



Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

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

Swiss Federal Office of Energy SFOE
Energy Research and Cleantech Division

Final report dated 10. December 2020

aliunid – when energy breathes

field test 2019-2020





Date: 10. December 2020

Location: Bern

Publisher:

Swiss Federal Office of Energy SFOE
Energy Research and Cleantech
CH-3003 Bern
www.bfe.admin.ch

Subsidy recipients:

BFH (Berner Fachhochschule)
Departement Technik und Informatik
Höheweg 80, 2502 Biel
www.bfh.ch

HSG (Hochschule St. Gallen)
Abteilung öffentliches Wirtschaftsrecht. IFR-HSG
Varnbuelstrasse 19, 9000 St. Gallen
www.unisg.ch

Empa
Urban Energy Systems Lab
Überlandstrasse 129, 8600 Dübendorf
www.empa.ch

ETH Zürich
Power Systems Laboratory
Physikstrasse 3, 8092 Zürich
psl.ee.ethz.ch

Authors:

Till Richter, aliunid AG, till.richter@aliunid.com
Dr. Dimitris Kouzoupis, aliunid AG, dimitris.kouzoupis@aliunid.com
Matthias Meier, aliunid AG, matthias.meier@aliunid.com
Tobias Rothen, BFH and Appmodule AG, tobias.rothen@appmodule.net
Anna Bossard, aliunid AG, anna.bossard@aliunid.com
Dr. David Thiel, CEO aliunid AG, david.thiel@aliunid.com
Prof. Dr. Andreas Danuser, CTO aliunid AG and BFH, andreas.danuser@aliunid.com
Dr. Bianca von Preussen, HSG
Dr. Natasa Vulic, EMPA, Natasa.Vulic@empa.ch
Dr. Hanmin Cai, EMPA, Hanmin.Cai@empa.ch
Dr. Stavros Karagiannopoulos, ETH Zürich, karagiannopoulos@eeh.ee.ethz.ch

SFOE project coordinators:

Philippe Mueller, Philippe.Mueller@bfe.admin.ch

SFOE contract number: SI/501841-01

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



Zusammenfassung

Im Jahr 2017 wurde die Energiestrategie 2050 in einer Volksabstimmung angenommen. Sie hat zum Ziel, den Energieverbrauch zu senken, die Energieeffizienz zu steigern und erneuerbare Energien zu fördern. Zudem verbietet sie den Bau neuer Kernkraftwerke. Aufgrund dieser Entscheidung wird erwartet, dass sich die Anzahl der dezentralen erneuerbaren Energiequellen bald stark steigern wird. Dies hat eine stochastischere Energieversorgung zur Folge und kann unkontrollierbare Lastflüsse im Verteilnetz verursachen. Für die Energieversorgung wie auch die Energienachfrage impliziert das unter anderem, dass ein flexibleres Lastmanagement erforderlich sein wird. Die Fähigkeit, Flexibilität im Verteilnetz identifizieren und nutzen zu können wird ein zentraler Faktor sein, um in Zukunft eine sichere Energieversorgung zu gewährleisten.

Um die existierende, teils unflexible Energieversorgung in ein zukunftsfähiges System zu transformieren, strebt man einen Mechanismus an, der es erlaubt, Energie «einzuatmen» (d.h. unregelmäßig produzierten Strom aus erneuerbaren Quellen zu speichern) und «auszuatmen» (d.h. Energie ins Netz einzuspeisen, falls ein Mangel besteht). Dieses Atmen soll einem ganzheitlichen Konzept folgen und eine dezentrale, sichere und flexible Infrastruktur als Grundlage haben, welche auf der Internet of Things (IoT) Technologie basiert.

Das Ziel des Projekts «wenn Energie atmet - Feldtest 2019-2020» ist es, die wichtigsten Aspekte eines solchen atmenden Versorgungsnetzes zu lokalisieren und zu beschreiben und die nötige IoT-Infrastruktur zu entwickeln. Schlussendlich wird die vorgeschlagene Technologie an verschiedenen Use Cases demonstriert.

In diesem Projekt wurden die folgenden Elemente und Methoden genutzt, um die Basis für die Transformation des existierenden Energienetzes zur «Energieversorgung neu» zu legen.

- Ein dezentrales System erfordert eine enge Zusammenarbeit entlang der gesamten Wertschöpfungskette der Energieversorgung. Bei einer ganzheitlichen Betrachtung des Systems können sich dabei neue Geschäftsmodelle¹ entwickeln, welche für alle beteiligten Parteien lohnend sind (vgl. Abschnitt 3.1).
- Es ist möglich, das Verhalten eines Systems nach verschiedenen Kostenfunktionen zu optimieren. Für Smart Grid Anwendungen sind häufig verwendete Kostenfunktionen beispielsweise minimale Kosten, maximaler Eigenkonsum, minimale Emissionen, minimale Spitzenlasten oder optimale Netzauslastung. In diesem Projekt war das Ziel, die CO₂-Emissionen zu minimieren, wobei berücksichtigt wurde, dass sich der CO₂-Gehalt des Strommix mit dem Einsatz der Kraftwerke in der Schweiz und im umliegenden Ausland kontinuierlich ändert. Um die Emissionen zu minimieren, wurde täglich eine 24-Stunden-Prognose des CO₂-Gehalts des Strommix gemacht (vgl. Abschnitt 3.2).
- Um eine Atmende Versorgung auf größerer Skala umzusetzen braucht es einen einfachen und intuitiven Weg, Flexibilität zu beschreiben. Dafür wurde ein Vier-Werte-Modell entwickelt, welches das Energiesystem mittels eines Sets von (hierarchisch angeordneten) «Grid-blocks» (GB) beschreibt. Ein GB repräsentiert ein Subsystem des Stromnetzes. Für jeden GB wird der momentane Konsum und die Produktion (unflexibel) sowie das Ein- und Ausatmen (flexibel) gemessen. Basierend auf diesem Modell wird die verfügbare Flexibilität von jedem Haus in Echtzeit erfasst. Die Anwendung des Vier-Werte-Modells schafft eine gemeinsame Wahrnehmung und Begrifflichkeit eines fragmentierten Systems, ist aber auch entscheidend für den Datenschutz. Die Flexibilität eines GB kann beschrieben werden, ohne Details über den Endnutzer preiszugeben (vgl. Abschnitt 3.3).
- Der Betrieb eines dezentralen Systems kann nur optimiert und gesteuert werden, wenn das Verhalten seiner Bausteine verstanden wird und mit einem relevanten Set von Echtzeitdaten modelliert werden kann. Die verwendete Methode des Model Predictive Control ist ein

¹ Geschäftsmodelle, welche die "Versorgung neu" ermöglichen, sind nicht Teil dieses Berichts.



vielversprechender Ansatz, um das dezentralisierte System als Ganzes, oder einzelne Komponenten individuell, zu steuern (vgl. Abschnitt 3.4).

- Die Grundlage eines atmenden Versorgungsnetzes ist eine robuste, bezahlbare und sichere Dateninfrastruktur, welche es ermöglicht, das Verhalten eines Systems in Echtzeit zu beobachten, zu analysieren und zu steuern. Die IT-Lösung, welche in diesem Projekt verwendet wurde, basiert auf SIOT, einer dezentralen und sicheren Cloud-Infrastruktur, welche in der Schweiz zusammen mit der Appmodule AG entwickelt wurde. Diese Lösung ermöglicht die Verbindung zu zahlreichen Sensoren und Aktoren, um Daten auszulesen und zu verarbeiten (vgl. Abschnitt 3.5 und 3.6).

Als Resultat des Projekts wurde eine dezentrale Infrastruktur bestehend aus 500 Echtzeitdatensystemen aufgebaut. Dies wurde erreicht, indem neue Hardware installiert wurde oder Schnittstellen zu bestehendem Equipment genutzt wurde. Nach Möglichkeit, wurden auch Geräte angebunden, welche flexibel gesteuert werden können (z.B. Wärmepumpen, Warmwasserspeicher und Hausbatterien). Mit dieser Hardware wurde der «Proof of Concept» für ein Atmendes Versorgungsnetz erbracht.

Die folgenden drei Use Cases für Nachfragesteuerung wurden untersucht:

- Wie kann man den Energieverbrauch eines Einfamilienhauses mit PV und einer Hausbatterie optimieren? Verglichen mit den beiden anderen Use Cases bietet die Hausbatterie den höchsten Grad an Flexibilität. Der eingesetzte Model Predictive Controller zeigte vielversprechende Resultate, in dem er zuverlässig den Strom in denjenigen Zeiten vom Netz bezog, in denen die Emissionen niedrig waren (vgl. Abschnitt 4.2).
- Wie kann man Warmwassertanks nutzen, um CO₂-Emissionen zu senken? Bis heute nutzen viele Versorgungsunternehmen eine Rundsteuerung um die Boiler basierend auf einem fixen Zeitplan zu steuern. Indem die Boiler stattdessen während eines optimalen Zeitpunkts in der Zeitspanne von 13 Uhr bis 23 Uhr geheizt werden, kann der CO₂-Fussabdruck der Warmwasserproduktion basierend auf der verwendeten Berechnungsmethode um ca. 25% verringert werden (vgl. Abschnitt 4.3)
- Wie kann man die Wärmepumpen in einem Quartier steuern, um Emissionen zu senken? Wärmepumpen während zwei Stunden pro Tag vom Netz zu nehmen ist eine interessante Möglichkeit um Lastverschiebung auch mit Wärmepumpen ohne Smart Grid Interface durchzuführen. Es konnte gezeigt werden, dass so die Netzbelastung um bis zu 10 kW gesenkt werden kann. Dies entspricht einer Lastreduktion der beteiligten Häuser um 30% in Zeiten einer hohen CO₂-Belastung des Strommix. (vgl. Abschnitt 4.4)

Im Verlauf des Projekts wurde ersichtlich, dass eine genaue Vorhersage der erwünschten Kostenfunktion (CO₂-Gehalt des Strommix in den beschriebenen Use Cases) essenziell ist, um eine wirksame Lastverschiebung zu ermöglichen. Weiter wird die bessere Modellierung des Wärmeverhaltens von Häusern das Flexibilitätpotential erhöhen, da dies dem Regler ermöglicht, Geräte öfter vom Netz zu nehmen, ohne den Komfort des Nutzers zu beeinträchtigen. Dadurch besteht auch bei den CO₂-Ersparnissen und anderen Optimierungszielen ein Steigerungspotential.

Zusammenfassend gesagt hat dieser Feldtest Player aus der gesamten Energiewertschöpfungskette zusammengebracht und die nötigen ökonomischen und technischen Kompetenzen geschaffen, um die Vision eines atmenden Versorgungssystems einen Schritt näher an die Realität zu bringen. Ein ganzheitliches anstelle eines individuell orientierten Energiemanagements ermöglicht es, die Ziele der Energiestrategie 2050 besser zu erreichen. Es besteht noch ein langer Weg, bis alle Häuser in der Schweiz so zuverlässig zusammenarbeiten wie die Zellen einer Lunge. Jedoch konnte in diesem Feldtest die technische Grundlage geschaffen werden, um dies zu erreichen und die vielversprechenden Resultate sind eine Motivation für die weitere Arbeit in diesem Gebiet.



Résumé

En 2017, les électeur·rice·s suisses ont accepté la loi fédérale révisée sur l'énergie, qui vise à accroître l'efficacité énergétique, à promouvoir l'utilisation des énergies renouvelables et qui interdit la construction de nouvelles centrales nucléaires. En raison de cette décision, la quantité de sources d'énergie renouvelables décentralisées devrait bientôt augmenter de manière significative, rendant la production d'électricité plus stochastique et provoquant des flux de charge incontrôlables dans le réseau de distribution. Les défis attendus, tant sur le plan de l'offre que de la demande d'électricité, nécessiteront des mesures importantes, dont le changement de stratégie pour équilibrer les charges. L'identification et l'utilisation de la flexibilité du réseau de distribution deviendront un facteur clé pour fournir à l'avenir une alimentation électrique fiable, à un coût raisonnable.

Pour transformer l'offre d'électricité existante, qui n'est partiellement pas flexible, en un système "à l'épreuve du temps", on envisage un mécanisme qui peut "inspirer" (c'est-à-dire stocker la production d'électricité renouvelable irrégulière en cas de surplus) et "expirer" (c'est-à-dire fournir de l'électricité en cas de pénurie d'énergie). La respiration doit être conçue dans un esprit intégré, être organisée de bas en haut et utiliser une infrastructure décentralisée, sûre et flexible, basée sur l'Internet des objets (IOT).

L'objectif du projet "when energy breathes - field test 2019-2020" est d'étudier et décrire les aspects essentiels d'un système d'approvisionnement énergétique qui respire et de développer l'infrastructure de l'IOT requise. Enfin, la "technologie de respiration" proposée fera l'objet de démonstrations dans différents cas d'utilisation.

