EV battery must exceed before departure. Particularly critical is the first charging stop of EV 7 (shown in the first plot), where, to meet the specified minimum battery energy threshold $E_{\text{feas } i,j}$, the EV must charge at 5.6 kW. This charging requirement ultimately determines the overall peak load of the fleet.

In a perfect-foresight optimisation, this situation would not have occurred because the next trip was evidently short and the EV could have completed it without needing to charge beforehand. However, this example more closely emulates real-world conditions, where EV drivers may prefer to sufficiently charge their vehicles whenever possible, even if it is not strictly necessary. This behaviour reflects a conservative approach to charging, ensuring that vehicles are adequately prepared for unforeseen needs, and it highlights the importance of accounting for driver preferences in optimisation models.

In all subsequent "optimiser" results f has been set to 0.9 or 0.95.

The following table provides a high-level overview of all setups and use cases explored in the simulations. The first column specifies the "contexts" assessed in the use cases, which primarily correspond to the type of uncontrollable "baseload" contributing to overall grid usage alongside the EV fleets. The next columns indicate which fleets were simulated in each use case. Not all possible baseload–fleet combinations were tested, as some would have not provided meaningful insights and an excessive number of simulations would have polluted the narrative and potentially confused the reader.

For each use case the table also indicates the type of solver used to derive the "load management" result, i.e. when either smart charging or battery management was applied to improve the objective function. A single "X" denotes simulations where only the MILP optimiser was used, while a double "XX" indicates that heuristic algorithms were also employed. These algorithms were taken directly from the live use cases and adapted to work in the digital twin environment.

Context / Baseload	Fleet type		Flexibility use case	Section
	Primeo Energie pool fleet	Quickpac delivery fleet		
No baseload	XX	XX	Fleet load peak shaving	3.4.1
		XX	Local (building) peak shaving	3.4.2
Münchenstein office building	XX	XX		
L5 grid in Hägendorf		X	DSO grid peak shaving	3.4.3
L5 grid in Münchenstein		X		
Only PV generation		X	Local PV consumption	3.4.4
Münchenstein office building + PV generation	X			

X = Only optimiser simulations

XX = Both optimiser and heuristic (i.e. based on live use case algorithms) simulations



3.4.1. Fleet load peak shaving

Before using the simulation environment to investigate the setups encountered in the live use cases, this first set of simulations considers the charging load of the fleets in their isolation. This also allows to generate more punctual insights that will be useful to understand also the later simulations.

Figure 58 shows the charging peak load for the QP fleet with a varying number of cars, where every EV has its dedicated EVSE at 11 kW. We tested different charging logics and to make results comparable, each peak is normalized in kW/EV.

The “charge immediately” approach shows the peaks obtained when no control is implemented: it thus provides insights on the natural synchronicity of the charging behaviour of this fleet. The values remain very high even for large fleets: for 50 EVs the average charging peak is 8.5 kW, which corresponds to a simultaneity factor of 77% and a total peak load of 425 kW. This is not far from a plausible scenario at the Högendorf site, where about 40 EVs used to be operated (albeit not all in the same way of the Mercedes eVito). For reference the recommended simultaneity factor by SIA 2060 for 11 kW chargers at commercial sites is 23%.




Figure 58. Normalized peaks of QP fleet charging load (kW/EV) for different fleet sizes and charging strategies. Charging rate of a single EVSE is 11 kW and simulation window is one week.

Applying charge control strategies visibly reduces the peaks: a plain load management system with no access to SoC information (orange line, “Algorithm 0”) can reduce the peaks by 40% for any medium-large fleet size. Adding SoC information significantly boosts the results (green line, “Algorithm 2.2”), achieving reductions of 64%; in absolute terms this translates to a reduction of 5.4 kW/EV or 268 kW for 50 EVs. The MILP optimisers can predictably achieve better performance, but interestingly the heuristic algorithm with SoC access can outperform an optimiser that lacks this information.

It’s worth noting that the charge management algorithms cannot further reduce the peak load for fleets larger than 20 EVs. This is not only due to the synchronous fleet behaviour, but also to the relatively constrained charging space available. The considerable amount of energy consumed in the long journeys of this fleet has to be supplied within the limited time window between late afternoon and early morning: this limits the number of EVs that can be charged sequentially, resulting very quickly in an almost linear increase of the peak load with the number of vehicles.



Figure 59 shows the same results for the PE fleet based in Münchenstein: in this case the EVSEs and EVs have different charging rates and the average fleet value is approximately 8.7 kW/EV.

All curves exhibit a very different behaviour from the QP fleet, with strong peak reductions with larger fleets. This is due to the very low synchronicity of PE fleet behaviour, characterized by sparse short journeys that require only brief charging sessions. For 50 EVs the “charge immediately” approach exhibits a simultaneity factor of 16%, which is even lower than the SIA 2060 guidelines.




Figure 59. Normalized peaks of PE fleet charging load (kW/EV) for different fleet sizes and charging strategies. Charging rates of the fleet vary between 6 and 11 kW; simulation window is one week.

Consequently, the margin of improvement for all charge management strategies is limited. Without SoC information, both the heuristic algorithm and the optimiser can only reduce the peak by 15% (for 50 EVs). When providing SoC information, the heuristics can unlock more meaningful peak reductions, up to 40%. Access to SoC information is particularly valuable for fleets such as the PE fleet: since their batteries are only minimally depleted, SoC data provides considerable flexibility in deciding when to start charging; without SoC information the management system must instead assume a worst-case scenario of 0% SoC. Nevertheless, the absolute peak reduction is still contained: for 50 EVs it results in about 0.7 kW/EV or 33 kW for the entire fleet.

- The simultaneity factor of uncontrolled charging can vary significantly depending on the fleet's utilisation intensity and timing. This affects not only the sizing of charging and grid infrastructure, but also the potential benefits of implementing smart charging.
- The high simultaneity of the QP fleet makes any charge management scheme highly effective in reducing peaks: dynamic charge control with access to SoC information can lower the peak of 20–50 EV fleet by up to 64%.
- The very light utilisation of the PE fleet results in small simultaneity factors and low uncontrolled peaks. Smart charging strategies can achieve further reductions, but the net benefits per EV tend to be minimal.



3.4.2. Local (building) peak shaving

In this section we assess the peak load reduction potential of different load management strategies in a local setup, where the objective is to reduce the monthly peak load measured at the grid connection point¹¹. Specifically we assess the following three cases:

- Main building of the Münchenstein site, whose loads come primarily from office spaces and the garage of the PE pool fleet
- Depot of the Hägendorf site, where the entire load comes virtually from the QP fleet (hence we neglect non-EV loads)
- A fictitious combination of the building load from Münchenstein with the charging load from the QP fleet

The flexible technologies considered to reduce the peak are:

- Only the EVs in the fleets via charge management systems
- Only stationary batteries of different sizes with a C-rate of 0.5 (EVs charge “immediately”)
- A combination of both

In each case, the employed energy management algorithms are:

- Optimal control: it's the result from the MILP optimisation process, which has full vision of the future building load and the fleet driving schedule. The SoC of all EVs and the battery is also always accessible. Importantly, we however limit the foresight power of the optimiser through the “Minimum desired SoC before new trip” constraint, which forces the optimiser to charge EVs in a very conservative way regardless of the length of subsequent journeys.
- Load-capping (Algorithm 0): standard load-shedding mechanism that lets EVs charge freely until a set maximum load is reached. In the case of a building baseload, the sum of all loads is checked against the threshold. This method has no view of future loads and has no access to the EVs' SoC. This approach is applied to batteries, as it is unreasonable to manage a battery without knowing its SoC.
- Dynamic control (Algorithm 2.3): also in this case the energy management system is assigned a target maximum load for the combined demand of the building and its assets. It then adjusts the load of EVs and the battery in response to the live-monitored aggregated load. While future loads remain unknown, the SoC of EVs is fully available to optimise the charging process.

The load-capping and dynamic control methods are also referred to as heuristic approaches, given their nature in comparison to the optimal solver. The target thresholds of these heuristic methods are set automatically in an adaptive manner: each setup is typically simulated multiple times to account for seasonal variability and the stochasticity of driving patterns. The peaks observed in previous simulated weeks are then used as the basis for determining the thresholds of subsequent weeks.

Additional dimensions explored include fleet size (10–50 EVs) and seasonal focus (summer vs. winter), as the latter affects the building's baseload. The main results reported correspond to a fleet of 50 EVs and represent annual averages, while the appendix provides additional solutions for different fleet sizes.

¹¹ All simulations were performed over a one-week time window with a 15-minute resolution. Both the heuristic and optimal solvers therefore aim to minimise the weekly peak load, which differs from the objective of reducing the monthly peak. This choice was motivated by 1) computational efficiency and 2) the observation that most relevant dynamics and patterns occur within daily or weekly cycles, while different weeks within a month tend to behave very similarly. Nonetheless, monthly simulations have also been tested: although the absolute peak values tend to be slightly higher (since a monthly peak corresponds to the maximum of four or five weekly peaks) the relative improvements achieved by the different algorithms and load management strategies remained virtually identical. The difference between the “weekly” and “monthly” results was well below the variability introduced by seasonality, cost uncertainty and algorithmic factors (e.g. the specific heuristic or threshold-setting criteria). Consequently, in all subsequent analyses the reported peaks are derived from weekly simulations but are treated as monthly peaks in the economic assessments.



Beyond the technical potential, this section also provides estimates of the economic savings that these solutions can bring. These estimates are based on the cost assumptions listed in the following table.

Peak shaving objective		
Peak load price	9 CHF/kW/month	Average of LV peak price for commercial users

EV charge management		
Load-capping (Algorithm 0)	0 CHF/EV/month	Most charging systems come with this functionality pre-installed
Dynamic control (Algorithm 2.3): material and installation	10 CHF/fleet/month + 8 CHF/EV/month	Based on cost estimate by VEIU Tech and average subscription price to access the vehicles' SoC. Uncertainty modelled with triangular PDF: 6-8-10 CHF/EV/month

Battery energy storage system		
Material costs	235 CHF/kWh	Cost estimate based on Primeo Energie's prior projects. Uncertainty modelled with triangular PDF: 170-235-300 CHF/kWh
Installation costs	Same as material costs	Based on Primeo Energie's prior projects.
Lifetime	15 years	Based on Primeo Energie's prior projects.
Resulting monthly cost	2.6 CHF/kWh/month	No discounting. Resulting uncertainty is a triangular PDF: 1.9-2.6-3.3 CHF/kWh/month

As indicated, we also modelled the uncertainty of a couple of key costs. While this does not affect the average results or the main insights of the analyses, it provides a visual representation of the uncertainty in the modelling and informs the reader about the significance of the conclusions.