Dans les projets, les éléments et méthodes suivants ont été utilisés pour définir les bases de la transformation du réseau énergétique existant en un "nouvel approvisionnement énergétique" :

- Un système décentralisé nécessite une étroite collaboration tout au long de la chaîne de valeur de l'approvisionnement en énergie. Si l'on considère l'ensemble du système de manière globale, de nouveaux modèles commerciaux peuvent évoluer qui rendront cette approche intégrée viable pour tous les partenaires concernés (voir également la section 3.1).
- Il est possible d'optimiser le comportement d'un système en fonction de différentes fonctions de coût (également appelées "tarifs"). Les fonctions de coût communes aux applications des réseaux intelligents sont le coût minimal, l'autoconsommation maximale, l'émission minimale, la charge de pointe minimale ou la qualité du réseau. Le choix de la fonction de coût la plus appropriée dépend des parties prenantes et de l'analyse de rentabilité. Les cas d'utilisation démontrés dans le cadre de ce projet visent à minimiser les émissions en carbone, sachant que sa présence dans l'électricité dépend de l'exploitation des centrales électriques en Suisse mais aussi dans les pays voisins (voir également la section 3.2).
- La mise à l'échelle de la vision d'un système d'énergie respiratoire nécessite une manière simple et intuitive pour décrire la flexibilité. Pour ce faire, un modèle à quatre valeurs a été développé. Il décrit un système énergétique à travers un ensemble de "Grid-blocks" (GB) (disposés hiérarchiquement). Un GB représente un sous-système de la grille. Pour chaque GB, la consommation et la production de courant (non effaçable) ainsi que l'inspiration et l'expiration (charges effaçables) sont mesurées. Sur la base de ce modèle, la flexibilité disponible de chaque maison est représentée en temps réel. L'application du modèle à quatre valeurs ne crée pas seulement une ontologie et une sémantique communes pour un système fragmenté, mais elle est également une manière pour assurer la confidentialité des données. La flexibilité d'un GB est décrite sans révéler aucun détail sur l'utilisateur final (section 3.3).
- Le fonctionnement d'un système décentralisé ne peut être optimisé et contrôlé qu'une fois que le comportement de ses éléments est bien compris et modélisé sur la base d'un ensemble pertinent de données en temps réel. Le contrôle prédictif par modèle est une approche prometteuse pour contrôler soit le système décentralisé dans son ensemble, soit individuellement certains de ses composants (voir aussi la section 3.4).



- La base d'un système d'énergie respiratoire est une infrastructure de données robuste, abordable et sécurisée, qui permet d'observer, d'analyser et de contrôler le comportement du système énergétique en temps réel. La solution informatique utilisée dans le cadre de ce projet est basée sur SIOT, une infrastructure en nuage décentralisée et sécurisée qui est développée en Suisse en collaboration avec Appmodule AG. La solution permet de se connecter à une variété de capteurs et d'actionneurs pour lire et traiter les données énergétiques (voir aussi les sections 3.5 et 3.6).

Le projet a permis de mettre en place 500 systèmes de données en temps réel basés sur une infrastructure décentralisée. Cela a été réalisé par l'installation de nouveaux matériels ou par l'interfaçage avec des équipements existants. Dans la mesure du possible, les dispositifs qui peuvent être utilisés pour une utilisation flexible (par exemple, les pompes à chaleur, les réservoirs d'eau chaude sanitaire et les batteries) ont été interfacés en plus. Ce matériel a été utilisé pour démontrer une preuve de concept d'un système énergétique qui respire.

Les trois cas d'utilisation suivants de la gestion de la demande ont été étudiés :

- Comment optimiser la consommation d'énergie d'une maison à famille unique équipée d'un système photovoltaïque et d'une batterie? Par rapport aux deux autres cas d'utilisation, la batterie offre le plus haut degré de flexibilité. Le modèle de contrôleur prédictif a donné des résultats prometteurs important l'énergie du réseau pendant les temps d'empreinte carbonique faible (voir la section 4.2).
- Comment utiliser les boilers domestiques pour réduire les émissions de carbone? Jusqu'à présent, de nombreux services publics utilisent la commande centralisée pour contrôler ces charges selon un calendrier fixe. En programmant les pompes à chaleur au moment le plus optimal entre 13h et 23h au lieu de 3h du matin chaque jour, l'empreinte CO₂ de la production d'eau chaude a été réduite de 25% en fonction de la métrique utilisée (voir la section 4.3).
- Comment faire fonctionner les pompes à chaleur d'un quartier pour réduire les émissions de carbone? Déconnecter les pompes à chaleur pendant 2 heures par jour est un moyen intéressant pour déplacer les charges en utilisant des pompes à chaleur sans interface de réseau intelligent. Le cas d'utilisation a démontré que la charge du réseau peut être réduite jusqu'à 10 kW. Ceci est équivalent à réduire les charges des maisons considérées de 30 % pendant les périodes de forte intensité carbonique (voir la section 4.4)

Au cours de la validation, il est apparu que la capacité à prévoir, avec précision, le tarif souhaité (intensité en CO₂ de l'électricité pour les trois cas d'utilisation), est cruciale pour un transfert de charge significatif. Dans le cas où les pompes à chaleur étaient contrôlées, il n'a pas été possible d'estimer de manière fiable les économies de CO₂ réalisées en raison de la nature du problème. L'interface prête pour le réseau intelligent facilitera la commande des pompes à chaleur lorsque cette norme sera plus répandue. De plus, une meilleure compréhension du comportement thermique des maisons, augmentera également le potentiel de flexibilité, en permettant au logiciel de contrôle de déconnecter plus souvent les appareils du réseau, sans compromettre le confort des clients. Cette plus grande marge de manœuvre se traduira par des économies de CO₂ plus importantes et donnera plus de flexibilité lors de l'optimisation du réseau pour d'autres objectifs.

En résumé, l'essai sur le terrain a rassemblé des acteurs de toute la chaîne de valeur énergétique et a créé les compétences économiques et techniques nécessaires pour rapprocher la vision d'un approvisionnement énergétique qui respire un peu plus de la réalité. Le fait de gérer l'énergie de manière globale plutôt qu'individuelle contribue à mieux atteindre les objectifs de la stratégie énergétique de 2050. Il reste encore beaucoup de chemin à parcourir avant que toutes les maisons en Suisse ne fonctionnent ensemble de manière aussi fiable que les alvéoles pulmonaires des gens. Toutefois, l'essai sur le terrain a permis de jeter les bases techniques nécessaires à cette fin et a donné des résultats prometteurs qui motivent la poursuite des travaux dans ce sens.



Summary

In 2017, the Swiss electorate accepted the revised Federal Energy Act, which aims to reduce energy consumption, increase energy efficiency, promote the use of renewable energy, while prohibiting the construction of new nuclear power plants. Due to this decision, the quantity of decentral renewable energy sources is expected to significantly increase soon, making electricity production more stochastic and causing uncontrollable load flows in the distribution grid. The expected challenges on both electricity demand and supply will require – among other actions – shifting strategies to balance out loads. Identifying and utilizing flexibility in the distribution grid will become a key factor to safely provide power supply at reasonable cost in the future.

To transform the existing, partially inflexible electricity supply into a “future-proof” system, aliunid envisages a mechanism that can “inhale” (i.e., store irregularly occurring renewable electricity production in case of surplus) and “exhale” (i.e., supply electricity in case of energy shortage). The breathing should be designed with an integrated mindset, should be organized bottom-up and should use a decentralized, secure, and flexible infrastructure, based on the Internet of Things (IoT).

The objective of the project «when energy breathes - field test 2019-2020» is to investigate and describe the essential aspects of an energy supply system that breathes and to develop the required IoT infrastructure. Finally, the proposed “breathing technology” is to be demonstrated in different use cases.

During the project, the following elements and methods were used to lay the basis for transforming the existing energy grid into a “new energy supply”:

- A decentral system requires close collaboration along the entire value chain of energy supply. When looking at the entire system holistically, new business models² can evolve that will make this integrated approach viable for all partners, from the power plant to the utility (see also Section 3.1).
- It is possible to optimize the behavior of a system according to different cost functions (also referred to as “tariffs”). Common cost functions for smart grid applications are minimal cost, maximal self-consumption, minimal emission, minimal peak load, or grid health. Which cost function is more adequate depends on the stakeholders and the business case. The use cases demonstrated within this project aimed at minimizing emissions, considering that the carbon intensity of the electricity depends on the operation of the power plants in Switzerland as well as in the neighboring countries (see also Section 3.2).
- Scaling the vision of a breathing energy system requires a simple and intuitive way to describe flexibility. For this purpose, a four-value model was developed, which describes an energy system through a set of (hierarchically arranged) “Grid-blocks” (GB). A GB represents a subsystem of the grid. For each GB, the current consumption and production (non-shiftable) as well as inhaling and exhaling (shiftable loads) are measured. Based on this model, the available flexibility of each GB is represented in real-time. Applying the four-value model does not only create a common ontology and semantics for a fragmented system, but it is also the key for data privacy. The flexibility of a GB is described without revealing any details of the end user (see also Section 3.3).
- The operation of a decentralized system can only be optimized and controlled once the behavior of its building blocks is well understood and modeled based on a relevant set of real-time data. Model predictive control (MPC) is a promising approach to control either the decentralized system as a whole or some of its components individually. For the purposes of this project, MPC was used to control the behavior of a single-family household (see Section 3.4), while distributed MPC formulations are part of future research.

² Details of a business model that supports a “new energy supply” is not scope of this report.



- The foundation of a breathing energy system is a robust, affordable, and secure data infrastructure, which allows to observe, analyze, and control the behavior of the energy system in real-time. The IT solution used within this project is built on SIOT, a decentral and secure cloud infrastructure which is developed in Switzerland together with Appmodule AG. The solution allows to connect to a variety of different sensors and actuators to read and process energy data (see also Sections 3.5 and 3.6).

As an outcome of the project, 500 real-time data systems based on decentralized infrastructure have been established. This was achieved by installing new hardware or interfacing to existing equipment. Where possible, devices that can be used for flexibility utilization (e.g., heat pumps, domestic hot water tanks and batteries) were additionally interfaced. This hardware was used to demonstrate a proof of concept of an energy system that breathes.

The following three use cases of demand side management have been investigated:

- How to optimize the energy consumption of a single-family house with PV and a battery? Compared to the other two use cases, the battery provides the highest degree of flexibility. The model predictive controller showed promising results, importing energy from the grid when the calculated footprint of electricity was low (see Section 4.2).
- How to use domestic hot water tanks to reduce carbon emissions? Until now, many utilities use ripple control to control these loads based on a fixed schedule. By scheduling boilers during the most optimal time between 1pm and 11pm instead of 3am every day, the CO₂ footprint of warm water production was reduced by 25% according to the adopted metric (see Section 4.3).
- How to operate the heat pumps of a neighborhood to reduce carbon emissions? Disconnecting heat pumps for 2 hours per day is an interesting way to shift loads using heat pumps without smart grid interface. The use case demonstrated that the load of the grid could be reduced by up to 10kW. This is equivalent to reducing the loads of the considered houses by 30% during times of high carbon intensity (see Section 4.4).

During the validation, it became evident that the competence of accurately forecasting the desired tariff (CO₂ intensity of electricity for all three use cases) is crucial for a meaningful load shifting. Moreover, developing a better understanding of the thermal behavior of houses will also increase the flexibility potential by allowing the control software to disconnect more often appliances from the grid without compromising customers' comfort. This increased room for maneuvering will result in higher CO₂ savings and gives more flexibility when optimizing the grid for other objectives.

In summary, the field test brought together players from the whole energy value chain and created the necessary economical and technical competences to bring the vision of an energy supply that breathes one step closer to reality. Performing energy management holistically rather than individually helps to better achieve the targets of Energy Strategy 2050. It is still a long way to go until all houses in Switzerland work together as reliably as the alveoli in people's lungs do. However, the field test laid out the technical basis for this to happen and demonstrated promising results that motivate further work in this direction.



Main findings

- The three investigated use cases showed that shifting loads of residential appliances to times of low CO₂ can be an adequate way to have a positive impact on carbon emissions.
- Close collaboration along the entire value chain of energy supply is needed to bring a decentralized approach into reality.
- Real-time data will be key in the future. The developed hardware proved to be an effective tool for data acquisition and control.



Contents

Contents	10
Abbreviations	14
1 Introduction	15
1.1 Background information and current situation	15
1.2 Purpose of the project	16
1.3 Objectives	17
2 Description of facility	18
2.1 Smart meter integration / aliunid field test	18
2.2 aliunid's demonstration site at BFH	20
2.3 aliunid's demonstration site at localnet.....	21
2.4 GOFLEX infrastructure and the Sion demonstrator	22
3 Procedures and methodology	23
3.1 Holistic system approach and its practical relevance	23
3.2 CO ₂ emissions related to electricity consumption	25
3.3 Grid-block structure and four-value model	30
3.4 Optimal and model predictive control	32
3.5 Real-time data infrastructure	33
3.6 SIOT: Swarm IOT	37
4 Results and discussion	42
4.1 Installation of real-time data infrastructure	42
4.2 Breathing in one building using MPC to minimize CO ₂	46
4.3 Breathing in multiple houses using DHW to minimize CO ₂	53
4.4 Breathing in multiple houses using HP to minimize CO ₂	59
5 Conclusions	62
6 Outlook and next steps	65
7 National cooperation	67
8 Publications	68
9 References	70
10 Appendix	72
10.1 Summary: Governing decentral energy systems	72
10.2 Summary: Identifying critical times and corresponding flexibility options	75
10.3 Summary: Minimal information for collaborative consumption control	80
10.4 Summary: Safe flexibility utilization in active distribution grids	87
10.5 Summary: Requirements customer interface	90
10.6 Summary: Requirements IoT security	99
10.7 Details for breathing experiment in BFH	110
10.8 Details of CO ₂ calculation	114