Münchenstein site with PE fleet

Figure 60 presents the average peak loads obtained when simulating the operation of 50 EVs sampled from the PE fleet at the Münchenstein site. The standard deviation is shown with error bars only for the highest bar in each column.

The blue bars show the peaks obtained without applying any load management strategy: the darker shade is the peak before adding the fleet load while the lighter shade represents the additional contributions from EVs. The large standard deviation emphasizes that in this setup the main driver for the peak is the building baseload: this was observed to be true especially in winter, with periods of intense heat pump activity alternating with weeks of no operation. The relative size of the two terms also highlights that the impact of pure smart charging strategies will be rather contained, since it can only address the light blue portion.

This is visible with the violet bars, which display the peaks resulting from applying the three smart charging algorithms to the PE fleet: the optimiser is able to entirely remove the additional contribution of EVs, basically "hiding" the entire charging load into the valleys of the building baseload. The heuristic algorithms do not manage to perform that well, achieving only minor reductions of the average peak load.




Figure 60. Average peak loads obtained for the Münchenstein site with 50 EVs from the PE fleet for different technological setups (x-axis) and algorithmic strategies (colour shades).

The yellow-red bars show the results when the EV management is supported by a stationary battery. The fact that batteries can also discharge the energy stored allows them to absorb even the peaks originating directly from the building baseload, hence leading to further reductions than the dark blue shade. The figure shows that an optimal management of EVs and batteries (red bars) can achieve this effect already with a 20 kWh battery. With heuristic algorithms (yellow-orange bars) the required size increases to 50–100 kWh. Predictably, the larger the battery the stronger the benefits: the cost assessment in the next section will look for the most cost-effective solution.

On average, the heuristic solvers achieve 60–65% of the peak reduction obtained with the optimal solver (excluding the “load-capping” cases with no or minimal battery capacity). The two heuristics for EV charge management sometimes display comparable peaks for this combined setup: this stems mostly from the difficulty in defining the thresholds used to guide the control logic, especially when the main driver of the peak is the building baseload with its large variability. The interaction between the battery and EVs is also complex to model and the weak performance of the heuristic approaches suggests that more advanced algorithms should be explored and tested: this was not a priority in the current project, since the combined setup was only simulated and had no real live-twin equivalent.

Finally, the green bars show the peaks resulting from using only the battery for load management, while allowing the EVs to charge immediately. As expected, this leads to similar results to the combined case, emphasizing that the additional contribution of EV charge management in this setup is minimal. Also in this case, the heuristic battery management achieves about 60–65% of the theoretical peak reduction obtained by the optimal solver. These results are consistent with the findings from the OrtsNetz pilot project, where the battery energy management algorithm achieved 55.6% of the theoretical ideal reduction under perfect foresight conditions.¹²

While the peak load reductions achieved by large battery systems are significant, the natural question is whether this solution is economically sensible.

Figure 61 compares the costs to implement the load management solution (x-axis) with the benefits in terms of monthly peak load reduction (y-axis). Only the results from the heuristic simulations are shown here for consistency, while the appendix provides the solutions with the optimiser. With the data presen-

¹² OrtsNetz P+D Project | ARAMIS



ted in this format, the net savings are represented by the diagonal isolines: the first black line shows the situation with no load management system (case with blue bars above) and each subsequent line represents additional 400 CHF/month of savings.




Figure 61. Monthly savings from peak load reduction (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein site with 50 PE EVs. Different colours indicate distinct technological or control setups (only heuristic management is shown). Small points represent results from individual stochastic simulations, while the rhomboids the average per setup. The rhomboids with the same colour depict the same setup with varying battery sizes (see Figure 60 for reference).

The colours follow the palette of the previous plot: the violet points represent the cases with only EV charge management and confirm the limited peak load savings observed above. The plain load-capping approach comes at no costs, so it still leads to positive net savings of about 100 CHF/month; the advance dynamic charging system however costs about 400 CHF/month, thus it is not profitable.

Combining the load-capping approach with a battery (yellow points) leads to the most profitable outcomes: adding a 100 kWh battery generates average savings of 200 CHF/month, with a few individual months potentially netting 600 CHF. Adding a battery to the dynamic charging system instead never achieves break-even.

Batteries on their own (green points) also do not perform particularly well. Somewhat surprisingly, the narrow additional contribution of EV load-capping is enough to shift the battery-only case from barely breaking even to achieving clear net savings.

These net savings should, however, be viewed in context. Considering that we are assessing a large office building with a garage of 50 EVs, even the best-performing case – load-capping combined with a 100 kWh battery – yields only about 4 CHF/EV/month. While the variability suggests that this profit is consistent and that it was achieved without a fine-tuned heuristic algorithm, the additional effort of setting up a full load management system may still be enough to discourage most commercial entities from pursuing it.



Hägendorf site with QP fleet

This section evaluates the impact of the same load management strategies at the Hägendorf site with 50 EVs sampled from the QP fleet. Since the baseload of the depot building is virtually negligible, we simplified the analysis by removing it entirely and focusing on a scenario where the load consists solely of the charging fleet. These results therefore correspond to the setup shown in Figure 58, but with a more detailed analysis and an accompanying economic assessment. The following subsection then considers the case where the QP fleet is located at a site with a meaningful baseload, such as the office building in Münchenstein.

Figure 62 displays the peak loads for the Hägendorf case with the QP fleet. In the no-management scenario (blue bar), the peak load is entirely due to the charging fleet and is visibly higher than the peak of the PE fleet with the same number of vehicles. It is therefore not surprising that the charge-management strategies (violet bars) achieve such substantial reductions, lowering peaks by 50–70% depending on the control algorithm.




Figure 62. Average peak loads obtained for the Hägendorf site with 50 EVs from the QP fleet for different technological setups (x-axis) and algorithmic strategies (colour shades).

Adding a battery to the charge management system results in only minor additional improvements (orange bars). This is because the battery takes over the load-management task once the EV charge management has already substantially flattened the curve. As a result, each additional kW of reduction requires the battery to shift multiple kWh and the benefits therefore scale only slowly with battery size.

Similar considerations apply to the battery-only management case (green bars). The difference from the previous case is that here the battery addresses the unmanaged fleet's natural charging profile, which is more synchronous and therefore peakier. As a result, batteries can deliver more kW reductions for the same kWh capacity, although they start from a much higher peak load. Large batteries can thus reduce the peak by 30–40%, but they still fall short of the reductions achieved by even the most basic load-capping charge management strategy.

In general, the more intensive usage of the QP fleet compared to the PE fleet makes it harder for batteries to meaningfully reduce the peak load, since this would require shifting and cycling a large amount of energy. In such cases, it is more effective to improve the load “at the source,” i.e. to directly flatten the fleet's charging profile before it develops into a bulky peak. By contrast, the leaner PE fleet – with its lower energy demand – even when unmanaged produces a less problematic charging load, both for the overall system and for any batteries that might be deployed to mitigate it.



Figure 63 displays the costs and savings for all the setups in this use case. As expected, the sole EV charge management strategies perform best, achieving solid 1500 CHF/month (or 30 CHF/EV/month) of net savings. The combined cases with a battery perform in a comparable way only for small sizes, but they get gradually worse for larger capacities.

As noted earlier, relying solely on batteries proves largely ineffective and yields minimal benefits.




Figure 63. Monthly savings from peak load reduction (y-axis) vs. monthly costs for load management solution (x-axis) for the Hägendorf site with 50 QP EVs. Different colours indicate distinct technological or control setups (only heuristic management is shown). Small points represent results from individual stochastic simulations, while the rhomboids the average per setup. The rhomboids with the same colour depict the same setup with varying battery sizes (see Figure 62 for reference).

In this setup, it is easier to compare the performance of load-capping and dynamic control strategies. The additional cost reductions achieved by SoC-aware dynamic charging are almost perfectly offset by its modelled costs. This limited advantage is understandable, since the objective function being pursued – restricting peak load based on a threshold – is precisely what standard load-shedding systems are designed to handle. In other words, we are comparing approaches of different complexity in a use case tailored to the simpler one. While access to SoC information does provide some benefits, such as improved scheduling of charging queues, these appear insufficient to outweigh the added costs. It should also be noted that most QP EVs return from their trips with an SoC of 30–40%, so the starting conditions of the advanced charging system are not drastically different from those of a load-capping system, which effectively assumes 0% SoC.

Münchenstein site with QP fleet (mixed case)

Finally, we test a mixed setup where 50 EVs from the QP fleet are connected to the same local grid of an office building analogous to the one in Münchenstein. Figure 64 presented the peak load obtained in this scenario.



The results are qualitatively similar to those observed in the previous case with only the QP fleet. This is because – as shown by the first blue bars – the overall peak load is primarily driven by the fleet’s charging demand. The main impacts of this new setup are a larger standard deviation, caused by the variable building baseload, and the consequently less consistent performance of the heuristic algorithms, particularly in the combined case.

Overall, the EV charge management systems still perform very well, reducing the peak loads to levels comparable with those of the office building alone. The additional contribution of batteries to the fleet management systems is negligible; when deployed on their own, batteries provide some gains but still fall short of the savings achieved by even basic charge control systems.




Figure 64. Average peak loads obtained for the Münchenstein site with 50 EVs from the QP fleet for different technological setups (x-axis) and algorithmic strategies (colour shades).

Figure 65 translates these results into the familiar cost-savings plot. Also in this case, the trends are very similar to those observed with the QP fleet alone; even quantitatively, the net savings achievable with the charge management systems (in violet) remain in the range of 1'500 CHF/month (or about 30 CHF/EV/month).

The only noteworthy distinction is that the combination of EV load-capping control with a small battery appears to perform slightly better, with average net savings approaching 1'700 CHF/month. However, the variability displayed by the stochastic simulations does not allow for a conclusive statement on this point.

The clear takeaway is that for a large, intensively used fleet such as the QP case, charge management systems are economically advantageous, and that the addition of a small, well-controlled battery can also be considered.




Figure 65. Monthly savings from peak load reduction (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein site with 50 QP EVs. Different colours indicate distinct technological or control setups (only heuristic management is shown). Small points represent results from individual stochastic simulations, while the rhomboids the average per setup. The rhomboids with the same colour depict the same setup with varying battery sizes (see Figure 64 for reference).