List of Figures

Figure 1: Top-down versus integrated control.....	16
Figure 2: Breathing electricity supply in a holistic view.	16
Figure 3: aliunid FD on which the aliunid swarm cloud is running (left) and FD-SAT hardware (right). 18	
Figure 4: BFH demonstrator.....	20
Figure 5: Single line diagram household to control the DHWT.....	21
Figure 6: Schema of facility.....	22
Figure 7: Community during the field test (October 2020). Red: utilities, blue: power producers, grey: universities, purple: IoT developer	24
Figure 8: (a) Net imports from Germany [12] and (b) Electricity production profile of Germany [13] on October 2, 2020.....	25
Figure 9: Illustration of the composition of the carbon intensity of electricity consumed in Switzerland.	26
Figure 10: CO ₂ intensity for the year 2019. Imports from neighboring countries lead to an increase during winter.	27
Figure 11: CO ₂ intensity according to equation (3.1) for the year 2019 indicating distinct daily patterns.	27
Figure 12: Forecasted and actual CO ₂ intensity during a 10-day period within the validation set.	28
Figure 13: Timeline of the (near-) real-time CO ₂ intensity algorithm.....	28
Figure 14: Illustration of the four-value model which is used to describe different GB in a consistent way to enable a breathing.	30
Figure 15: Aggregation of the four values in a hierarchical structure of GBs.	31
Figure 16: Smart grid.....	33
Figure 17: Reference model IoT functionality.	35
Figure 18: Central cloud (left) and cloud with edge computing (right).	36
Figure 19: Swarm of private clouds.....	37
Figure 20: PRIC and its «inner life» based on messaging and micro-services.	38
Figure 21: SIOT system.	40
Figure 22: DVPN – inner life.....	40
Figure 23: DVPN operation.	40
Figure 24: DVPN access.	41
Figure 25: Example of a measurement plan for one utility participating in the aliunid field test.	42
Figure 26: Recorded data single house.	44
Figure 27: Recorded data substation.	44
Figure 28: Recorded data transformer station.	44
Figure 29: visualization of the four-value model (a) in the field test app and (b) in the planned release.	45
Figure 30: Block diagram of breathing experiment.	46
Figure 31: MPC versus baseline controller, sunny day.....	48
Figure 32: MPC versus baseline controller, cloudy day.....	48
Figure 33: Measurements of breathing experiment in BFH demonstrator for a period of one week. ...	50
Figure 34: MPC versus baseline controller for a 60-hour window.	51
Figure 35: Illustration of the DHWT switching algorithm.	53
Figure 36: DHWT running times (red) of “house (a)” as extracted through k-means clustering of reactive power.	54
Figure 37: Total active power, background subtracted total active power and derived median active power value of DWHT of “house (a)”.	54
Figure 38: Aggregated power consumption of “house (a)”. White parts correspond to missing data... 55	



Figure 39: Derived standard DHWT power consumption profile for the same period.	55
Figure 40: Synthetic DHWT consumption of “house (a)” shifted according to the experimental setup described in Section 4.3.1.	56
Figure 41: CO ₂ forecast in the period of the simulation.	56
Figure 42: DHWT shifting experiment in “house (a)”. The grey regions represent the periods within which boiler operation was allowed.	57
Figure 43: DHWT shifting experiment in “house (b)”. The grey regions represent the periods within which boiler operation was allowed.	57
Figure 44: Illustration of the HP switching algorithm.	59
Figure 45: Load curves of 10 houses during “breathing with HP” for (a) one day and (b) 7 days.	60
Figure 46: Changes in the accumulated power compared to the previous minute.	61
Figure 47: Average daily consumption (blue lines) and PV production (orange lines) profiles by weekday- weekend (columns) and months representative of the four seasons (rows) with minimum and maximum ranges (shaded areas). The largest and smallest load mismatch between PV production and consumption occurs in January during weekdays and in July during weekends, respectively (red arrows). Regularly scheduled loads are also observed (black arrows). When day-lighting hours are short, an additional evening peak follows the sunset times (shown for January and October, vertical dashed lines).	77
Figure 48: (a) Effective PV consumption (y-axis) at each PV production level (x-axis) without load shifting (grey) and load-shifting windows of specified duration (colors). Remaining unused PV production (hatched), following perfect load-shifting (up to 24 hours) is also shown. (b) Daily effective PV production- to-consumption throughout the year for the 40% annual production case. Minor differences are observed between 12 and 24-hr load shifting windows since the diurnal mismatch is already addressed.	78
Figure 49: Illustration of flexibility envelope adapted from [20] where t_0 denotes next time step. The left figure depicts upper and lower energy bounds that allow derivation of flexible power and corresponding available duration. The right figure maps the time-varying power and duration into 3-dimensional space and the resulting red curve represents a slice of flexibility envelope at t_0	81
Figure 50: Structure of overall experimental implementation. Interactions with system operator are not implemented and are depicted to illustrate the integration into system management.	81
Figure 51: Experimental results with space heating control. The dashed black line indicates the thermal comfort zone.	85
Figure 52: Experimental results with domestic hot water tank control. The dashed black line indicates the thermal comfort zone.	86
Figure 53: Proposed overall methodology for the offline training stage.	88
Figure 54: Methodology for the real-time, online-stage.	88
Figure 55: Baseline and flexibility bounds communicated to GB2.	89
Figure 56: Intelligentes Messsystem.	90
Figure 57: Kundenschnittstelle CII.	91
Figure 58: LED, welche mittels Blink-Impulse Wirkleistungen anzeigt.	92
Figure 59: Mess-Kopf (links), welcher eine optische Schnittstelle nach IEC62056 (rechts) anbindet Zudem gut sichtbar (rot umrandet): 2x LED zur Anzeige von Wirk- und Blindleistung.	93
Figure 60: Kundenschnittstelle umgesetzt mit M-Bus.	94
Figure 61: Kundenschnittstelle gemäss DSMR/P1 (Kopie aus Dokument [28]).	95
Figure 62: SmartMeter CII «ideal» auf Basis Modbus/TCP über WIFI und Ethernet.	96
Figure 63: Anschluss an Inhouse Netzwerke (IHN).	97
Figure 64: Integration bestehende SmartMeter an CII «ideal».	97
Figure 65: Grundmodell für die Aspekte der IT-Sicherheit.	100
Figure 66: Übersicht atmende Versorgung ausgehend vom Prosumer (rechts im Bild).	101



Figure 67: Schematische Darstellung der Funktionalität eines IoT-basierten Smart Grids.	102
Figure 68: 4-Werte-Modell als Basis zur Steuerung eines Energienetzes.	103
Figure 69: Klassischer Cloud-Ansatz (links) versus Schwarm-Ansatz (rechts).	104
Figure 70; Struktur einer Private Cloud PRIC.	105
Figure 71: Physische Struktur von SIOT mit DVPN.	105
Figure 72: Comparison of simulated and measured PV production.	111
Figure 73: (a) Electricity production per category according to Schweizerische Elektrizitätsstatistik 2019 [11]. (b) Electricity production according to the data published on ENTSO-E.	116
Figure 74: (a) Electricity production per category according to [11]. (b) Electricity production according to <i>corrected</i> data from ENTSO-E.	117
Figure 75: Energy production per category, from [11].	118

List of Tables

Table 1: Overview of devices used within the aliunid field test for real-time measurements.	19
Table 2: Equipment description of BFH demonstrator.	20
Table 3: equipment description case “breathing using DHWT”.	21
Table 4: Real-Time Requirements related to smart grids.	34
Table 5: Overview of real-time data infrastructure of aliunid (status 29.10.20).	43
Table 6: CO ₂ emissions and energy consumption for two optimal control problem instances versus baseline controller.	49
Table 7: Pareto front values for simultaneous minimization of CO ₂ emissions and energy consumption from grid.	49
Table 8: CO ₂ emissions and energy consumption in closed-loop.	52
Table 9: Results of the DHWT load shifting simulation.	56
Table 10: Results of the DHWT load shifting experiment.	58
Table 11: Summary of the measured CO ₂ emissions during the experiment.	61
Table 12: Next Steps regarding aliunid field test objectives	66
Table 13: Next steps working packages	66
Table 14: National field test partners.	67
Table 15: List of publications.	68
Table 16: List of components included in the emulator of a residential building.	83
Table 17: Bedeutung des Begriffs „Echtzeit“ in Smart Grids	101
Table 18: Comparison of production of run-of-river plants as published on ENTSO-E with the mean expected production according to the SWKS [SWKS2019].	117
Table 19: CO ₂ intensity in gCO ₂ eq/kWh per production type and country used for the calculation of the CO ₂ intensity of electricity consumed in Switzerland (Section 3.2). Empty fields are not required in the calculation.	119



Abbreviations

BAT	Batteries	PV	Photovoltaic
BEV	Battery Electric Vehicle	q	State of Charge
BFH	Bern University of Applied Sciences	Q	Battery Capacity
C	Consumption	QP	Quadratic Program
CE(S)T	Central European (Summer) Time	RES	Renewable Energy Sources
CHP	Combined Heat and Power	RMSE	Root Mean Square Error
CP	Control Points	RP	Result Points
DDOS	Distributed Denial of Service Attacks	SES	Schweizerische Elektrizitätsstatistik 2019
DER	Distributed Energy Resource	SFH	Single Family Household
DHW(T)	Domestic Hot Water (Tank)	SFOE	Swiss Federal Office of Energy
DN(O)	Distribution Network (Operator)	SIOT	Swarm Internet of Things
DR	Demand Response	SL	shiftable loads
DSO	Distribution System Operator	SVM	Support Vector Machine
DVPN	Distributed Virtual Private Networks	TSO	Transmission System Operator
E	Exhale	UWEK	Umwelt- und Energiekommission
ES2050	Energy Strategy 2050	VPN	Virtual Private Networks
ETH	Federal Institute of Technology	ZEV	Zusammenschluss Eigenverbrauch
EV	Electric Vehicle		
EVU	Energy Supplier		
FD(-SAT)	Fog Device (Satellite)		
GB1, 2	Grid Block 1, 2		
HES-SO	University of Applied Science and Arts		
HP	Heat Pump		
HSG	University of St. Gallen		
I	Inhale		
IoT	Internet of Things		
IP	Information Points		
LV	Low Voltage		
MP	Measurement Points		
MPC	Model Predictive Controller		
MV	Medium Voltage		
NC	Net energy consumption		
NSL	non-shiftable loads		
P	Production		
POC	Proof of Concept		
PRIC	Private Cloud		



1 Introduction

Authors: Till Richter, David Thiel, aliunid

1.1 Background information and current situation

To limit the average global warming to less than 1.5 °C, Switzerland decided to transition from a fossil-based to a renewables-based energy system [1]. In 2017, the Swiss electorate accepted the revised Federal Energy Act. This act aims to reduce energy consumption, increase energy efficiency, promote the use of renewable energy, and prohibit the construction of new nuclear power plants [2]. Substituting the current nuclear power with renewable energy sources in the future will make the production more stochastic and lead to uncontrollable load flows in the distribution grid. Decarbonization efforts also include the increased usage of heat pumps (HP) and battery electric vehicles (BEV) which will result in a more volatile electricity consumption [3].

The expected challenges on both electricity demand and supply require shifting strategies to balance out the loads [4]. Without such demand and supply shifting strategies or massive investments in the grid infrastructure, overvoltage, and unbalance issues at noon hours with high solar PV power injects will violate the grid code [5].

Demand management or demand response (DR) is a method to adjust the electrical loads in the distribution grid according to better match the power supply, by controlling flexibility of customers. There are no significant investments in power infrastructure needed since DR is mainly relying on existing equipment to solve local congestion problems. The residential sector represents about 30% of the total electricity consumption in Europe and is therefore an interesting target for demand management [6].

Flexibility in the distribution grid is generated by either increasing or decreasing the power consumption of the controlled households through demand management. Devices that can be used for demand management without impacting the users comfort include HPs, domestic hot water tanks (DHWT), batteries (BAT) and BEV. Furthermore, the curtailment of photovoltaic (PV) production is a discussed strategy [5].

One disrupting technology which recently gained a lot of momentum is internet of things (IoT). IoT describes the network of “smart things” (e.g., objects with embedded sensors and actors) which exchange data over the internet in real-time. The development of IoT lowers the costs for deploying DR and makes a large-scale deployment in households worth considering [6].

Several research and experimentation projects have investigated and demonstrated demand management in the residential sector [6]. Besides research projects, several startup and commercial companies are already active in the field of demand management and energy management systems. Home energy management systems are targeting end customer needs to optimize self-consumption. Other commercial solutions in the field of demand management are mainly addressing needs of commercial customers (B2B) or provide solutions to manage ZEV (Zusammenschluss Eigenverbrauch). One important use case is the optimal control of charging infrastructure in a neighborhood, to intelligently schedule charging of different cars and avoid load peaks. Existing commercial solutions rely almost exclusively on central cloud computing (e.g., AWS, Azure, IBM or many other), which raises security and privacy concerns.

Technological advancements in demand management and IoT technologies allow now for demand management solutions that target and optimize the distribution grid [7]. A commercial solution along the entire value chain of the energy system, which is based on a distributed IoT system and can interface existing infrastructure is so far not available (or not known to aliunid at the time of writing).



1.2 Purpose of the project

The traditional electricity supply is top-down controlled and relatively inflexible (Figure 1). The vision of aliunid is a "breathing electricity supply", which can "inhale" (or store irregularly occurring renewable electricity production in case of surplus) and "exhale" (or supply electricity in case of energy shortage). Energy that breathes is not limited to the optimization of the individual end customer, but is oriented to the overall system from household, via the grid, to the power plant. The goal is a system that can measure, analyze, and influence energy flows in real-time (e.g., every 1 to 5 seconds). The vision of a breathing system is not limited to electricity and will be extended in the future to gas, heat, and hydrogen (sector coupling).