In conclusion, it should be emphasized that this peak-shaving use case neglects other important benefits of batteries, such as the ability to absorb excess PV generation from a potential on-site plant and discharge it during the night. Similarly to the discussion on the two charge management systems, in this setup we are evaluating the value of batteries against an objective (peak shaving) that does not fully capture their overall capabilities. In Section 3.4.4 we will therefore investigate cost-minimisation scenarios in the presence of a local PV plant.

- The shape of the peak profile strongly influences which strategy is most effective for peak shaving. If peaks are thin and infrequent, batteries can reduce them efficiently while offering additional benefits. When peaks are instead “bulky”, batteries typically lack the capacity to do so cost-effectively. In such cases, it is more convenient to address the source of the peaks directly, for instance via smart charging solutions if the peaks stem from large uncontrolled charging loads (such as in the case of the Quickpac delivery fleet). A similar reasoning likely applies to other peak sources, such as heat-pump operation.
- While batteries are technically more effective at flattening load profiles in mixed-use contexts such as office buildings, their economic return is minimal. Unless the building owner has strong (monetized) incentives to reduce the peak, deploying a battery is only justifiable when it also serves other purposes, such as increasing PV self-consumption.



- The economic benefits of peak-load reduction through smart charging are significant for intensively used fleets like Quickpac's delivery fleet. Both load-shedding and SoC-aware smart charging enable reliable returns of around 30 CHF/EV/month.
- Dynamic charging with access to EV SoC data does not show a strong economic edge over simple load-shedding when the only objective is peak reduction: a natural outcome since load-shedding is specifically designed to prevent peak overshoots. As with batteries, the added complexity of SoC-aware dynamic charging is better justified when it supports more advanced flexibility services, such as participation in grid-support programs.

3.4.3. DSO grid peak shaving

In this simulation series, we adopt the DSO perspective and assess whether the ongoing electrification of fleets may pose risks to the grid, particularly with respect to transmission capacity. In the local setup, we assumed that a business operates a fleet of a given size according to its needs and then seeks to optimise its charging operation (in this case, by reducing the peak load). In the L5 grid use case, by contrast, the size and number of fleets contributing to the load lie outside the control of the DSO. For this reason, we adopt an “evolutionary” perspective, examining how the situation develops as the number of fleets in the grid increases and identifying the point at which it becomes critical.

To this end, we take the current load of two representative grid nodes and superimpose the potential additional load of multiple EV fleets. The grid nodes selected for this evaluation are the L5 feeders exiting the L4 substations located directly upstream of the project sites, namely Münchenstein and Hägendorf. We then apply the same load management strategies considered in the local setups, ranging from purely EV charge management to the deployment of large battery storage systems.

Before presenting the quantitative findings, Figure 66 shows a qualitative comparison of the load profiles of the fleets and the two L5 feeders. Both fleets consist of 50 EVs and the uncontrolled charging load (red) immediately highlights the more intensive usage of the QP fleet compared to the PE fleet. The load-shift potential (green) is derived from the total fleet charging power connected at any given time and represents the maximum load that could be shifted to that moment. Interestingly, for the QP fleet this potential drops to zero during the central hours of the day, since all EVs are typically out in operation. This is not the case for the PE fleet, as a share of vehicles is always parked in the garage.

When comparing with the grid loads, we observe that the QP peaks coincide unfavourably with the evening peak of the Hägendorf grid, especially in winter. At the same time, the largest valley of the Hägendorf profile occurs during midday hours (likely due to local PV generation), precisely when no QP vehicles are available for charging. Consequently, the only meaningful mitigation strategy is to shift charging sessions to the early morning hours (2–6 am).

A similar strategy would also be effective if the QP fleet were located at the Münchenstein site; however, in this case the initial situation is less critical, since the two uncontrolled peaks do not overlap.




Figure 66. Typical daily load profiles and spread of the L5 feeders used in the DSO grid assessment (top two plots) and of the two project fleets (each with 50 EVs, bottom two plots). For the fleets, the maximum load shift potential is shown in green.

By contrast, the uncontrolled load of the PE fleet peaks during the day, coinciding directly with the Münchenstein grid peak. A simple shift of charging loads to the night can already mitigate this stress on the grid. If the PE fleet were located in Hägendorf, a more nuanced strategy would be required: early-morning charging sessions should be avoided, midday sessions maximised, and the remaining charging load shifted to the early hours of the following morning (2–6 am).

In principle, all four cross-scenarios could be explored. However, from a DSO perspective, only the more critical evolutions merit closer investigation and early intervention. Figure 66, together with the results from the local peak-shaving use case, shows that the uncontrolled load of even a large PE fleet does not generate significant peaks. For this reason, in the DSO grid analysis we focus solely on the risks posed by the introduction of QP-like fleets.



Hägendorf L5 grid with multiple QP fleets

Figure 67 presents the average monthly peak loads resulting from gradually connecting additional QP fleets (50 EVs each) to the Hägendorf L5 feeder. The light blue bars show the evolution without any control strategy: the situation escalates quickly, and with 5 fleets we regularly observe peaks above the maximum load ever recorded in the available data series (dashed line), which we may take as a proxy for the feeder's operational limit.

The violet bars report the new peak loads when charging is controlled. We tested both heuristic and optimal approaches, however, in this L5 setting the design of heuristic logic and thresholds proved problematic, yielding results that were neither consistent nor particularly informative.¹³ We therefore present the optimiser results to enable a fair comparison across scenarios. While these may lean optimistic, the local peak-shaving use case indicated that a well-designed control strategy can approach the optimal outcome closely enough to keep the insights relevant.




Figure 67. Peak loads resulting from the gradual introduction of QP-like fleets (50 EVs each) in the Hägendorf L5 grid. Uncontrolled peaks are shown in light blue, while optimal charge management in violet. Only for the scenario with 5 fleets, the additional impact of batteries is explored: either alone (green) or combined with smart charging strategies (orange).

Since the addition of 5 uncontrolled fleets already becomes critical for the grid, Figure 67 extends this scenario by evaluating the impact of introducing utility-scale battery storage. Following the same colour scheme as before, the green bars show the peak-load reductions achieved when batteries are added without managing the fleet charging load. Installing 1 MWh of storage already reduces the likelihood of reaching the current operational limit, but at least 5 MWh are required to bring the system back to today's grid load levels.

¹³ A natural threshold choice for a heuristic is the current grid peak. However, this creates two issues: first, it makes battery handling asymmetric, as the optimiser actively minimises the peak, whereas the heuristic merely avoids exceeding today's peak. Second, this "EV-blind" threshold scales poorly: it performs reasonably with 1–2 fleets but degrades with 5 or 10 fleets, because the historical peak ceases to be a meaningful reference as the system evolves. In any case, a strict threshold-based heuristic is unlikely to be the way that decentralised charging will support DSOs. More plausible mechanisms include local peak pricing, time-of-use and dynamic tariffs, and occasional direct load curtailments under critical conditions. For these reasons, we use the optimiser as a benchmark for technical potential, leaving policy-based heuristics to future work.



When batteries are combined with smart charging (orange bars), the additional contribution is only marginal. This is because, as observed in the local peak-shaving analysis, batteries are ineffective once the underlying load profile has already been flattened. In practice, this would be comparable to deploying batteries today in a non-critical grid with no new loads expected, solely to reduce the existing peak loads: the impact is negligible and this is clearly not the best use case for battery storage.

To better compare the different solutions, it is useful to weigh their economic benefits against their deployment costs. For this purpose, an economic measure of grid relief (or stress) is required. However, accurate estimates of grid expansion costs are difficult to obtain and depend heavily on local specificities, which makes generalization impossible. Moreover, grid reinforcements typically follow a step-function logic: works are not undertaken until the grid load reaches a critical level and, when they are, capacity is usually oversized since material costs represent only a small share of the overall investment. Grid expansion costs are therefore a poor metric to capture the economic implications of grid stress in our stochastic simulations, where in some months the apparent operational limit is exceeded and in others it is not.

A more continuous indicator is better suited and for this reason we adopt the monthly peak-load pricing already used in the local peak-shaving case. The only difference is that the average peak tariff at medium voltage is assumed to be 8 CHF/kW/month instead of 9.

Figure 68 compares the different solutions within this economic framework as the number of QP fleets in the grid increases from 1 to 10. The black arrows indicate the rising monthly peak costs when no load management is applied. The rapid escalation observed in Figure 67 translates here into a cost surge from about 16'000 CHF/month to nearly 40'000 CHF/month.




Figure 68. Monthly costs for grid peak load (y-axis) vs. monthly costs for load management solution (x-axis) for the Högendorf L5 grid with a growing number of QP fleets (50 EVs each). Each arrows series represents the economic development under the same load management approach; all results are obtained with optimal control.

The violet arrows indicate the improvements achieved by smart charging alone. Naturally, management costs increase – especially when multiple fleets must be coordinated – but the benefits clearly outweigh the costs: with 5 EV fleets, the additional costs amount to only about 2'500 CHF/month. Adding a small



battery (yellow arrows) appears to deliver comparable benefits, but larger batteries increase costs without reducing peak-related expenses enough to justify the investment.

Acting with batteries alone (green arrows) proves to be a detrimental strategy. Using again the 5-fleet example, while a 5 MW battery can technically bring the load back to today's levels, the associated investment costs of roughly 13'000 CHF/month are not recovered and cannot compete with the performance of smart charging alone.

It is important to emphasise once again that this analysis considers batteries under the assumption that their sole application is peak shaving. In practice, batteries would also be operated for other use cases, likely unrelated to grid relief. The economic assessment of such cases lies outside the scope of this project. Nonetheless, the technical results presented here still offer valuable insights into the battery scales required to achieve given levels of peak-load reduction, independently of their economic viability.

Münchenstein L5 grid with multiple QP fleets

Next, we analyse the case where QP fleets are connected to the L5 feeder in the Münchenstein grid. The qualitative trends in Figure 69 are very similar to those observed for the Hägendorf grid in Figure 67. A key difference is that the uncontrolled introduction of the first few fleets does not cause major disturbances in the average peak load, since the additional charging demand occurs during off-peak times for the Münchenstein grid – namely in the evening. However, as more fleets are added, the average peak load increases sharply and clearly exceeds the current maximum observed load (dashed line).




Figure 69. Peak loads resulting from the gradual introduction of QP-like fleets (50 EVs each) in the Münchenstein L5 grid. Uncontrolled peaks are shown in light blue, while optimal charge management in violet. Only for the scenario with 5 fleets, the additional impact of batteries is explored: either alone (green) or combined with smart charging strategies (orange).

This highlights the risks posed by additional demand accumulating at what are currently off-peak times: interventions may be delayed because no criticalities appear at first. On the other hand, if the increase in load does remain limited (e.g. a maximum of two additional fleets), these are favourable situations



because grid reinforcements may be avoided altogether. Ultimately, it is the DSO's responsibility to anticipate, as far as possible, when and where new loads may be connected and whether they threaten safe grid operation.

If the additional load is managed with smart charging strategies (violet bars), the peaks remain well below the current operational limit, as the entire off-peak valley during the night can be filled by delaying the charging sessions of the QP fleets.

As before, complementing smart charging with batteries (orange bars) brings only marginal additional benefits, whereas batteries alone (green bars) appear to have a stronger impact. Interestingly, with 5 MWh of storage, batteries alone can perform almost as well as with smart charging schemes. This is related to the more "peaky" nature of the Münchenstein load profile, where batteries can reduce the midday peak by shifting a comparatively smaller amount of energy.

Finally, Figure 70 presents the economic results for this use case. Also here the qualitative trends resemble the outcome of the L5 Hägendorf grid in Figure 68.

The relatively small impact of the first fleets is evident from the clustering of many "1" and "2" fleet results around the "Today's costs" isoline. Despite the significant peak-load reductions, deploying batteries alone (green arrows) does not appear economically viable, as investment costs consistently exceed the achievable savings. Smart charging strategies remain the preferred solution, either on their own (violet arrows) or in combination with small batteries (yellow/orange arrows), which could in parallel be employed for other applications.




Figure 70. Monthly costs for grid peak load (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein L5 grid with a growing number of QP fleets (50 EVs each). Each arrows series represents the economic development under the same load management approach; all results are obtained with optimal control.



- Intensive EV fleets with synchronous behaviour (such as the QP delivery fleet) can significantly alter the grid load profile and increase the average peak.
- All else being equal, controlling the charging load – whether through direct or indirect measures – is likely the most economical pathway to contain the fleet-induced peak increase.
- Utility-scale batteries may also help contain the peak increase, particularly when the excess load is concentrated within a few hours of the day. However, their economic viability generally requires that they are also operated for other applications beyond peak shaving (e.g. voltage stabilization, PV absorption, ancillary services).
- In cases of criticality, DSOs may still favour grid reinforcement measures such as transformer upgrades, given their reliability and established practice. However, where cost or space constraints limit this option, more digital and decentralized approaches – such as managed EV charging – remain a feasible and cost-effective alternative.

3.4.4. Local PV consumption

Both previous simulation series focused on peak shaving. While this is a key indicator of potential grid stress, it does not translate directly into economic costs for the end user. We have also seen that batteries generally perform poorly when evaluated solely against peak-shaving benefits. This final simulation series therefore shifts the perspective to a cost-minimisation problem for electricity supply, considering also the availability of a local PV plant.

The scale of this analysis is the “local building” setup of the Münchenstein and Högendorf sites, each in the presence of a single fleet of 40 EVs. Both sites feature large rooftops with the potential to host up to 300 kW of PV capacity if fully covered. In the simulations, we examine cases with 0, 150, and 300 kW of installed PV capacity. The PV generation profile comes from the monitored generation load of the PV plant installed at the Münchenstein site. The nominal capacity of the reference plant is 20 kW and for this investigation we proceeded to scale it to the target level.

The cost components and assumptions underlying the simulations are as follows:

Energy and grid usage working price	23 Rp./kWh	Swiss average for a large commercial consumer (C4) ¹⁴
Peak power tariff	9 CHF/kW/month	Swiss average for a large commercial consumer (C4)
Feed-in tariff	13 Rp./kWh	Average of the 30 largest Swiss DSOs ¹⁵

The technological solutions investigated are the same as in the previous analyses: fleet charging load management, the installation and operation of a battery, or a combination of both. Only for this series of simulations we do additionally consider V2X technology, which is particularly suited to maximising PV own-consumption. There is still large uncertainty surrounding the commercial costs of V2X-enabled charging stations and the associated upgrades; based on the latest public announcements, we assumed a total cost differential of about 2400 CHF/EV compared to a non-dynamic wallbox, which, given a lifetime of 8 years, results in about 25 CHF/EV/month.

In terms of algorithms, only the optimiser has been applied, since no cost-minimisation heuristic was developed within the project. In any case, the “unbiased” behaviour of the optimiser provides the fairest

¹⁴ [Strompreise Schweiz](#) | [EiCom](#)

¹⁵ [pvtarif](#) | [VESE](#)



comparison between technological solutions, even if the resulting savings may lean somewhat to the optimistic side.

Importantly, the economic evaluation does not assess whether installing a PV plant is itself profitable, but rather which form of load management is most convenient under a given setup, defined by building baseload, fleet usage, and PV generation.

Münchenstein site with PE fleet

This setup reflects a situation very close to the live use case, since the building baseload, PE fleet usage, and PV generation profile are all derived directly from the same site. The only changes are the expansion of the electric fleet (40 EVs, all modelled as the service pool fleet) and an increase of the PV plant's nominal capacity to 150 or 300 kW.

Figure 71 presents a detailed breakdown of the monthly electricity costs for the case with a 150 kW PV plant. The blue bars show the expenditure for energy and grid usage proportional to the kWh taken from the grid, while the yellow bars indicate the earnings from feeding excess PV generation back into the grid. The orange bars represent the additional costs from the monthly peak load and the black points show the resulting net monthly expenditure. The large standard deviations illustrate that costs vary significantly across the year: as noted above, the building baseload in Münchenstein is strongly affected by heat pump operation and the addition of a PV plant further reinforces this seasonal imbalance in electricity procurement.

EV load management alone changes the situation only marginally: as shown in Figure 60, the peak load is reduced only slightly, while PV own-consumption improves from 60% to 70%, yielding savings of around 400 CHF/month. Thanks to the possibility of cycling energy between day- and night-time, V2X technology significantly improves the efficacy of the EVs, increasing the PV own-consumption rate to 82% and reducing monthly expenditure by 900 CHF/month compared to the base case.

Stationary batteries can emulate the gains of V2X, especially in terms of PV own-consumption, where they can achieve a rate of 86%. However, by having more freedom of operation, batteries are more effective in flattening the peak load (as also illustrated in Figure 60). This is particularly true when employing the optimiser, but even in a real setup a battery is more likely to have the necessary headroom to intervene at short notice during any peak occurrence. This allows for example a 300 kWh battery to achieve energy savings in the range of 1'300–1'500 CHF/month.

It is interesting to evaluate the “kWh ratio” between a V2X fleet and a stationary battery that can achieve the same own-consumption rate. In principle, a V2X-enabled fleet offers a significantly larger energy capacity: in the 40-EVs PE fleet example it is about 2'000 kWh. However, to increase PV own-consumption while preserving safe fleet operation, V2X vehicles must adhere to the following constraints: 1) they must always ensure a minimum SoC, 2) they must be plugged in during times of excess PV generation, and 3) they can discharge only if they can recharge with solar generation before their next trip, i.e. they do not depart early in the morning. Approximately, these constraints translate into reality factors of 0.5 (we assumed a minimum 50% SoC must always be guaranteed), 0.6, and 0.5 respectively. Taken together, these create an “operational reality factor” of about 0.15, reducing the theoretical flexible storage potential from 2'000 kWh to about 300 kWh. Indeed, as shown in Figure 71, the V2X fleet performs comparably to the 300 kWh stationary battery.




Figure 71. Breakdown of monthly electricity costs with different load management solutions at the Münchenstein site with 40 PE pool EVs and a 150 kW PV plant. Right y-axis and stars indicate the PV own-consumption rate achieved in each case.

Also in this assessment we evaluate which solution is economically more convenient by accounting for the cost of managing the charging process and/or the cost of procuring and operating a battery. Figure 72 presents the results for the different technological setups under three levels of PV installed capacity.

The first case with 0 kW PV mirrors the system analysed in Figure 61 without a PV plant: in the absence of PV, the only possible savings come from peak reduction, which was precisely the objective in Figure 61. Unsurprisingly, the results are very similar: none of the load management setups unlocks any meaningful net savings.¹⁶ Also V2X does not achieve cost reductions given its general unsuitability to reduce the peak load.

When PV capacity is added to the system, the energy procurement costs shown on the y-axis visibly decrease, even without any active management (black points). That follows from what was illustrated in the previous figure, where with 150 kW of PV capacity around 60% of the solar generation is own-consumed even under unmanaged conditions. This analysis does not assess whether installing additional PV capacity is economically beneficial in itself, since it focuses on identifying which load management strategies are most advantageous within a given setup. Nevertheless, it is worth noting that the average monthly electricity bills decrease by more than 3'000 CHF/month when adding 150 kW of PV and by a further 2'700 CHF/month with the next 150 kW.

In these PV-intensive setups, batteries become more effective and can generate additional net savings compared to the no-management scenario. In other words, the further increase in own-consumption (and peak shaving) helps to offset the investment in the battery system. By contrast, optimising only the EV charging process never yields substantial savings. This is partly because the total amount of energy that can be shifted is inherently small, given the light usage of the fleet, and partly because most uncontrolled charging sessions already occur around midday, thus naturally coinciding with solar gener-

¹⁶ Technically, Figure 61 used the heuristic algorithm, whereas here we apply the optimal solver. For consistency, the appendix reports the local peak-shaving results also with the optimiser: in that case, medium-sized batteries (100–200 kWh) achieve modest savings of around 200 CHF/month, which are comparable to the small gains visible here for the 150 kWh battery case. The cost-optimisation results in this section are actually slightly worse, since the additional battery cycling increases total energy purchases while reducing the peak. In any case, the very limited scale of the savings and the uncertainty surrounding them suggest that no load management system is worth considering under these conditions.



ation without the need for active management. The V2X solution does yield clear savings in energy and power expenditure, but it requires significant upfront infrastructure investments, resulting in virtually no net economic benefits compared to the no-management scenario

Since managing the charging process provides only limited benefits, adding its operational costs to a system already equipped with batteries tends to reduce overall savings. Therefore, larger battery systems that directly address the building baseload offer the best value for money. For both the 150 and 300 kW PV cases, a 300 kWh battery appears to provide the highest return on investment, with average net savings of approximately 530 and 720 CHF/month respectively against the no-management scenario.




Figure 72. Net monthly costs for electricity procurement (y-axis) vs. monthly costs for load management solution (x-axis) for the Münchenstein site with 40 PE EVs and PV plants of different capacity (individual plots). Points with the same colour refer to the same setup with varying battery sizes (respectively 150, 300, and 450 kWh).