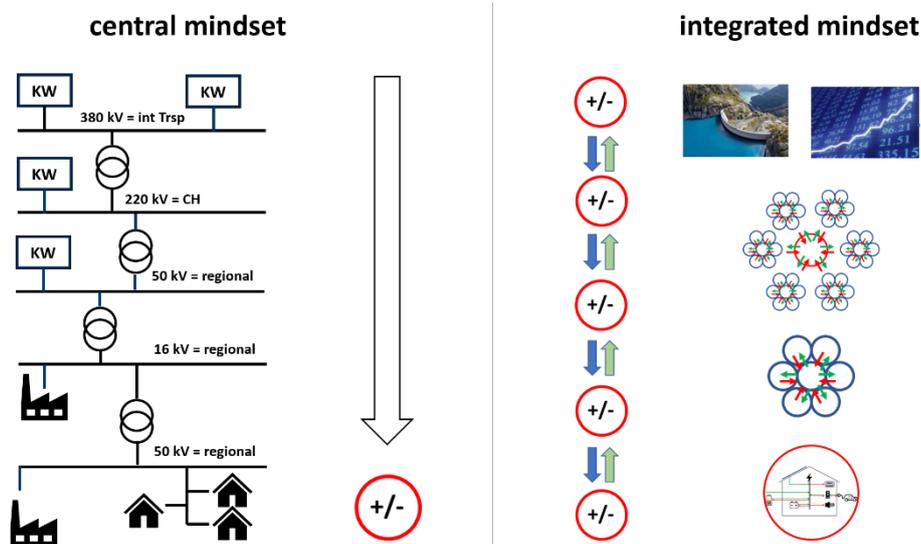


Figure 1: Top-down versus integrated control.

A new, climate-friendly energy supply should be able to be implemented safely and affordably through the envisioned breathing technology. The implementation of an energy system that breathes requires new approaches as shown in Figure 2. The "central mindset" of top-down control from the power plants down through the grid to the end customer in the city is to be replaced by an "integrated mindset": Autonomous energy systems produce in a decentralized fashion and regulate themselves bottom-up. The superordinate system levels will be accessed only in the event of shortages or surpluses. The respective superordinate system level serves as a virtual battery and supplies or obtains the energy and flexibility required or offered by the underlying level in real-time. At the top level, Swiss pumped storage power plants and stock exchange access ensure that the aggregated overall system can "breathe".

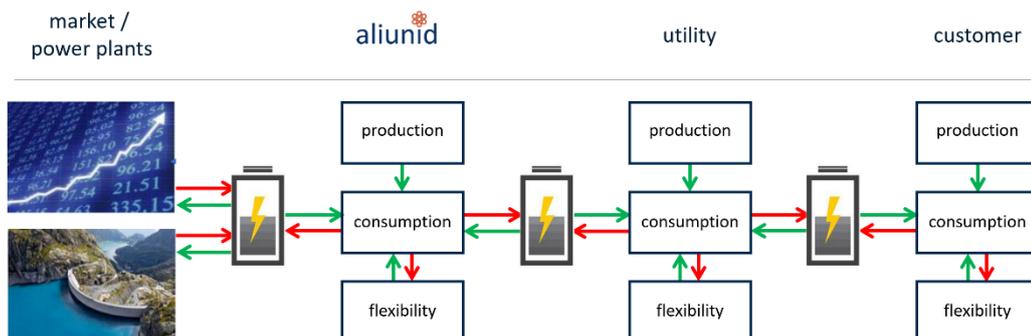


Figure 2: Breathing electricity supply in a holistic view.



1.3 Objectives

The main objective of the project was to conduct a nation-wide field test that investigates the potential of a new energy supply that “breathes”. In the period from April 2019 to November 2020, this field test was carried out at around 20 utilities throughout Switzerland. This document summarizes the outcomes of aliunid’s work in the last two years. The aim of the field test is to show that the ability to breathe is relevant to the feasibility of the energy strategy 2050 (ES2050) and that it can be implemented safely and efficiently.

The main tasks of the project were divided into different work packages that were identified as prerequisites for an energy system that breathes. A short summary of these work packages, which are submitted as separate reports, can be found in the appendix (Section 10).

To make the scope of the project more concrete, the consideration of three use cases was agreed upon, each addressing one of the following questions:

- How to optimize the energy consumption of a single-family house with PV and a battery?
- How to use domestic hot water tanks to reduce carbon emissions?
- How to operate the heat pumps of a neighborhood to reduce carbon emissions?

These use cases have been documented in Chapter 4 and their findings are further discussed in Chapter 5. Furthermore, the work of aliunid aims to address the following aspects:

- Representation of the flexibility of a household using a suitable model in real-time (see Sections 3.3 and 4.1.3).
- Calculation of a real-time tariff for which the flexibility of a household or a distribution grid shall be optimized (see Section 3.2).
- Competence for robust, affordable, and secure acquisition of real-time data (see Section 3.5).
- Organization of data flows in a modern decentralized swarm cloud (see Section 3.6).

As far as flexibility is concerned, the report documents the effective inhaling (storing energy) and exhaling (releasing energy back to the grid). The actual measured values from the three breathing use cases are compared to each other. Based on these findings, considerations are made on how to commercialize breathing solutions. Whether a breathing system is viable and feasible is further discussed in Section 10.2.

In summary, aliunid plans to scale the findings of the conducted experiments to estimate the potential of an energy system that breathes. This shall answer how such a system can contribute to the goals of the ES2050 (see Section 10.2).



2 Description of facility

Author: Till Richter, aliunid

2.1 Smart meter integration / aliunid field test

One core competence of aliunid is built around flexible and inexpensive data acquisition. The underlying hardware and software infrastructure is further described in Sections 3.5 and 3.6. This section explains which hardware is currently used to get real-time load data.

Figure 3 shows the heart of aliunid's real-time data acquisition infrastructure: The fog device (FD) and the fog device satellite (FD-SAT). The FD is a more powerful client compared to the FD-SAT. Both devices serve in aliunid's system architecture as a gateway to different sensors and actors. In the field test, aliunid is currently using the approaches listed below. Table 1 gives a detailed overview of which devices are used within the aliunid field test.

- Meter behind the meter: aliunid uses 3 phase energy meters from Carlo Gavazzi which are installed after the smart meter from the utility and connected via Modbus to the FD. While this meter is providing high quality data, the installation must be carried out by a certified electrician which makes scaling more difficult and expensive.
- MeterSensor: aliunid developed an optical plug and play sensor, which can interface with the infrared port of smart meters. While this interface is universally available on the smart meters, the field test showed that most utilities have either deactivated or encrypted this port, which makes deployment more challenging.
- BlinkSensor: To overcome the challenges of the MeterSensor aliunid developed a basic impulse counter that can read energy consumption via the S0 interface, available in every smart meter.
- To quantify inhaling and exhaling in the house, aliunid developed interfaces to different PV inverters, battery inverters and charging infrastructure for electric vehicles.

Section 10.5 gives a more detailed overview on customer interfaces and lists the advantages and disadvantages of the different approaches. One of aliunid's key competences is interfacing many different devices and bringing them into a reliable, secure cloud environment.

Besides interfacing all kind of different devices, aliunid has also developed an in-depth experience on how to best connect FD or FD-SAT to the internet. This is not a trivial task since WLAN and LTE connection is often weak in the basement, where most smart meters are installed. While using the customer's WLAN is a cheap solution, it is not always reliable, and connection issues may occur if the customer decides to, e.g., temporarily disconnect their WLAN router from the power supply. On the other side, LTE may be more stable, but it is also more expensive due to the recurring costs of the data service.



Figure 3: aliunid FD on which the aliunid swarm cloud is running (left) and FD-SAT hardware (right).



Table 1: Overview of devices used within the aliunid field test for real-time measurements.

Brand and Model	Description	Real-time data (overview)
Carlo Gavazzi EM340 EM330 EM210	<ul style="list-style-type: none"> - Direct 3-phase measurement, 65A per phase - Converter measurement, 3-phase, 100A/150A/200A/250A/400A/600A/800A per phase - Measurement with Rogowski coils 4000A per phase 	<ul style="list-style-type: none"> - Frequency - Consumption / production - Active power - Reactive power - Apparent power - Voltage - Current <p>Available for all three phases</p>
MyStrom Plug	WLAN adapter plug	<ul style="list-style-type: none"> - Active power - Consumption - On / off switching - Temperature
Sonnenbatterie eco8.0	Rechargeable lithium-ion battery. Intended for home energy storage, to maximize self-consumption of solar energy	<ul style="list-style-type: none"> - Consumption / production - Voltages - State of charge <p>Automatic / manual control of charging and discharging</p>
SolarEdge SExx and all other Inverters using the SunSpec protocol	PV Inverter	<ul style="list-style-type: none"> - Production - Current - Voltage - Frequency
SolarLog	Energy management system for renewable energy that can interface with many PV inverters	Depends on what is connected to the SolarLog
Shelly 2.5 Shelly 1PM Shelly 3EM Shelly Dimmer	Small-sized sensor and relay with WLAN connectivity	<ul style="list-style-type: none"> - Active power - Consumption - On / off switching



2.2 aliunid’s demonstration site at BFH

The demonstrator of aliunid at the Bern University of Applied Sciences (BFH) is used to perform breathing experiments in a single household. Goal of the experiments is to demonstrate how the developed breathing technology can help to reduce CO₂ emissions³. The core of the control algorithm consists of a numerical optimization routine which continuously calculates the optimal trajectories of the shiftable load over the next 24 hours. A brief introduction to optimal control for the reader to be able to follow the results of Section 4.2 is provided in Section 3.4 and a detailed description of the experiment in the Appendix (Section 10.7).

The high-level structure of the demonstrator is shown in Figure 4. A lithium battery establishes a connection between the solar panel installation on the roof of the building, the electricity grid, and the lab. A proprietary battery controller maximizes self-consumption by using the produced energy to primarily cover the electricity needs of the lab and charge the battery. Excess of electricity is sold to the grid operator at a fixed tariff. When energy production is not enough to cover the consumption of the lab, electricity is purchased from the grid based on a standard dual-tariff system.

The implemented control algorithm bypasses the automatic mode of the battery controller and minimizes CO₂ emissions by forecasting PV production, fixed load consumption and CO₂ content of the purchased energy. More precisely, the software decides how to optimally charge and discharge the battery to ensure that any external energy demand is covered during periods of low CO₂ footprint. More details on the methodology and implementation are found in Sections 3.4 and 4.2. A summary of the equipment is provided in Table 2.

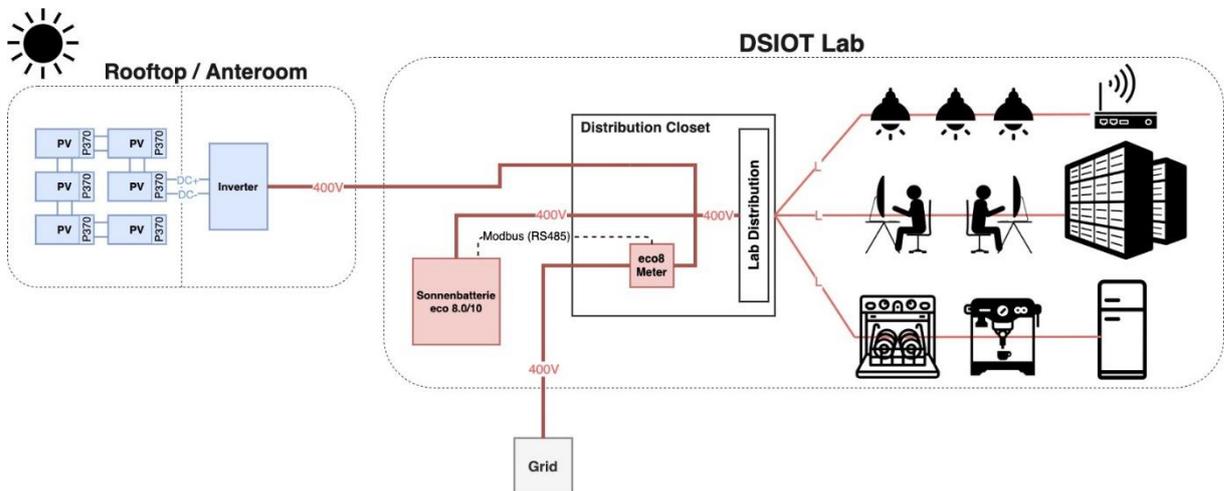


Figure 4: BFH demonstrator.

Table 2: Equipment description of BFH demonstrator.

Type	Description
Flexible load	SonnenBatterie eco 8.0, 10kWh capacity
Fixed load	Computers, hot water tank emulator (1.8kW, 3 hours overnight), lab and kitchen equipment
Energy production	Jinko Cheetah 60M 320Wp (34 panels), 10.25 kW maximum DC power

³ Or electricity costs with only small modifications.



2.4 GOFLEX infrastructure and the Sion demonstrator

In the third use case, aliunid is demonstrating how to control the heat pumps of a neighborhood to reduce CO₂ emissions. For this proof of concept, aliunid relied on existing infrastructure in the customers of Oiken.

Oiken, the local utility in Sion, is one of aliunid's partners and stakeholders. Together with the University of Applied Science and Arts Western Switzerland (HES-SO), Oiken has built the Swiss GOFLEX demonstration site. GOFLEX was an EU funded research project, which developed an integrated solution for cost-effective demand response in distribution grids [9]. For this purpose, gateways were installed, allowing to control big loads at the end consumers' side. Since aliunid does not only have the competence to breath with proprietary IoT hardware, the existing infrastructure in Sion was controlled remotely from aliunid's server via MQTT. The setup is illustrated in Figure 6. aliunid selected 10 houses with heat pumps to perform the breathing experiment. During the experiments over two weeks, the participating houses used between 20 and 60 kWh electricity for heating and warm water. Such a deviation can be explained by the various factors on which heating demand depends on, such as a building's isolation, size etc.