Hägendorf site with QP fleet

In this final investigation, we analyse the Hägendorf site with 40 EVs sampled from the QP delivery fleet. As in previous Hägendorf analyses, the low building baseload is neglected to provide clearer insights into the performance of the load management systems in this setup.

Figure 73 shows the electricity cost breakdown for the configuration with 150 kW of installed PV capacity and different load management strategies. Since in this case the load consists entirely of the QP fleet's charging demand, the pure charge management solution already delivers visible savings, both through peak shaving (lower orange bar) and increased own-consumption. At first, this may seem counter-intuitive, given that the QP fleet is typically away during midday hours and therefore unable to directly capture solar generation. A closer inspection, however, reveals that the optimiser takes strong advantage of weekend periods: it avoids charging vehicles on Friday evening (and when possible it limits charging on Thursday evening) to instead schedule charging during the solar hours of Saturday and Sunday. This approach enables around 20% of the charging demand to be shifted into solar hours, increasing the average autarky rate by 17% and the own-consumption rate by 36%.

While a realistic control algorithm would necessarily act more conservatively and likely achieve smaller gains, this result shows a potential strategy that – given adequate weather and trip forecasting – could be effectively implemented for fleets that are typically in operation during solar hours.




Figure 73. Breakdown of monthly electricity costs with different load management solutions at the Hägendorf site with 40 QP delivery EVs and a 150 kW PV plant. Right y-axis and stars indicate the PV own-consumption rate achieved in each case.

Given this fleet pattern, adding V2X technology leads to only very minimal improvements. The rise in own-consumption visible in Figure Figure 73 stems mostly from an extreme attempt by the optimiser to reduce the peak load by discharging the few EVs that are occasionally not used the next day and can thus recharge using PV energy. But in a more realistic setup, the benefits from this scheme are virtually negligible.

Using only batteries does not lead to comparable gains and the general reason remains the relatively small battery size compared to the large amount of uncontrolled energy that would need to be shifted. In terms of peak costs, the large and concentrated peaks generated by uncontrolled charging can be shaved only marginally by batteries, as shifting a few kW of power requires moving a substantial amount of kWh. However, batteries perform more competitively in terms of PV own-consumption: the 300 kWh battery provides a service comparable to that achieved through EV load management alone. This is consistent with the scale of the system: given that the average daily consumption of the 40 QP EVs is about 1'500 kWh, the EV-management strategy shifts nearly one full day of charging to sunny weekend hours (i.e. around 1'500 kWh); similarly, the 300 kWh battery can shift its full capacity each day, amounting also to approximately 1'500 kWh over a work week. The main difference between the battery-only and EV-management approaches therefore lies in how they address peak loads.

The combined configuration benefits from the advantages of both solutions, at least until it reaches the saturation point of available excess PV energy, which for a 150 kW PV plant appears to occur around 300 kWh of battery storage.

Figure 74 compares the economic performance of these load management strategies across all setups with 0, 150, and 300 kW of installed PV capacity. Also in Hägendorf, adding PV capacity reduces the net expenditure for electricity procurement even without any active energy management (while still neglecting the associated investment costs). However, the savings are slightly lower than in Münchenstein – about 2'700 CHF/month and 5'400 CHF/month for 150 and 300 kW of PV, respectively – due to the lower natural overlap between solar generation and the uncontrolled load profile.



The case without PV generation resembles the setup analysed in Figure 63, where the objective was to minimise local peak load. Without solar generation, the cost-minimisation objective effectively converges to peak-load reduction, making the results directly comparable.¹⁷ In both analyses, the battery-only solution provides virtually no benefit, whereas EV management successfully shifts most of the charging load away from peak hours, hence reducing the associated costs. Adding batteries in the combined setup does not appear to be economically justified, making the EV-only management approach the most cost-effective option, with net savings of around 1'600 CHF/month (or 40 CHF/EV/month).




Figure 74. Net monthly costs for electricity procurement (y-axis) vs. monthly costs for load management solution (x-axis) for the Hägendorf site with 40 QP EVs and PV plants of different capacity (individual plots). Points with the same colour refer to the same setup with varying battery sizes (respectively 150, 300, and 450 kWh).

In the setups with PV plants, batteries can deliver some net reductions in electricity bills; however, as anticipated in Figure 73, their performance remains well below that of EV-management solutions. Achieving comparable results, particularly in terms of peak shaving, requires large battery capacities and once the associated investment costs are considered, the overall economic assessment does not favour their use. Combined solutions also do not appear to be cost-effective. Nor is V2X, which delivers the same benefits as dynamic EV charging but at approximately three times the cost. Consequently, dynamic EV-only management emerges as the most economical approach, yielding savings of about 2'100 CHF/month (or 52 CHF/EV/month) for both the 150 and 300 kW PV setups.

- Batteries can leverage their ability to shift energy between day and night, generating consistent savings in all scenarios with PV capacity.
- The effectiveness of EV charge management solutions is highly dependent on the fleet usage intensity and the uncontrolled charging profile: the more intensively the fleet is used, the more cost-effective charge management becomes. Even when few vehicles are available for charging during midday hours, well-tuned systems can still take advantage of weekends to achieve meaningful savings.
- V2X technology can boost energy savings from PV self-consumption, especially if the fleet

¹⁷ The only difference is that the results in Figure 63 are slightly less favourable, as they were obtained using the heuristic algorithm, while the analyses presented in this section rely on the optimal solver.



has extended idle times during sunlight hours, as in the Primeo Energie pool fleet in Münchenstein. Whether this is worth the higher upfront charging infrastructure costs depends on the cost differential and the realised savings, but in the investigated fleets this did not appear to be the case.

- The peak-power component of electricity bills can be surprisingly significant for fleet-centred businesses such as delivery companies. Reducing these costs translates directly into peak-shaving strategies, which are most effectively addressed through smart charging systems. In this contexts batteries can provide support, but they generally struggle to shift the bulk of the charging load.
- The best return observed from EV charge management in our simulations was 52 CHF/EV/month for the intensive QP fleet under PV-generation setups.

3.4.5. Summary of digital twin results

Figure 75 summarises the main results of the digital twin simulations. For each use case, a selection of representative setups is presented to illustrate the overall profitability of EV flexibility. Solely EV-only-management solutions are shown, i.e. without battery involvement (i.e. the numbers correspond to the results typically shown in violet in the previous plots). For the local peak shaving case, results are displayed for all three algorithms employed: load-capping (Algorithm 0 from the live use case), dynamic control (Algorithm 2.3) and optimal control using the MILP solver.

The reported savings represent the net economic benefits of EV load management compared to the unmanaged scenario, accounting for the costs of installing and operating the charge management system. For the optimal case, the same costs as the dynamic control setup are applied, since it can be considered the best-case scenario of that method. These figures can be cautiously compared with the “Total Earnings per EV” presented in Figure 56, with the distinction that the system operator costs have already been deducted here. In other words, the values shown here represent the net earnings to be shared between the fleet operator and the charge system operator.

Nevertheless, the findings from the simulations and the live use cases are fully consistent. The use cases involving the PE fleet barely generate any profit and in some cases even result in negative outcomes. The only exception observed in the live use case concerned ancillary services, which however were not simulated in the digital twin environment.

On the contrary, charge control of the QP fleet can deliver meaningful income for the operators. Reducing the local peak can generate savings in the range of 30–40 CHF/EV/month, consistent with observations from the live use case. Simple load-capping achieves comparable net savings but with lower system complexity, suggesting that this application may not represent the most suitable niche for dynamic charge control with advanced EV-data-integration.

Reducing the peak load on the L5 grid can also be profitable, although in this case it is much harder to predict the actual load profiles and the type of incentives that DSOs may introduce to address grid congestion. The live use case based on current load conditions showed earnings of around 10–15 CHF/EV/month, while simulations under a critical scenario – with five QP fleets connected to the same L5 feeder – indicate a potential added value of 25–35 CHF/EV/month.

Finally, the local PV self-consumption case delivers the most compelling results. The PE fleet setup in Münchenstein closely reflects the findings from the live use case, showing negligible or even slightly negative financial gains.




Figure 75. Summary of net economic benefits for a few key illustrative use cases as obtained from the digital twin simulations. The different columns correspond to different use cases and representative setups; the different colours to differ algorithmic methods.

Conversely, the setup with the QP fleet (with the low baseload of Hägendorf excluded) exhibits the highest potential economic value, with net savings reaching up to 52 CHF/EV/month. These results underscore that specific local conditions (fleet utilisation, PV installed capacity, economic incentives) have a paramount impact on the business case for EV fleet flexibility.

- The simulation results confirm all findings from the live use cases.
- The fleet usage intensity plays the most decisive role in determining the economic viability of EV flexibility. None of the simulated use cases showed meaningful profits for the lightly used Primeo Energie pool fleet (ancillary services being the only exception in the live use case).
- The live use cases restricted to the Münchenstein location naturally produced low savings, most likely underestimating the average revenue potential of EV fleet flexibility.
- For instance, the local PV consumption case generated significant savings when simulated with the Quickpac fleet in Hägendorf, achieving up to 52 CHF/EV/month in net gains.
- While 52 CHF/EV/month represents a notable reduction in fleet-related energy costs, it remains relatively modest for a business whose total fleet operating costs range between 700 and 1'000 CHF/EV/month. This implies that any smart-charging provider engaging with such businesses must ensure that the proposed system is exceptionally lean: plug-and-play, reliable and free of burdensome setups, maintenance or user intervention.



4 Conclusions and outlook

4.1 Deployment findings

Vehicle system integration

- The vehicle interface over the OBD channel provided the basic data (e.g. SoC, 12V voltage level, charging plug status) for all car models in the project fleet. However, a big variability in the availability of more advanced metrics (e.g. drawn power, HV battery current) current still exists between car models.
- Web-based API data interfaces as an alternative exhibit even more scattered data and control accessibility throughout manufacturers and models. Data availability also shows large fluctuations.

Vehicle system operation

- The vehicle systems ran reliably with only two effects impacting the operation.
- On the one hand, modern vehicles would have a modern CAN architecture making it necessary to actively request data from the vehicle systems. This would then keep all systems alive and drain the 12V battery. The possibility of charging the 12V battery from the (fully charged) high-volt battery could not be observed for multiple vehicle models, which would lead to inaccessibility of the car.
- On the other hand, theft alarm settings would inhibit the proper use of the OBD socket for special situations on a vehicle sub-model. This shows how very car-specific setups can interfere with a modern data fetching system.