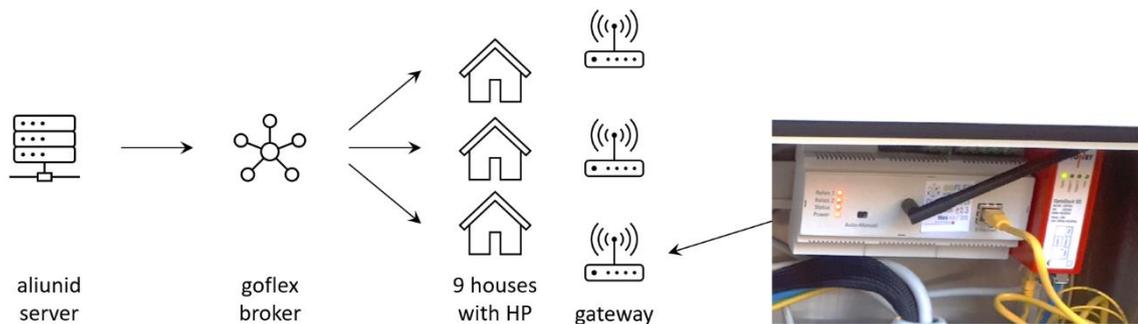


Figure 6: Schema of facility.



3 Procedures and methodology

3.1 Holistic system approach and its practical relevance

Author: David Thiel, aliunid

A breathing energy supply addresses the entire system of energy supply in its technical, economic, and socio-economic diversity. To cope with the complexity involved, the field test "when energy breathes" is based on the following principles:

Community along the entire value chain:

The extent of the expertise required to implement an energy supply that breathes has led aliunid to build up a community of currently 16 energy companies, power producers, IoT developers and leading universities. Figure 7, Section 7 and Table 14 describe this community in more detail, which ultimately includes:

- All levels in the value chain of the Swiss electricity industry (except TSO) from the flexible pumped storage power plant to the local energy service provider. Since EVU / VNB within the community are often involved in the cross-commodity networks, the aspect of sector coupling is also addressed.
- Companies of different sizes, from international energy traders, to metropolitan municipal utilities, to agile local utilities.
- Companies in the German-, French- and Italian-speaking parts of Switzerland. For aliunid's innovation work, the different ways of thinking and working are extremely valuable.

System approach:

The implementation of the Energy Strategy 2050 requires new, holistic approaches. These are not limited to the Internet of Things (IoT) but require a general rethink of the entire system of energy and network industry, real-time data systems, regulation, product design and new business models. Such a change will not be successful without the inclusion and consent of end users. New, digital electricity products must bring real benefits, which are widely understood and accepted.

Utilities are generally in public ownership. Therefore, the actions of a utility are determined by the interaction of technical, economic, and political influences. They always have a "legacy", which ranges from rolled out electricity meters, to cross shareholdings, to long-term electricity contracts and their own small power plants. The implementation of a climate-friendly, secure, and affordable energy supply will only be possible if solutions can be found that persuade most of the stakeholders.

Practical relevance by solving real existing problems:

The close collaboration with different parties along the value chain motivated innovation close to the customer and addressed existing problems of producers, utility companies and their customers. The direct and continuous interaction with end customers is intensive and instructive. Even if many problems are still unsolved, a high sensitivity for the practical relevance of solutions has been developed. For example, a breathing energy supply is only possible if inexpensive solutions can be found for the secure handling of real-time energy data throughout Switzerland. Digital electricity products only have a chance of being implemented if the energy supply companies can see added value in them. A demand side management solution alone will not be sufficient to have an impact: The technology must be embedded in a viable business case and tailored to smoothly fit into the existing ecosystem, while changes in the regulatory framework and market rules are also crucial (see Section 10.1). If the complexity of a solution is too high, it will hardly be relevant in practice. The challenge lies in building a simple enough system, which can connect the protected area of the grid control center with the home energy management systems in the basement of Swiss homes. Only when the end customers embrace the breathing energy supply and order the related products, the technology will become relevant.

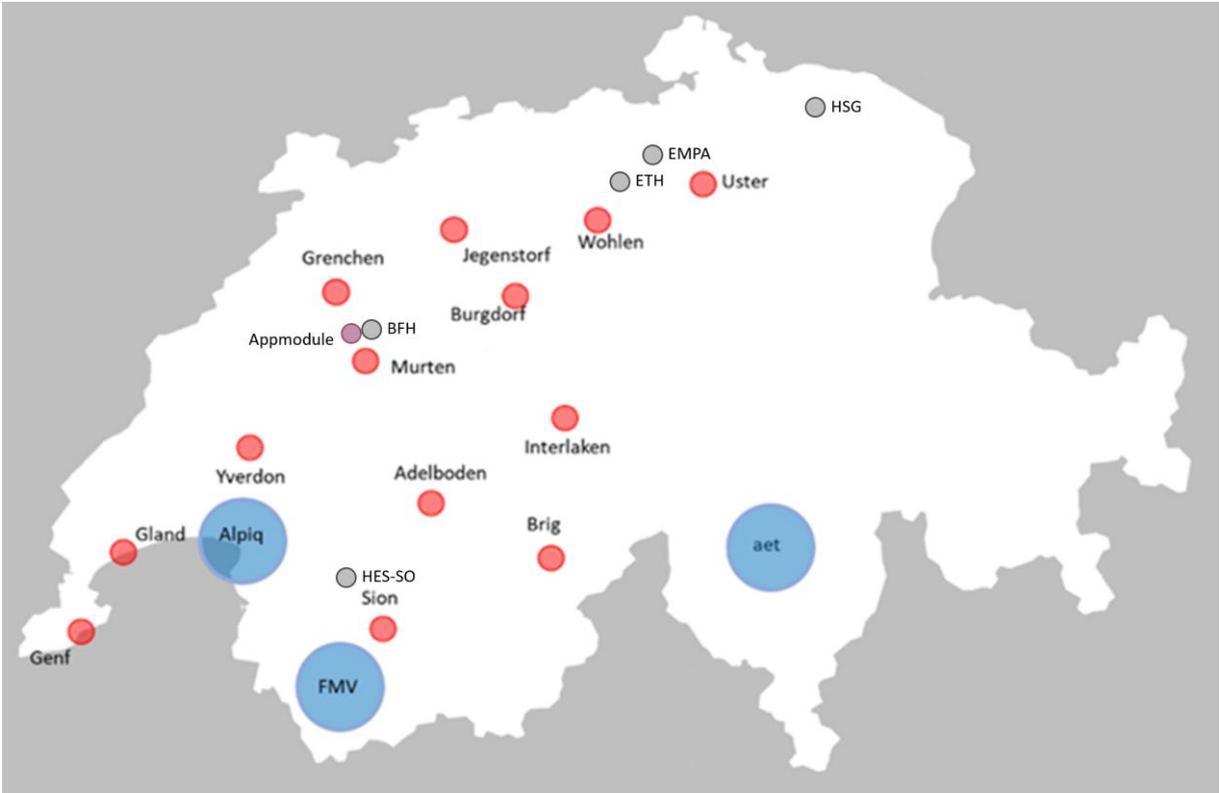


Figure 7: Community during the field test (October 2020). Red: utilities, blue: power producers, grey: universities, purple: IoT developer



3.2 CO₂ emissions related to electricity consumption

Author: Matthias Meier, aliunid

3.2.1 Motivation

Electricity production always involves carbon emissions⁴. The amount of emissions however highly depends on the type of production plant. While the (life cycle) emissions of CO₂ equivalents of hydro power plants and nuclear power plants may lead to emissions of only around 10 gCO_{2eq} per kWh of produced electricity, a coal power plant may emit well above 1000 gCO_{2eq} per kWh [10]. Knowledge and prediction of the varying CO₂ emissions related to electricity consumption can therefore potentially help to minimize the carbon footprint of consumers.

Since Switzerland's domestic electricity production is mainly composed of hydro- and nuclear production [11], the corresponding carbon emissions are rather low. This view does however not represent the carbon emissions that are caused by the electricity *consumption* of Switzerland. This is due to the fact that - especially in winter times - Switzerland strongly depends on imports from its neighboring countries, e.g., Germany, which produces a large fraction of its electricity with carbon intensive coal power plants.

Imports do however not only fluctuate on a seasonal scale but also on a daily scale. And so does the composition of electricity with respect to generation source. Figure 8 (a) shows the net import profile at the German border for a day in October. Figure 8 (b) shows the electricity production profile of Germany for the same day. The latter nicely shows how the shares of solar (yellow) and wind energy (light blue) - and consequently the share of coal - vary strongly throughout the day.

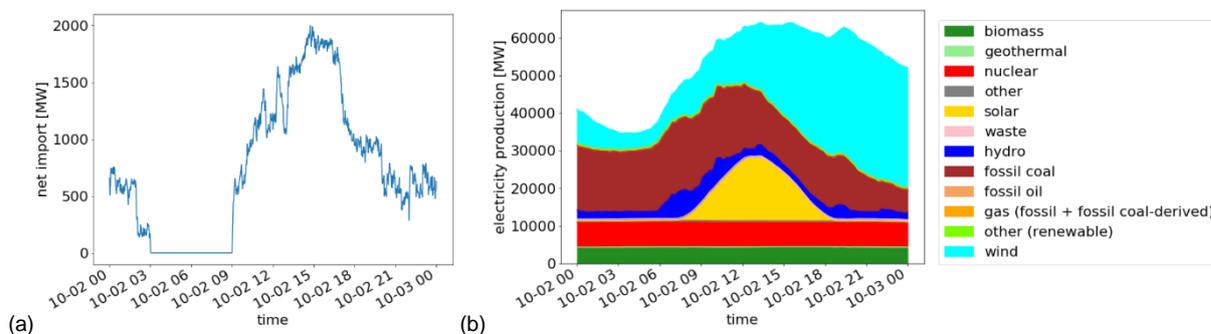


Figure 8: (a) Net imports from Germany [12] and (b) Electricity production profile of Germany [13] on October 2, 2020.

Therefore, from a CO₂ reduction perspective, it not only matters how much electricity is consumed but also *when* electricity is consumed. By providing both a real-time calculation of the carbon impact of electricity consumption as well as the ability to flexibly shift electrical loads, aliunid may help to reduce the demand of carbon intensive electricity.

3.2.2 Calculation approach for the CO₂ intensity

As indicated in the previous section, the carbon intensity⁵ (gCO_{2eq}/kWh) of the consumed electricity in Switzerland is a combination of the carbon emissions associated to domestic production and carbon emissions associated to imported electricity. This is illustrated in Figure 9.

⁴ The terms “carbon”, CO₂ and CO_{2eq} are used interchangeably to refer to CO₂ equivalents in this report.

⁵ The term “carbon intensity” has established itself in literature for quantities with the unit [gCO_{2eq}/kWh] and has no connection to the physical quantity “intensity” [power/area].

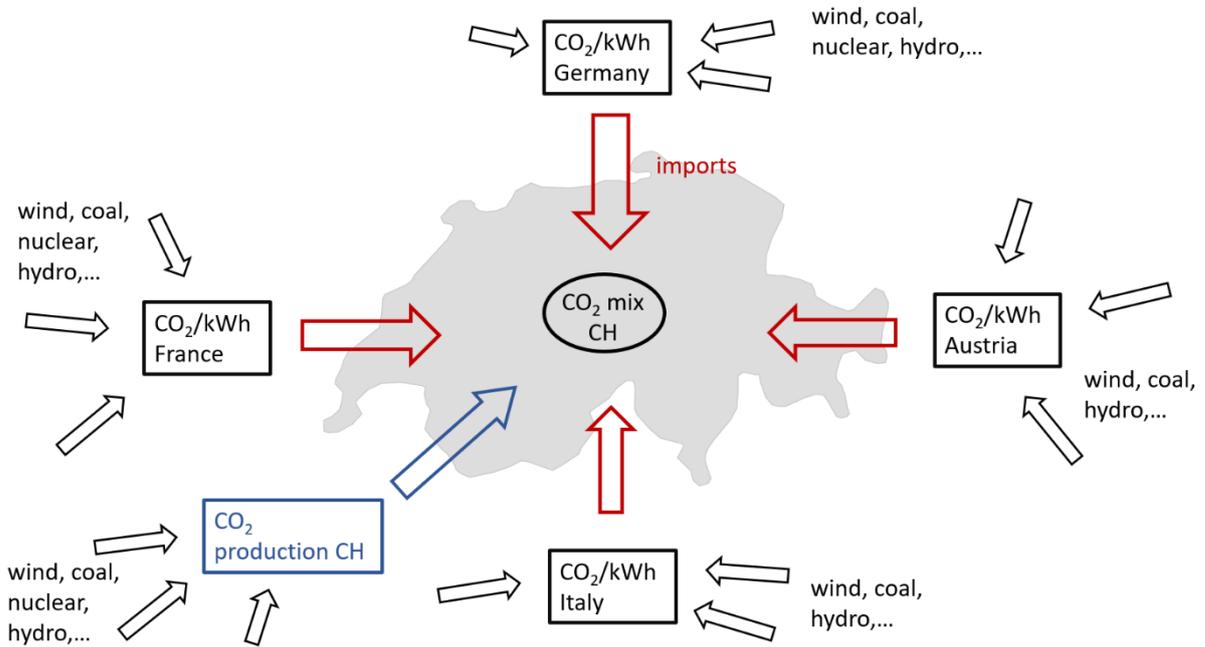


Figure 9: Illustration of the composition of the carbon intensity of electricity consumed in Switzerland.

The exact way how to account for imported electricity is however ambiguous [14] [15]. One of the ambiguities stems from the fact that cross-border flows at the Swiss borders are not only motivated by the electricity requirements of Switzerland, but Switzerland also strongly acts as a transfer platform for electricity trade between neighboring countries, e.g., between Germany and Italy. It is therefore not obvious what part of carbon emissions shall be accounted to Switzerland. These ambiguities as well as details about the chosen approach are discussed in the Appendix 10.8.

The equation used to calculate the carbon intensity of electricity is the following:

$$x_{CH} = \frac{\sum_m I_{CH,m} \cdot production_{CH,m} + \sum_j x_j \cdot import_{CH,j,net}}{(\sum_m production_{CH,m} + \sum_j import_{CH,j,net})} \quad (3.1)$$

where

x_{CH} = CO₂ intensity in gCO_{2eq}/kWh of Switzerland

x_j = CO₂ intensity in gCO_{2eq}/kWh of neighboring country j

$I_{CH,m}$ = CO₂ intensity in gCO_{2eq}/kWh of production type m in Switzerland

$production_{CH,m}$ = production of type m

$import_{CH,j,net}$ = net import from country j to Switzerland (zero if net export)

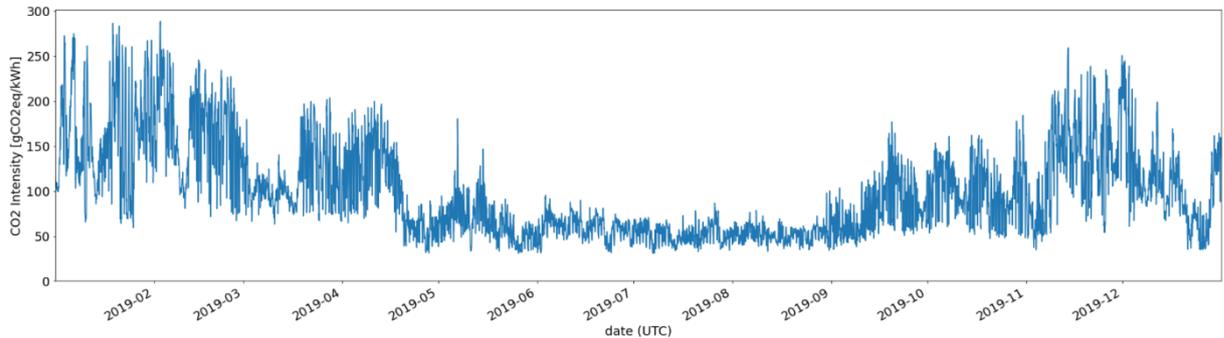


Figure 10: CO₂ intensity for the year 2019. Imports from neighboring countries lead to an increase during winter.

To apply this formula, aliunid needs information about the production and its carbon emissions of Switzerland as well as information about imports from neighboring countries. Furthermore, information is needed about the carbon intensities of the imported electricity. For the latter, we assume that the carbon intensity of imported electricity is equal to the carbon intensity of the electricity *produced* in the corresponding country⁶, i.e.

$$x_j = \sum_m I_{j,m} \cdot production_{j,m} \quad (3.2)$$

where

$I_{j,m}$ = CO₂ intensity in gCO_{2eq}/kWh of production type m of country j .

The relevant data of electricity production per production type used for Equations (3.1) and (3.2) can be retrieved from ENTSO-E [13] and the data for cross-border flows can be retrieved from [12]. The ENTSO-E production per type data from Switzerland however is incomplete. To compensate for this, aliunid adjusts the data as described in Section 10.8.3.

The carbon emissions per production type are taken from different sources in the literature⁷. The resulting carbon intensity fluctuates between around 50-100 gCO_{2eq}/kWh in summer up to almost 300 gCO_{2eq}/kWh in winter (Figure 10). Additionally, distinctive daily patterns are observed (Figure 11).

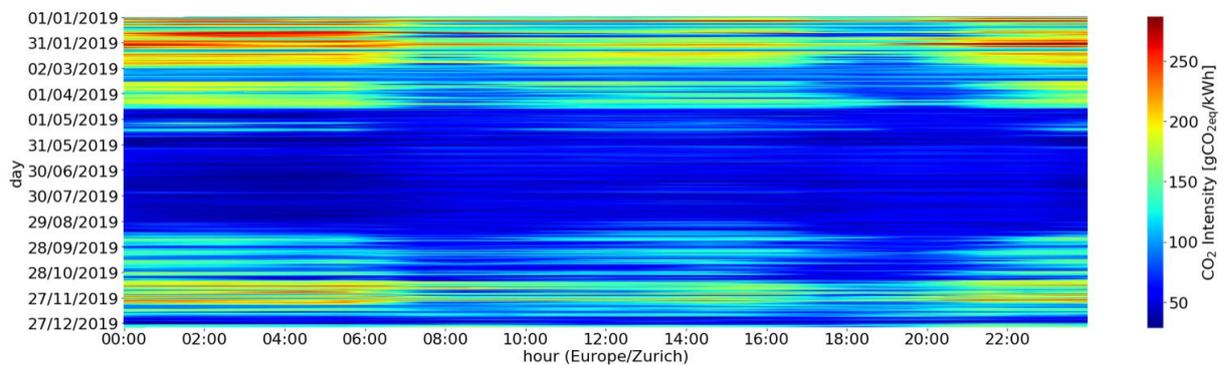


Figure 11: CO₂ intensity according to equation (3.1) for the year 2019 indicating distinct daily patterns.

⁶ For details, see Appendix 0.

⁷ For details, see Appendix 10.8.4.



3.2.3 Forecast of the CO₂ intensity

The chosen forecast model is an auto-regressive model with external input and is trained on historical data. The external input consists of the day-ahead aggregated electricity generation forecast published on ENTSO-E for the countries CH/IT/DE/FR/AT.

To get appropriate coefficients, the model is fitted to data from January 2016 to August 2019 and validated on data from September 2019 to July 2020 (80% and 20% of the available data respectively). The resulting model reaches a coefficient of determination R^2 of 0.81 in the training set and 0.65 in the validation set. Its symmetric mean absolute percentage error in the validation set is 19.3%. Figure 12 shows the result of the forecast in a 10-day period within the validation set. Most importantly for CO₂ optimization purposes, the locations of peaks and valleys are forecasted quite reliably.

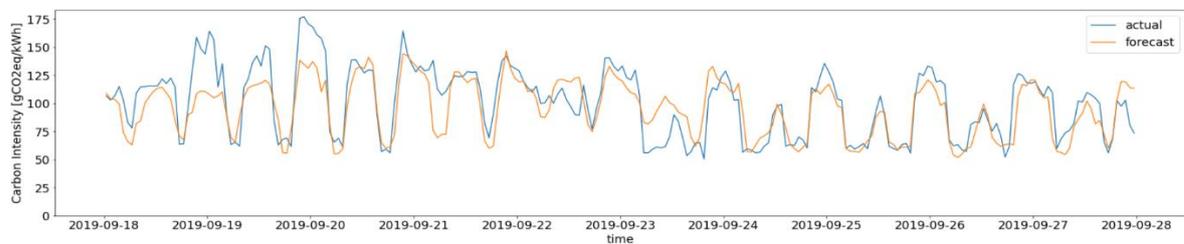


Figure 12: Forecasted and actual CO₂ intensity during a 10-day period within the validation set.

3.2.4 Implementation in (near-) real-time

To calculate the CO₂ intensity of Switzerland and its forecast, aliunid depends on data from ENTSO-E and Swissgrid. The electricity data from ENTSO-E is published with a delay of at least two hours and the Swissgrid data is published with a delay of 20 minutes. The closest to real-time one can achieve is therefore a calculation with a delay of two hours. The ENTSO-E day ahead generation forecast which is required for the CO₂ forecast used in this report is published once a day at 18:00 (CET/CEST) for the next day. This means that with the used model aliunid can make a forecast once a day after 18:00.

The simplified procedure of the algorithm used for load shifting optimization is illustrated in Figure 13. Every day after 18:00 (CET/CEST), when the day-ahead electricity forecast is published by ENTSO-E, the database of the actual CO₂ intensity is updated (according to equation (3.1), 'History date_x') and the CO₂ forecast for the next day is calculated ('Forecast date_{x+1}').

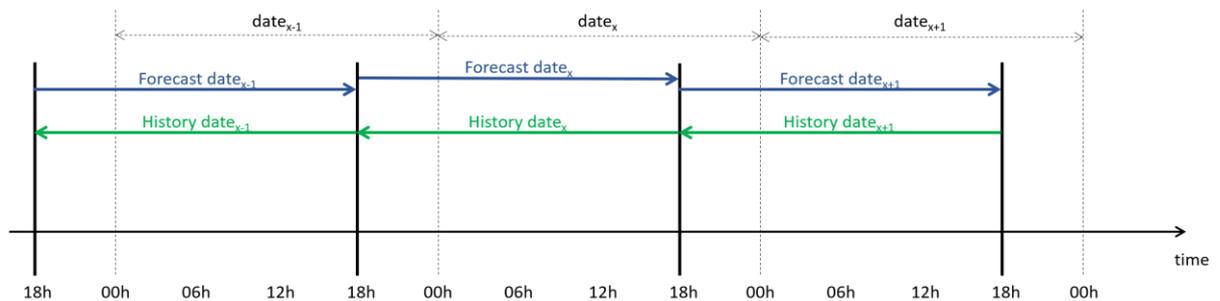


Figure 13: Timeline of the (near-) real-time CO₂ intensity algorithm.



3.2.5 Effect of load shifting on CO₂ mitigation

In a future scenario where electricity from renewable sources is abundant and capable of covering all the demand during certain periods of the day, the benefit of load shifting is obvious. As long as renewables, even at their peak times, can only cover a fraction of the electricity demand, the effect of load shifting on CO₂ mitigation is however not as obvious as it might seem at first glance. This is due to the limited coupling between consumption and production. In general, three broad timescales and mechanisms can be identified [16]:

- Short term balancing impact
- Systematic energy trading impact
- Long-term infrastructure impact

If the consumption deviates from the predicted level due to unexpected shifting of demand, this will – provided that the deviation is significant – lead to balancing actions of the grid operator. The impact of this balancing on carbon emissions can vary from positive to negative – depending on the specific actions and the type of power plants involved.

If the load shifting is however realized systematically, the altered demand profile will start having an impact on energy trading. The nature of the impact then depends on the precise mechanisms of the market. In a simplified model where power plants - due to fuel costs and carbon taxes - have increasing marginal costs with increasing carbon intensity, carbon intensive plants with the highest marginal costs can be pushed out of the market if during times of elevated demand and low shares of renewables the demand is decreased by load shifting (merit-order effect). Carbon intensive power plants will not be able to compensate for this if the demand is instead increased during times of low demand and high shares of renewables since they will not be able to enter the market under these conditions.

Finally, this will lead to a lower rentability of carbon intensive power plants and will influence the infrastructure in the long-term by making them less attractive for investment.

The simplified model just described has however its weaknesses – especially in Switzerland. While low-carbon technologies like wind, solar, hydro run-of-river and nuclear may indeed have lower marginal costs than fossil plants, this is not the case for hydro storage. However, the CO₂ metric used for the scope of this report (Section 3.2.2) does not reflect this fact. It may therefore favor load shifting towards times of high demand, naturally correlating with high shares of electricity from hydro storage, which may then unwantedly favor fossil plants at the same time. Furthermore, with the current level of carbon pricing, the order of fossil plants according to marginal costs does not exactly match their order according to carbon intensity – a fact that further complicates the assessment of the impact of load shifting on carbon emissions.

Therefore, for the special purpose of load shifting, it might in principle be desirable to define a metric that is specifically tailored to maximize the impact on carbon emissions – as opposed to a metric only reflecting the carbon emissions related to the current production.



3.3 Grid-block structure and four-value model

Author: Till Richter, aliunid

Scaling the vision of a breathing energy system described in Section 1.2 requires a scalable system architecture, considering that thousands of houses would be connected to the system. For this purpose, aliunid is proposing to introduce a Grid-block (GB) structure. Between the different GBs four different values are exchanged. The patent of this invention is pending at the time of writing [17].

The four-value model of aliunid is the most simple and intuitive way to describe the flexibility within a GB. The underlying idea is to disaggregate the net flow in four values, as illustrated in Figure 14, by separating non-shiftable from shiftable loads.

- aliunid internally defines that all *shiftable loads* must be measurable and controllable. Shiftable loads that are commonly installed in single-family households are HP, BAT, DHWT and BEV.
- *Non-shiftable loads* cannot be measured or controlled, in particular because of negative impact on the end consumer. For example, the electricity consumption of a stove cannot be shifted, because the end consumer would not be able to prepare food as planned.

Common ontology and semantics are a prerequisite for a scalable solution, especially if a large variety of installed hardware configuration in the field is considered. The four-value model of aliunid is a simple and pragmatic framework to map the complex interdependencies of a breathing energy system in a formal and consistent way. For every GB, the following four values are measured:

- Production: for prosumers, this is almost exclusively the PV power produced on the rooftop.
- Consumption: this is the aggregated power of non-shiftable loads in a GB.
- Inhale: this is the power consumption⁸ of all devices that can be controlled. If a device inhales, it consumes additional electricity. Inhaling can be achieved by either load scheduling of devices like HP, DHWT or charging of BAT and BEV.
- Exhale: if a device exhales, power is flowing back to the grid. Currently exhaling can only be achieved if a GB contains a BAT or a BEV with the option of bi-directional charging.