Charge control system

- Even though the charge control happened through a very standardised interface (Type 2 charge plug), significant differences in the handling of control signals were still attested. EVSEs would sometimes rate a certain feedback from the EVs differently and EVs would in turn occasionally react inconsistently to advanced charge control sequences. This shows that advanced control sequences are not yet sufficiently standardised and that OEMs continue to develop their own proprietary implementations of charging processes.

Local power metering

- While a large variety of metering setups was used within the scope of the project, none of the solutions appeared ready to deliver the full range of benefits expected by an advanced customer seeking to optimise its energy operations. Direct access to meter data and its forwarding through a user network remain cumbersome and manufacturer-specific solutions are still required to utilise already acquired data. Here too, standardisation seems to be lacking for the advanced and further use of existing metering data.

DSO power metering

- Thanks to access to more advanced tools, the integration of live data monitoring and acquisition for the level 6 transformer proved to be more straightforward.
- The main challenge encountered concerns the potentially restrictive attitude of some grid operators toward sharing such grid data. Smaller operators in particular often show limited interest



in using this information for potential grid-relief applications. This is likely due to the additional effort required to (1) install extra equipment in transformer cabinets and (2) establish a structured method for publishing this data in a condensed form.

- Level 5 grid data was also easy to obtain and process thanks to the open and collaborative attitude of Primeo Energie and the SCADA system they already operate to monitor those grid levels.

EV fleet operation

- Fleet operations are not only tailored to the specific business needs, but also depend on the local building layout and the standard practices of the fleet management team.
- The project demonstrated that even small adjustments to the charging schedules can have an impact on the business processes and thus raise the need for dedicated procedural changes within the company.

Fleet booking data

- Access to fleet booking data through a web REST API proved highly effective and provided the necessary information for advanced charge control.
- However, it seemed that the effort required to set up such a system would only be considered in pool-fleet environments, where a vehicle booking system is already a core operational need.
- High availability of pool cars can also reduce the need for pre-booking, resulting in only short-notice information about upcoming trips.
- In the absence of a booking system, static working schedules appear to be the most practical option for a broader range of users.

IT system operation

- The chosen architecture with an MQTT data broker proved highly effective. The ability to classify data points for access control and for different data consumers met all requirements. This setup also enabled secure and reliable data transfer from edge devices to the cloud, which also supported instant event forwarding for ancillary service requests.

Change management

- Having an advanced charge management system that optimises the charging process introduced additional complexity for the users in the project fleets. The dynamically adjusted charging schedules made it difficult to maintain a clear understanding of vehicle status and expected readiness.
- Supporting tools such as online dashboards can partly restore this missing overview; however, they must be complemented by a change management process to ensure that all users become comfortable with the new operations. This represents a management challenge for the deployment of such advanced charge control systems.



4.2 Use case findings

Local (building) peak shaving

- Local peak shaving has been the most common and widespread application of charge management systems. This is mostly due to the current regulation and the easier setup of such a solution. The project showed that at this time this remains the most profitable use case if the fleet has a significant size and usage.
- Despite some uncontrollable loads, the charge management system in Högendorf managed to reduce the building's peak load by about 20–40 kW while managing 40–50 kW of charging power, all without compromising fleet readiness.
- Simulations show that dynamic charge management systems leveraging vehicle and operational data can unlock greater benefits than standard data-agnostic load-capping approaches; however, these advantages do not necessarily offset the higher implementation costs and the resulting economic return may remain limited.

DSO grid peak shaving

L7 peak shaving with fast charging

- The setup with fast charging stations in the L7 grid level showed a typical load-drop case since these fast charging stations would insert high power peaks into that grid.
- Limited availability of vehicles for charging, however, reduced the overall effectiveness of the case: fast charging sessions occurred at highly variable times of day, and the controlled fleet was not always available to compensate.

L5 grid peak shaving

- On the L5 grid level the baseload shows a much more static behaviour that is still very location-specific. A time-based “peak-avoidance” approach can help to flatten the overall load and lead to a more efficient grid usage.
- In Högendorf this strategy could reduce loads by about 30–40 kW during the targeted peak intervals. Although the impact on the L5 grid was not detectable due to the limited fleet size, the results indicate that larger EV pools can meaningfully mitigate such peaks.
- Simulations indicate that connecting EV fleets with strongly synchronous behaviour (such as delivery fleets) can pose risks to safe grid operation. While conventional measures like grid expansion may remain the preferred option for DSOs, we show that smart charging schemes can significantly address the issue and are a viable and cost-effective alternative.
- Designing an effective “peak shaving” incentive is however challenging, as load profiles vary widely across regions and seasons, and the corresponding monetary elasticity is difficult to quantify. Dynamic tariffs are a step in this direction, but current designs remain rather generic: neither grid-specific nor adaptive to weather forecasts.

Ancillary services

- Ancillary service provision remains the “royal discipline” of energy flexibility: potential revenues are high, but so is the associated technical complexity.
- EV fleets currently lack both sufficient overall availability and the reliability required for precise control processes.
- Secondary control (SRL/aFRR) still entails overly demanding constraints that most charging setups cannot meet.



- Tertiary control (TRL/mFRR) appears more feasible from both a technical and regulatory perspective. However, the minimal bid size of 5 MW still raises the need for non-trivial asset pooling in order to fill the missing availability blocks or to run a dedicated bidding strategy.
- Capturing the revenue potential of an EV fleet primarily depends on maximising its available power capacity, which in turn requires a finely tuned charge control strategy tailored to the specific fleet usage profile.
- Without V2X technology, EV fleets can barely provide products other than negative control, limiting their turnover potential compared to stationary batteries.
- As of today, the most promising ancillary service use case is the provision of negative tertiary control through daily tenders: this could generate a gross annual revenue of up to 370 CHF/EV, which would however need to be shared among the pool manager, the charging system operator and the fleet owner.

Local PV consumption

- Managing charging sessions to maximise direct solar utilisation can be an effective strategy, but the actual gains depend strongly on the load-shifting potential of the EV fleet. This potential is determined by when the vehicles are available for charging, which in many cases may be “counter-cyclical” to the typical PV generation profile.
- Primeo Energie’s pool fleet in Münchenstein almost naturally charges during sunlight hours; therefore, the benefits of active charge management are limited. Nevertheless, live tests with the fleet resulted in a 10% increase in own-consumption and a shift of about 80 kWh/week from grid feed-in to direct use. However, given the small spread between feed-in and grid tariffs, the corresponding economic gain per EV remains minimal (about 2.70 CHF/EV/month).
- More promising results were obtained in the simulations for Quickpac’s fleet in the presence of a large PV array (150–300 kW). Despite being mostly in operation during daylight hours, Quickpac’s delivery EVs could take advantage of excess solar generation during weekends, achieving savings of up to 52 CHF/EV/month (with the optimiser).
- At the time of the project, grid and feed-in tariff designs provide incentives to maximise self-consumption only for direct PV owners. Therefore, EV fleets have little motivation to react to the solar capacity or output of neighbouring systems. Upcoming changes with V-ZEV and LEG frameworks will broaden these opportunities, albeit with lower incentives due to intermediary fees.

Dynamic tariffs optimisation

- Dynamic tariffs represent an efficient and implementation-friendly way to incentivise energy flexibility.
- Within the project, the VARIO dynamic tariff was tested; however, it proved to still be rather “static” on a day-to-day basis, thus failing to fully exploit the available flexibility potential.
- The incentives created by the price spread and the resulting benefits remain rather modest, particularly for fleets with limited energy turnover.
- Dynamic tariffs appear suitable mainly for addressing broader, system-level grid objectives rather than local or regional effects. This reflects the current “utility-wide” tariff designs; introducing more granular, spatio-temporally dynamic tariffs seems unlikely given the already challenging customer onboarding process.



4.3 General conclusions

4.3.1. Implementation assessment

- Medium sized fleets of 10–40 EVs can have a significant impact on the distribution grid – dependent on their usage patterns and synchronicity.
- Live charge control of EVs is possible but most key interfaces are not standardised, which leads to some uncertainty and relative unreliability.
- Flexibility use cases that require a “proof of activation”, such as ancillary services, bring along additional metering and accounting needs. While this can be financially justified for critical or “premium” ancillary services, other more local use cases would need to be financially compensated for such a scheme or simplified to go with another verification solution.
- Many flexibility use cases still lack a clear framework for service activation and remuneration. This is further compounded by the absence of technical standards needed to support such a framework.
- In most cases reliability is key and it’s core to the value creation. Current EV charge control setups cannot guarantee the required reliability level yet.
- A standardisation of data interface and regulated fees would be an big enabler of EVs energy flexibility.

4.3.2. Flexibility evaluation

- The success of flexibility use cases is first and foremost dependent on the fleet utilisation levels, i.e. on their overall energy throughput and associated expenditure. Driving schedules play a secondary, albeit important, role.
- Low fleet utilisation leads to very narrow charging windows, which significantly limit the flexibility potential and the associated economic value. A notable exception is represented by negative ancillary services, which can take advantage of the long idle times of this type of fleet.
- Most of the flexibility use cases eventually result in shifting loads from morning and evening peaks to midday/early afternoon or night times.
- Peak shaving cases are generally quite profitable but very situation-specific and hard to generalise. The observed diversity in fleet usage patterns and grid congestion levels indicates that there is no universal incentive scheme that can appeal equally to all fleets and ensure the best outcomes in every context.
- Local (behind-the-meter) peak shaving is a simple use case with clear pricing incentives that can already be profitable today; however, standard load-shedding mechanisms that do not interact with EVs can capture 70–80% of this value at little to no cost.
- The common lack of a clear overview of distribution grid congestion levels and operational limits prevents an accurate assessment of both current and future grid criticalities; consequently, the relief potential and economic value that EV fleets can provide today or in the future can barely be quantified and properly valued by grid operators.
- Less strict use cases that allow for occasional inconsistencies, such as dynamic tariffs and PV self-consumption, are much easier to comply with and to implement.



4.3.3. Economic validation

- After considering the multiple actors involved in each use case, the final net revenue or savings for fleet operators is rather low, typically in the range of 1–15 CHF / EV / month.
- From an economic perspective, local peak shaving proves to be the most valuable flexibility use case, particularly for intensive and highly synchronous fleets.
- Dynamic tariffs are promising, but the price spread is still rather low, especially for the investigated tariff and limited-use fleets.
- The monthly peak power tariff appears somewhat outdated and it may provide suboptimal incentives. In most grid configurations, what truly matters is not the individual customer's maximum peak, but the aggregated load levels during periods of high network congestion. The single-customer peak power tariff likely persists due to its simplicity and its indirect benefits for grid stability. However, with rising concerns over grid congestion and the increasing availability of flexible, grid-supportive assets, an update of this tariff structure should be considered.
- Strong business- and system-dependencies on OEMs currently dominate the charge control interface, thus business models built on these systems carry a high level of risk.
- From a purely financial standpoint, stationary batteries appear to be a much more expensive alternative for flexibility services than decentralised control of EV fleets; however, they bring major advantages in terms of simpler interfaces, more common installation practices, higher reliability, and constant availability.
- While EV-fleet flexibility may in some contexts represent, in principle, the most economical and efficient solution for a grid operator, the generated added value – after it has been transferred down the chain until the fleet operator – often remains too small to justify investing in advanced charge management. This seems to also suggest that the forthcoming, long-dreaded expansion of grid capacity and flexibility is not yet fully “priced in”: end consumers do not yet see this transition reflected in their bills and therefore do not perceive sufficient incentives to participate in flexibility schemes.

4.3.4. Use case comparison

Use case	SoC / drive data need	Complexity	Criticality of delivery	Expected operator revenue	Expected benefit for the grid
Local peak shaving	Low	Low-Medium	Medium-High	High	Low-Medium
DSO grid peak shaving	Low	Medium	Medium-High	Medium	High
Ancillary services	High	High	High	High	High
Local PV consumption	Medium	Medium	Low	Medium-High	Medium
Dynamic tariff optimisation	Medium-High	Low	Low	Medium	Medium-High



4.4 DSO's take

Decarbonization – and the associated electrification – constitutes one of the most profound transitions in the power sector and places substantial operational and strategic demands on DSOs. Effective congestion management in distribution grids requires new approaches to mitigating demand peaks. While conventional grid reinforcement is often the default response – bringing clear benefits – it ties up significant financial and human resources. Increasingly, expansion is constrained by site availability and suitable land, including for battery energy storage systems (BESS).

This is precisely where energy and charging management systems – such as VEIU Tech's solution – can provide near-term relief. When paired with congestion signals, for example dynamic grid tariffs reflecting day-ahead load forecasts, these management systems can unlock their grid-supporting potential. Dynamic tariffs operate as a price-based alternative to direct control actions and avoid hands-on intervention in customer installations; accordingly, Primeo Energie already offers dynamic optional tariffs in its service area.

Not all peaks are predictable (e.g., sudden spikes from DC fast charging). Dynamic tariffs alone won't mitigate those events. The experimental setup demonstrates how such peaks can be addressed by adding more "intelligence" within the DSO's network. However, integrating third-party energy management solutions with the DSO's SCADA and metering environments still faces regulatory, technical, and tariff-design hurdles. Establishing end-to-end, standardized interfaces must therefore be a priority.

Field experience with Primeo Energie's fleet shows that dynamic charging based on SoC and vehicle-usage data clearly outperforms less "intelligent" charging and self-consumption solutions that are common in the grid.

Finally, the pilot illustrates how EV fleets can be activated as flexibility in balancing markets. Beyond today's communication heterogeneity on the vehicle side, the sector still needs workable coordination between DSOs and the TSO, as well as fit-for-purpose market mechanisms to bring distributed flexibility from grid level 7 into practice.

4.5 Outlook

Modelling and assessment suggestions

- The heuristic algorithms used in the simulations for the joint control of EV fleets and stationary battery showed rather inconsistent performance, particularly under varying building and PV load conditions. The performance gap between these setups and the perfect-foresight simulations remained significant. This indicates that, from a control perspective, such configurations may be sensitive and require a fine calibration in both simulated and real environments. Therefore, the control algorithms for these joint setups should be further improved and tested to ensure the most synergistic use of both flexible assets.
- The economic viability of smart management systems will be strongly influenced by the introduction of new types of dynamic tariffs (e.g. those incorporating day-ahead energy prices or PV generation forecasts) and by the forthcoming implementation of dynamic hourly feed-in tariffs. These developments are worth exploring, as they may be part of the decisive changes that ultimately unlock the case for grid-friendly dynamic charging.
- The potential of V2X has not been explored, as its current implementation costs still appear too high relative to the expected benefits, while also being affected by the same standardisation and data-access barriers discussed in this work. Nevertheless, the resolutions of those barriers, a reduction in V2X deployment costs and the aforementioned framework developments



may also unlock the economic case for this technology, which therefore merits further investigation.

General recommendations

- Standardisation of vehicle control interfaces is strongly needed in order to improve applicability and reliability, and to reduce access costs for EV flexibility. Without it, the widespread adoption of advanced charge control will remain difficult.
- Access to meter data and other inputs necessary for providing grid relief should also be standardised and simplified. This is particularly relevant in cases where data is already being collected but not yet packaged for publication.
- Simplifying grid flexibility products should be a key objective, for example through dynamic tariffs and/or DER flexibility platforms with more relaxed constraints to better integrate EV flexibility.
- The incentives for energy flexibility need to be adjusted to provide a meaningful benefit for fleet operators, sufficient to offset the complexity of adapting their operational processes.
- A more detailed understanding of grid status and operational limits could help tailor flexibility products to local conditions and make flexibility dispatch more effective.



5 Appendix

5.1 Back-of-the-envelope use cases evaluation

In the following table, we present a simplified “back-of-the-envelope” assessment of the economic performance of the different use cases for the participating EV fleets. Its purpose is to provide a convenient framework that reduces the complexity of fleet operations and use case economics to enable quick profitability estimates. The resulting numbers may therefore differ somewhat from the empirical results reported in the main document, as those are based on more detailed and realistic analyses.

The table also includes the economic assumptions used during the preparation of the project submission proposal. This allows to highlight in a simplified way the main insights gained from the project and how these have impacted the initial assumptions and expected economic returns of the various use cases. In other words, the table serves as a “before-” and “after-project” comparison of the economic potential of dynamic smart charging.

All figures refer to fleets composed of 40 EVs. The assumptions made in the project proposal were based on fleet patterns similar to those of the Quickpac fleet and are indicated in the table with the label “ref”. We highlight in green and red the assumptions that have changed most significantly compared to the *a priori* estimates before the project start (in green for positive and in red for negative impacts on the dynamic charge control business case).

Use case	Fleet	Theoretically available flexibility (power)	% of flexibility accessible / useful	Actually accessed flexibility (power)	Hours / month	Price/savings per flexibility unit	CHF / month / fleet	CHF / month / EV
Local (building) peak shaving	ref	400 kW	18%	70 kW		9 CHF/kW	630	16
	QP	400 kW	33%	132 kW	-	9 CHF/kW	1190	30
	PE	400 kW	12%	48 kW		9 CHF/kW	430	11
DSO grid peak shaving	ref	400 kW	18%	70 kW		7.50 CHF/kW	520	13
	QP	400 kW	33%	132 kW	-	8 CHF/kW	1060	26
	PE	400 kW	12%	48 kW		8 CHF/kW	380	9.60
SRL- (capacity provision)	ref	400 kW	25%	100 kW	720	50 CHF/MW/h	3600	90
	QP	400 kW	not addressed	-	-	-	-	-
	PE	400 kW	20%	80 kW	490	14 CHF/MW/h	550	13.70
SRE- (energy activations)	ref	same as SRL-			60 (8.3%)	100 CHF/MWh	600	15
	QP				-	-	-	-
	PE				17 (3.5%)	320 CHF/MWh	440	11
TRL- (capacity provision)	ref	same as SRL-			720	50 CHF/MW/h	3600	90
	QP				-	-	-	-
	PE				490	13 CHF/MW/h	510	12.70
TRE- (energy activations)	ref	same as SRL-			60 (8.3%)	100 CHF/MWh	600	15
	QP				-	-	-	-
	PE				44 (9%)	130 CHF/MWh	460	11.40



TRL+ (capacity provision)	ref	400 kW	25%	100 kW	720	50 CHF/MW/h	3600	90
	QP	400 kW	25%	100 kW	245	3 CHF/MW/h	73	1.80
	PE	400 kW	practically never	-	-	-	-	-
TRE+ (energy activations)	ref	same as TRL+			60 (8.3%)	100 CHF/MW/h	600	15
	QP				-	-	-	-
	PE				49 (20%)	70 CHF/MW/h	340	8.60

Use case	Fleet	Theoretically available flexibility (energy)	% of flexibility accessible / useful	Actually accessed flexibility (energy)	Days / month	Price/savings per flexibility unit	CHF / month / fleet	CHF / month / EV
PV local consumption	ref	800 kWh	50%	400 kWh	12	6 Rp/kWh	288	7.20
	QP	1200 kWh	100%	1200 kWh	4	10 Rp/kWh	480	12
	PE	120 kWh	25%	30 kWh	10	10 Rp/kWh	30	0.75
Dynamic grid tariff optimisation	QP	1200 kWh	40%	480 kWh	20	11 Rp/kWh	1060	26
	PE	120 kWh	50%	60 kWh	20	11 Rp/kWh	130	3.30
Spot market day-ahead arbitrage	ref	800 kWh	50%	400 kWh	20	30 CHF/MW/h	240	6
	QP	1200 kWh	40%	480 kWh	20	50 CHF/MW/h	480	12
	PE	120 kWh	50%	60 kWh	20	60 CHF/MW/h	72	1.80

Assumptions and considerations:

- For the local (building) peak shaving use case the key input is the share of the fleet that contributes to the overall peak load and whose charging can be shifted to another time. In the case of Quickpac the fleet is actually causing the peak, therefore shifting most charging sessions brings significant relief. However, not all sessions can be shifted as this could create a new peak, and operational constraints must also be considered. For the PE fleet, the share of EVs charging during the office building peak load is low due to the limited charging needs.
- The DSO grid peak shaving use case is assessed in a similar way to the local peak shaving case, although the price incentives tend to be lower.
- For the SRL-/TRL- cases, we did not investigate the potential of the QP fleet. In principle, this fleet does not have suitable driving and parking patterns, as the times when it could provide negative control are extremely limited (only a few hours in the evening before the scheduled charging for the next day). Nevertheless, we cannot exclude that offering negative control during a few night-time blocks could generate some interesting revenue.
- For the TRL± prices, a weighted average of weekly and daily tenders was used. Daily tenders are considerably more remunerative, but the likelihood of being awarded capacity is lower.
- In the case of TRL±, the PE fleet's contribution is negligible, as its charging load is minimal and thus offers limited potential for providing positive control. The QP fleet, on the other hand, can provide this service, but only while charging (approximately 12 hours per day between 16:00 and 04:00).
- SRL± was not considered, as the results would be similar to TRL+, which already proved to be not particularly remunerative.