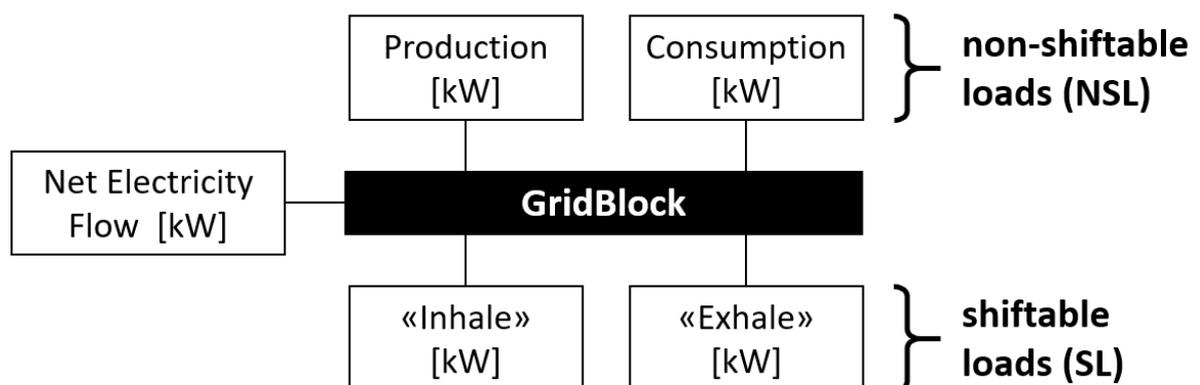


Figure 14: Illustration of the four-value model which is used to describe different GB in a consistent way to enable a breathing.

⁸ If the four-value model is extended in the future to energy, sector coupling will provide additional levers for inhaling and exhaling.



From the law of conservation of energy, it follows that the sum of all four values is equal to the net electricity flow:

$$P_{NEF} = P_P + P_C + P_I + P_E$$

To enable a breathing energy system, these four values are transmitted in real-time (e.g., every 5 seconds) and the way how to optimally breathe is derived. However, the transmitted information also contains sensitive information: the load curve contains information about personal behavior and needs to be protected accordingly.

For this purpose, a hierarchical structure is proposed as shown in Figure 14. The four values of every GB can be summed up. By doing so, the information needed for breathing is transmitted, while drawing conclusions on the personal life of the inhabitants of an SFH is not possible anymore.

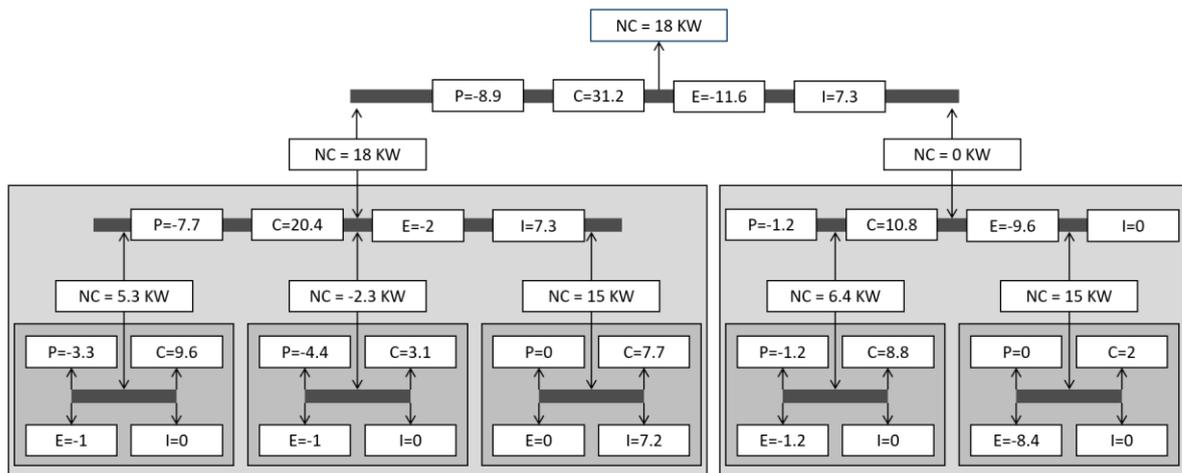


Figure 15: Aggregation of the four values in a hierarchical structure of GBs.



3.4 Optimal and model predictive control

Author: Dimitris Kouzoupis, aliunid

3.4.1 Introduction

Model predictive control (MPC) is an advanced technology for the control of dynamical systems [18]. One of its main advantages over traditional control techniques, such as PID control, is that it offers an intuitive way to formulate the objective to be minimized and the constraints of the problem. This is especially advantageous in applications where the objective is not to follow a given set-point or trajectory (i.e., regulation/tracking problems) but something more “abstract”, such as profit maximization.

In the context of this report, MPC is used to repeatedly compute the optimal behavior of the shiftable loads in 24-hour windows. More details on the implementation can be found in Sections 4.2 and 10.7.

3.4.2 Building blocks

In MPC, an optimal control problem (OCP) is solved at each sampling time to calculate the input to the controlled system. An OCP is a special class of mathematical optimization problems that consists of:

- *A model*: a mathematical description of the controlled system (for example the battery in the BFH demonstrator), establishing a relation between its inputs, its internal states, and its outputs (see Section 10.7.1 for details).
- *A cost function*: the objective to be minimized, such as the CO₂ emissions of Section 3.2.2 (see Section 10.7.2 for details).
- *Constraints*: a description of the physical and any other desired constraints on the inputs, states, and outputs of the system (see Section 10.7.3 for details).

The solution of an OCP is the optimal trajectory of the inputs to the system for a predefined time-window that minimizes the given cost function. To compensate for uncertainties, such as for example deviations from the predicted non-shiftable load, the optimization problem is re-solved at each sampling time based on the most recent information.

Note that the solution of such optimization problems in real-time is not a trivial task from a computational point of view. Fast and reliable numerical algorithms and software to solve such problems is an active field of research as well as a lucrative domain for many companies that develop commercial software packages.

3.4.3 Use of MPC in field tests

MPC has been used in a large range of control applications in the field of power systems. For the purposes of this report, aim of aliunid is to:

- design an MPC controller to investigate how breathing in a single household can reduce CO₂ emissions and costs.
- find a combination of open-source, commercial and self-developed software that solves such problems reliably, while maintaining reasonable (and scalable) deployment costs.



3.5 Real-time data infrastructure

Author: Andreas Danuser, aliunid

3.5.1 Introduction

aliunid designs a series of products to allow for a new way of managing energy systems – the «breathing energy network».

The foundation of such a breathing energy system is a data infrastructure, which allows to observe, analyze, and control the behavior of the energy system in real-time. And this is exactly the goal and mission of so-called «Smart Grids», as illustrated in [19] (Smart Grid Roadmap Schweiz).

This section introduces the key aspects and requirements for the real-time data infrastructure behind a smart grid and outlines various implementation options, which finally leads to the solution chosen by aliunid, i.e., SIOT. The SIOT system, however, will be subject of the next section.

3.5.2 Key attributes and requirements

Given by the nature of a smart grid energy system (see Figure 16), the related data infrastructure has the following key attributes:

Tasks

The data infrastructure must do the following:

- It shall observe/measure, by means of suitable sensors and energy meters, the energy flows at different points of the grid. The goal of these measurements is to build a “digital twin” of the energy grid.
- It shall analyze such flows to understand what is going on and to determine an optimal operation of the grid, which includes functionalities to forecast and optimize energy flows.
- It shall influence/steer the energy flows in the smart grid by means of controlling production and/or consumption of energy as well as charging or discharging energy storages.

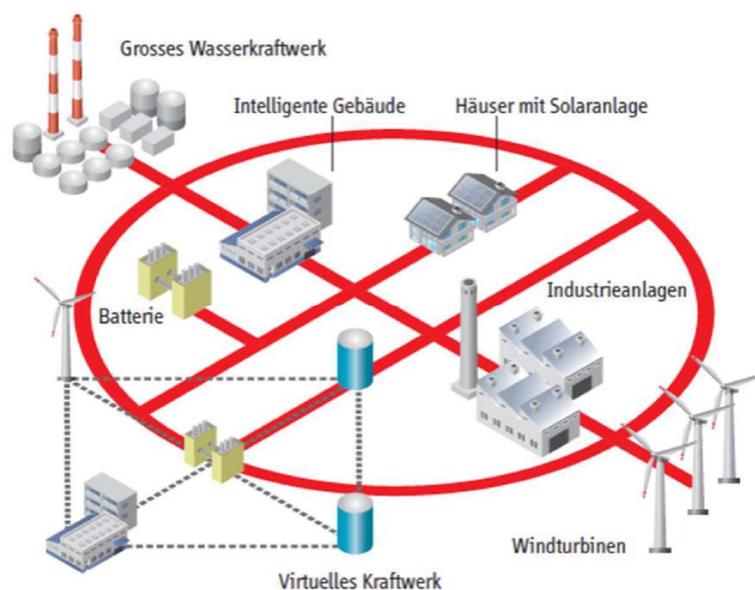


Figure 16: Smart grid.



Distribution - «ubiquitous» and «pervasive»

The distribution grid spans over large areas (e.g., whole cities). Therefore, the measurements and controls of the energy flows will be distributed as well – over large distances.

Building a meaningful digital twin of the energy grid will involve many measurement points. So far, the meters of an energy system were used primarily for billing. Smart grid enabled energy systems will include a lot more measurement points, in the distribution grid (NE7), at the edge to the user/customer as well as inside the customer premises. Correspondingly, the number of actuators controlling the energy flows is large as well.

This is referred to as «ubiquitous» and «pervasive» - a key attribute of IoT systems; and it is quite different from today's energy grids, which engage measurements primarily at the edge of the NE7.

Real-time

Since the data infrastructure shall control the smart grid based on real data measured in the grid, such data must be available for analysis soon after the measurement. The time between measurement and availability of the data is referred to as “latency”. In addition, the expressiveness of the data is much better, when having measurement data very often. This is referred to as the frequency of the measurements, or just “frequency”.

Table 5 shows the real-time specifications for smart grids, as they are required by the breathing energy grid of aliunid: Measurements shall be made every 1 to 10 seconds, and the latency shall be less than 1 second. Analyses based on these measurements shall be made every 10 to 30 seconds and actions typically will be initiated approximately every 1 minute, whereas the effect of the action will be detected after maximum 10 to 15 seconds.

Table 4: Real-Time Requirements related to smart grids.

Activity	frequency	latency
Measurement	1“– 10“	0.1“-1“
Analysis	10“– 30“	1“- 5“
Action	approx. 60“	10“- 15“

Security

Energy networks are mission-critical and failure, although happening sometimes, is basically no option. Therefore, the data infrastructure behind smart grids must work properly and as expected. These security demands can be grouped as follows:

Infrastructure availability:

The computing nodes as well as the links of the data infrastructure shall not fail; if they fail, a suitable backup solution must be available. Failure of computing node can occur as a defect of a hardware device or as a malfunction of the software, e.g., by infiltrating malware/malicious code.

Data security:

Measured data are the basis on which the data infrastructure decides how to influence/control the smart grid. Wrong or manipulated data will lead to a malfunction of the smart grid. Therefore, the data infrastructure must be able to answer the following questions:

- Do the data originate from a trusted source, or were they sent by a malicious person (“man in the middle”)?
- Have the data been altered on the way?
- Are the measured data plausible, or might there be a measurement error? The system must detect anomalies regarding the measured data.



- Can an undesired party read, understand and/or alter the data and thus attack the system?

Data protection

The following questions relate to protecting the privacy of the people connected to the grid as well as to keep critical core aspects of the grid confidential:

- Can somebody read and understand my data and thus conclude my private life (e.g., absence vs. presence)?
- Can somebody read and understand the data related to infrastructure and conclude the structure and status of a mission-critical system?

Life-cycle management

Smart grids are large systems, which are not built in one step, but evolve over a long period of time. Their lifetime is long, even eternal since they undergo permanent changes and improvements. This imposes stringent requirements to the underlying data infrastructure. The components of such infrastructure must be based on sound technical standards, which will be valid for a long time and which can be migrated to future technology versions. The data infrastructure will also have to work “eternally”, of course thanks to a clever technology migration strategy.

3.5.3 Data model and system architecture

The requirements introduced above cannot be addressed by several security and communication protocols. We rather identify two major elements, which allow to successfully build a real-time data infrastructure for smart grids:

- Suitable data models
- Functional decomposition

The challenges of large and changing software systems have been addressed by various methods, like object-oriented software construction and micro-service based reactive system designs. Large and complex systems need carefully designed data structures and a system architecture following the principles of object-oriented design and reactive (event-driven) micro-service systems.

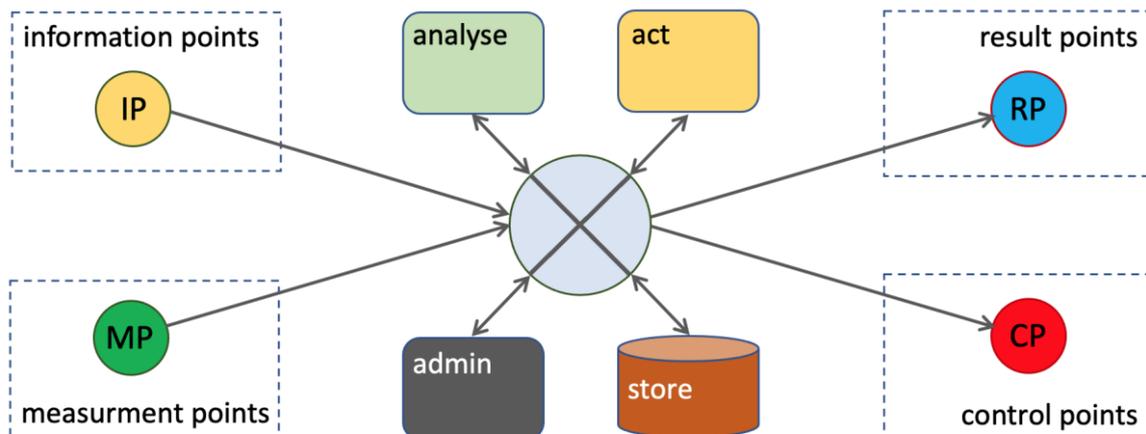


Figure 17: Reference model IoT functionality.