- For the PV local consumption use case the key factor is the average amount of charging energy that can be shifted from non-daylight to sunlight hours.
 - In this case the two fleets behave quite differently: for the PE fleet only a small share of vehicles can take advantage of this opportunity, but they can do so almost every day with meaningful PV generation (estimated at about 12 working days per month on average). The QP fleet, on the other hand, cannot participate in this scheme during regular operations but can virtually shift the entire Friday (and part of Thursday) charging load to the weekend in case of sunny weather. The entire fleet can therefore participate, although only for approximately four days a month.
 - In both cases it is important to highlight that such savings are only possible if shifting the EV charging load directly increases the self-consumption rate of the local PV plant. For instance, in the actual setup in Münchenstein the PV plant is relatively small compared to the overall building load, resulting in an almost constant 100% self-consumption rate. In such cases, shifting the PE fleet's charging load brings no additional benefit under this use case.
 - The resulting savings for the QP fleet of about 12 CHF/EV/month may appear lower than the 52 CHF/EV/month obtained from the digital twin simulations. However, the latter reduction also included the effect of local peak load reduction, which according to the back-of-the-envelope table may account for an additional 30 CHF/EV/month. The advantage of these use cases is that the two effects are not conflicting and can likely be achieved simultaneously.
- The dynamic grid tariff optimisation use case was not investigated prior to the start of the project; thus no *a priori* economic estimate for QP* is provided. The estimates for the two fleets consider 1) the share of EVs that have an incentive to shift the start of their charging sessions to capture part of the tariff spread and 2) the average magnitude of that price spread. This optimisation can in principle be performed every working day, i.e. about 20 days a month.
- A final potential use case involves taking advantage of a “spot contract”, where the energy price is based on day-ahead spot market prices. In this case, the EV fleet could achieve some level of arbitrage by shifting charging loads according to market price variability. The optimisation logic is similar to that of the dynamic grid tariff case: only part of the fleet can actively benefit from these price spreads by adjusting their charging times. Contrary to dynamic tariffs, however, the average daily price spread is smaller, resulting in overall lower savings potential per EV.
- Even if not included in the table, an interesting use case considered in the project submission was the live compensation/adjustment/balancing of the balance group, with the goal of avoiding penalties resulting from physically imbalanced positions. This use case was not tested or simulated during the project due to the difficulty of precisely defining how such a scheme would be implemented and operated in practice. During discussions with relevant professionals, the following technical and economic challenges were highlighted:
 - Current poor visibility of the live balance of a balance group. Setting up a system that enables complete real-time monitoring and intervention is complex and costly. Current approaches rely primarily on forecasts and frequently updated short-term forecasts (particularly for PV generation).
 - If information about an imbalance in the time window T–T+00:15 becomes known, the balance group can still choose to settle it via the intraday market until approximately T-00:05. Therefore, a decentralised energy resources (DERs) must offer an intervention that is as reliable and cheaper than simply correcting the imbalance through a last-minute intraday trade.
 - If an asset can intervene with high reliability, short notice, and fast ramp-up, it is usually more attractive to offer this flexibility to classical ancillary services, thereby making it available to the entire Swissgrid control area instead of being restricted to a single balance group.




- The possibility of relying on post-scheduling adjustments to settle part of the imbalances further reduces the incentive for this use case.

5.2 Digital Twin

5.2.1. Fleet load peak shaving

In the following figure, we use the simple peak-shaving use case with PE fleet alone (i.e. without any additional building load) to determine the “battery-equivalent” capacity needed to reduce the peaks of an uncontrolled fleet to the same level as a perfectly managed charging load. The solid line represent the normalised peak load (kW/EV) obtained for different fleet sizes with heuristic (green) or optimal (purple) control; the dashed lines instead display the peak loads obtained by optimally managing stationary batteries of the stated size without controlling the fleet. The three proposed battery size “scales” have all a C-factor of 0.4. The figure shows that a battery sized with 1.35 kWh per EV in the fleet can already approximate the performance of the heuristic optimiser.




5.2.2. Local (building) peak shaving

The following figures present the results obtained for the local peak-shaving use cases using the optimal solver. In the EV-management scenarios, no distinction is made between “load-capping” and “dynamic charging,” since operating the charging process with perfect foresight inherently assumes full access to the SoC. For the same reason, when assigning costs to the different load management solutions, the EV charge management cases use the cost assumptions for the SoC-aware dynamic charging solutions. For the battery systems, the same cost parameters used in the main report with the heuristic solver have been retained.


We also present the results obtained for different fleet sizes, i.e. with 10 or 30 EVs per fleet. In all plots the points refer to the results of single stochastic simulations, while the rhomboids indicate the average of each load-management setup. Rhomboids with the same colour depict the same setup with varying battery sizes (20, 50, 100, 200, and 500 kWh).



Münchenstein site (office space baseload) + 50 PE EVs, optimal solver:




Münchenstein site (office space baseload) + 30 PE EVs, optimal solver:






Münchenstein site (office space baseload) + 10 PE EVs, optimal solver:




Hägendorf site (no baseload) + 50 QP EVs, optimal solver:






Hägendorf site (no baseload) + 10 QP EVs, optimal solver:




Münchenstein site (office space baseload) + 50 QP EVs, optimal solver:






Münchenstein site (office space baseload) + 30 QP EVs, optimal solver:



Münchenstein site (office space baseload) + 10 QP EVs, optimal solver:







5.2.3. Local PV consumption

The following figures present the detailed breakdown of monthly electricity costs obtained in the “Local PV consumption” simulations with a PV installed capacity of 300 kW. All results are derived with the optimal solver.

Münchenstein site (office space baseload) + 300 kW PV + 40 PE EVs, optimal solver:



Hägendorf site (no baseload) + 300 kW PV + 40 QP EVs, optimal solver:





5.3 Complete Primeo Energie's pool fleet user survey

Attached we provide the full results from the user survey carried out with the drivers of Primeo Energie's pool fleet in Münchenstein.

Responses Overview Active

Responses

104



Average Time

08:07




Duration

260 Days




1. Was ist die typische Fahrt, die Sie mit den Fahrzeugen der Poolflotte unternehmen?

Fahrt zu einem/r Kunden/in	37
Fahrt zu einer Netzeinrichtung	38
Fahrt zu einer Veranstaltung/Sitzung	26
Servicefahrt	3




2. Wie oft nutzen Sie ein Poolfahrzeug?

nie	3
< 1x pro Woche	40
1x pro Woche	29
2-3x pro Woche	30
>4x pro Woche	2




3. Wie oft fahren Sie mehr als 100 km pro Buchung?

nie	51
1x pro Monat	43
2-3x pro Monat	4
>4x pro Monat	4




4. Welchen Autotyp wählen Sie am liebsten für ihre Fahrt?

VW eUP	10
VW eGolf	23
VW ID.3	35
BMW i3	10
Skoda Enyaq	19




5. Wie oft buchen Sie spontan (<1h vorher) ein Poolauto?

nie	12
selten	38
manchmal	39
meistens	13
jedes Mal	2



6. Fühlen Sie sich durch die Nutzung von Elektrofahrzeugen eingeschränkt?


ja	4
manchmal, aber ich kann damit umgehen	22
nein	78



7. In welchem Mass werden folgende Einschränkung wahrgenommen (1 wenig einschränkend, 5 sehr einschränkend):


1 2 3 4 5

begrenzte Reichweite
umständlicher Ladevorgang
langsameres Aufladen während der Fahrt im Vergleich zu Tankstelle
langsameres Laden in der Garage
Unsicherheit über den Ladezustand beim Start




8. Mussten Sie bei einer Ihrer Fahrten eine öffentliche Ladestation benutzen?

ja	25
nein	79




9. Wenn nein: Wüssten Sie, wie man einen öffentlichen Ladevorgang durchführt?

ja	18
nein, aber ich würde es herausfinden	49
nein	12




10. Wenn ja: Warum haben Sie es getan?

weil die Ladung nicht reichte	21
um sich mit mehr Puffer sicherer zu fühlen	4




11. Interessieren Sie sich für das Konzept, dass E-Fahrzeuge mit einer flexiblen Ladesteuerung die Belastungen im Stromnetz entschärfen können?

ja	59
ja, bin aber skeptisch bei der Umsetzung	21
nein, es besteht kein Bedarf	22
nein, dies ist nicht zuverlässig genug	2




12. Haben Sie seit Beginn des QuickOpt-Projekts (Q2 2024) eine Unannehmlichkeit bei der Nutzung der Poolflotte festgestellt?

ja, kritisch	0
ja, lästig	4
nein, aber davon gehört	15
nein, nie	84
Other	1




13. Würden Sie sich wohl fühlen, wenn Sie während des Buchungsvorgangs zusätzlich die ungefähre Fahrtstrecke (Distanz) eingeben müssten?

ja	32
ich weiss diese vorher nicht genau	42
nein	29




14. Würden Sie in einem solchen Fall einfach den Maximalwert eingeben, um 100% SoC (Ladezustand) zu haben oder würden Sie sich mit einem entsprechenden SoC basierend auf ihrem eingegebenen Wert mit einem Sicherheitspuffer wohlfühlen?

immer 100% vollgeladen	25
Fahrtstrecke + Puffer reicht mir	78




15. Wären Sie damit einverstanden, Ihre Fahrten mit einem SoC von 80-90% zu beginnen?

ja	96
nein	7




16. Hätten Sie das Gefühl, dass dadurch die Häufigkeit von Schnellladungen zunehmen würde?

ja	19
nein	29
ich weiss es nicht	54




17. Welche Mindestreichweite würden Sie als notwendig erachten, um mit einer solchen Konfiguration sicher unterwegs zu sein? (inkl. Puffer)

● 100 km	31
● 150 km	49
● 200 km	23




18. Fahren Sie selbst ein Elektroauto?

● ja	17
● nein	86



19. Könnten Sie sich vorstellen, ein Auto nur zu bestimmten Stunden des Tages aufzuladen und ausnahmsweise aufzuladen, wenn Sie die Ladung durch eine mobile App auslösen?

● ja	12
● nein	5



20. Könnten Sie sich vorstellen, ein Auto nur dann aufzuladen, wenn Sie aktiv eine Abfahrtszeit festlegen?

● ja	9
● nein	8