3.5.4 Reference model IoT functionality

Figure 17 introduces a reference model for the functionality of IoT systems and the real-time data infrastructure of a smart grid:

Data measured by sensors (at measurement points – MP) and information obtained from other sources (information points – IP) are made available to various functions: data may be stored, they are analyzed followed by actions, which lead to commands issued at control points (CP) or to results forwarded to other destinations (result points - RP). Finally, all data must be administered, which for instance includes directory services. This reference model shall serve to discuss and qualify the different approaches towards an IoT system suitable for building the data infrastructure of a smart grid.

3.5.5 Implementation options

Figure 18 illustrates two different approaches on how to set up a cloud computing system. Typically, an IoT system is built today as a cloud infrastructure using the Internet as communication network: measurement points (green) send measured data to the cloud, and control points (red) receive commands and data from the cloud. Likewise, additional information, e.g., weather forecasts, is collected (yellow) and results are forwarded (blue).

The central cloud hosts all functions introduced in the reference model, like administration of all data points, analyses, initiating actions and storing data.

This approach works well, and it is often easy to use. However, there are serious drawbacks. The connectivity between data points and the cloud may be interrupted, or it may be *slow* leading to long latencies. The *amount of data to be processed* at the central cloud may become large, and data protection is *difficult*.

These drawbacks have recently led to the introduction of a concept referred to as “edge computing”: Computer nodes are placed near to where the measurement and control points are located. These computers are tasked to execute time-critical analyses and control functions, i.e., they eliminate the latencies imposed by the potentially long links between datapoints and the central cloud. Moreover, they make the system more robust in case of link failures. Nevertheless, these systems still rely on a central cloud, which is a limiting factor for this system architecture.

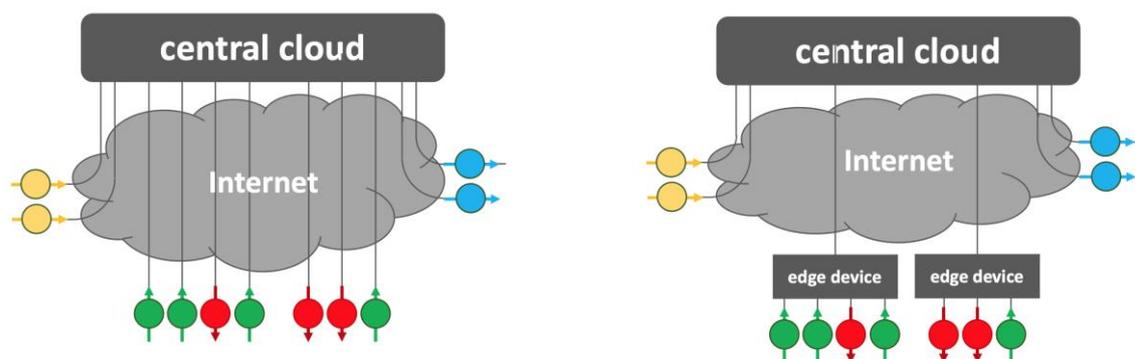


Figure 18: Central cloud (left) and cloud with edge computing (right).



3.6 SIOT: Swarm IOT

Author: Andreas Danuser, aliunid

The previous section introduced the cloud-based IoT system, which was recently extended by the concept of edge computing to overcome some of its disadvantages. In this section, we introduce SIOT, which totally revolutionizes the way how IoT systems may be implemented.

Foremost, the letter «S» in SIOT stands for “swarm”, i.e., swarm IoT.

The «S» also characterizes other key attributes, like «s» for “secure”, “simple”, “structured” – attributes as they have been required in the previous section.

Finally, the «S» also indicates the origin of this invention, which is Switzerland!

3.6.1 Swarm

Unlike the cloud-based approach, where a central cloud is responsible to provide data storage and functionality of the IoT system, SIOT consists of a set of many “small” clouds, referred to as «Private Clouds» and abbreviated by «PRIC» (see Figure 19 below).

A PRIC is an entity, which works autonomously, and which includes all functions as defined in Section 3.5 (see reference model for IoT functionality in Figure 17). Data measured at measurement points and other information from measurement points are entered into the PRIC, where they are analyzed and stored. The PRIC administrates all data points, sends commands to control points, and results to result points.

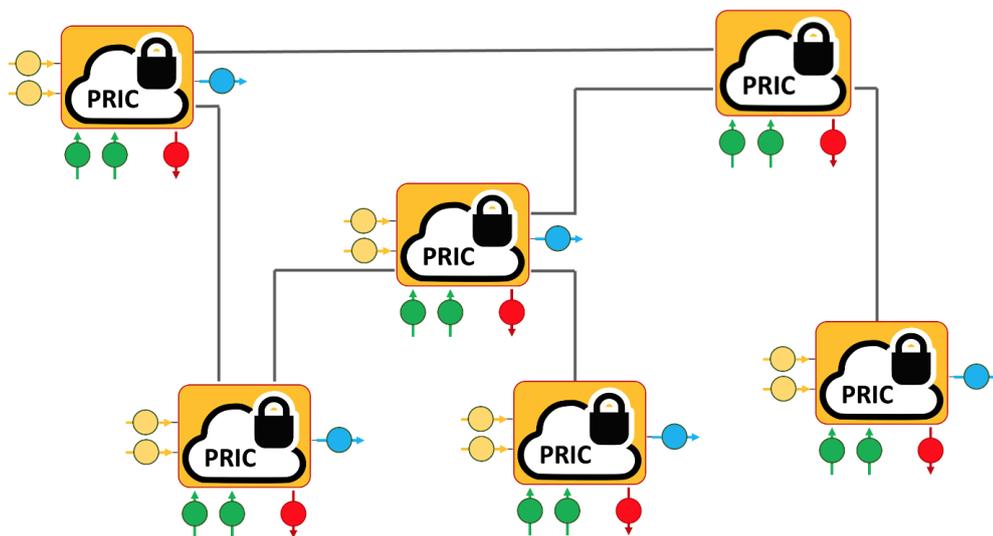


Figure 19: Swarm of private clouds.

These PRICs exchange data at their own discretion, i.e., each PRIC controls what data is sent to or received from which other cloud.

A set of such PRICs, where each PRIC communicates with other PRICs, is referred to as a swarm of PRICs. Therefore, this operational model is called “Swarm IoT” or «SIOT» and it is summarized by the slogan «... and there is no cloud».



3.6.2 PRIC and its environment

Figure 20 illustrates the “inner life” of a PRIC: The central point acts like a “town square” – a communication platform based on data messaging (e.g., using MQTT). Various function modules, referred to as «micro-services» (e.g., a particular analysis function) “pull” data from this town square, perform their functions (e.g., analyses, forecasts, action initiation and administration), and “push” the results to that same place. Micro-services may also act as interface services, which fetch data from measurement points and information points, or which send data to control points and result points. A key attribute of this town square is rigid rights and roles management: the system can exactly control, which micro-service can pull from or push to the town square.

Data is stored locally at the PRIC. A PRIC offers database resources for storing time-series of data as well as structure information, like directories and status tables of resources available at the PRIC.

The description above emphasizes, that the architecture of a private cloud is purposely engineered to achieve the key attributes of an IoT system, as required in the previous section:

- Functional decomposition into micro-services allows for simple implementation of the functionality.
- The entire functionality is provided by a (potentially large) set of micro-services, which are implemented following the functional programming model and which are “fired” upon availability of data (reactive functional programming).
- Data and functionality are encapsulated in a private space like it is the case with object-oriented software systems. Data, functionality, and communication is private and professionally managed, which is emphasized by the padlock of the PRIC symbol.

3.6.3 SIOT system

A swarm of PRIC as shown in Figure 19 looks appealing, but it only shows a logical structure. To make it really happen, we need a suitable system of computing nodes and communication links, which can execute the tasks of the PRICs and which securely transport the data between the sources and destinations. We need an underlying infrastructure, an «SIOT system». Figure 21 shows a related version of such a system.

Several PRICs may run on a server, referred to as a «group device» (lower right corner). Is it not this configuration kind of like a cloud system, what we would decline as a suitable solution? No, this is not the same, since every PRIC still has full autonomy and privacy. Each PRIC operates its communication system and has its data base fully private. However, failure of such a group device would harm several running PRICs; therefore, related redundancy measures must be applied to avoid server outage.

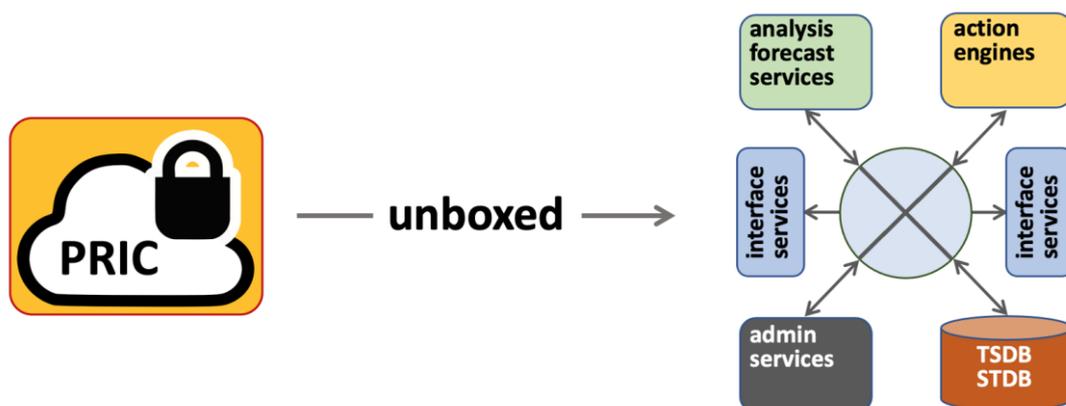


Figure 20: PRIC and its «inner life» based on messaging and micro-services.



A much better but also more expensive approach is to dedicate a (small) server specifically for one PRIC. Such a server is referred to as a «fog device», and it is typically located on-premises with the user. The name «fog device» emphasizes, that the cloud has come near to the user. It can also be understood as if the fog device is a cloud, which “dropped down” to the ground.

Each group and fog device has a physical connection to the Internet, or, more generally speaking, to an IP network. In most cases, fog devices are placed in a private IP network, which is connected via an access router to the Internet, which may even have a firewall to protect the network. Thus, the fog device has a private IP address and cannot easily be reached from the public Internet. This is a big advantage on the one side (the infrastructure is protected from malicious accesses), but it causes difficulties to network PRICs, on the other side.

Therefore, an additional system, a distributed VPN for the devices, referred to as DVPN, is used – described next.

3.6.4 DVPN

The DVPN is not one box, but an entire system of networking devices (routers, switches), servers and communication links.

When professionally engineered, a DVPN can be built with a high degree of redundancy and used to additionally protect the entire SIOT system from malicious access, such as distributed denial of service attacks (DDOS).

Each device running one PRIC (fog device) or many PRICs (group device) is connected to the DVPN, specifically to one of the routers in Figure 22, and establishes a VPN connection to one of the VPN servers (VPN-S).

The DVPN runs redundant services, which allow PRICs to be “found” and addressed (directory services), which again is under strict access control.

In SIOT, each PRIC has a unique identification, consisting of a «domain» and a «subdomain». The domain ID is worldwide unique, and the subdomain is unique to the domain. The scheme is like IP addresses, which consist of a network and a host address.

Each PRIC, uniquely identified by [domain/subdomain], can be reached via the DVPN under a specific IP address. Figure 23 shows, how two PRICs can find each other and connect via DVPN (see path 1).

The DVPN is also used for connecting user devices, such as a mobile phone or workstation, to the appropriate PRIC. This is shown by path 2, where a new element comes to the scene, i.e., the proxy device, which also part of the DVPN. This is explained below.

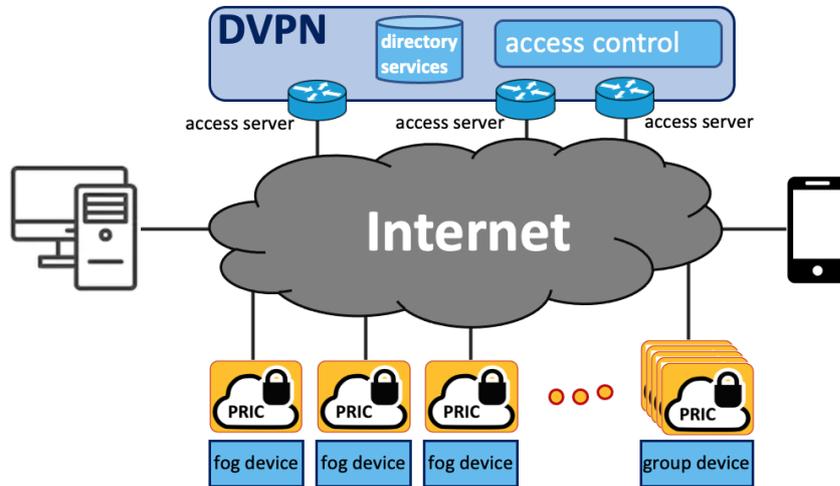


Figure 21: SIOT system.

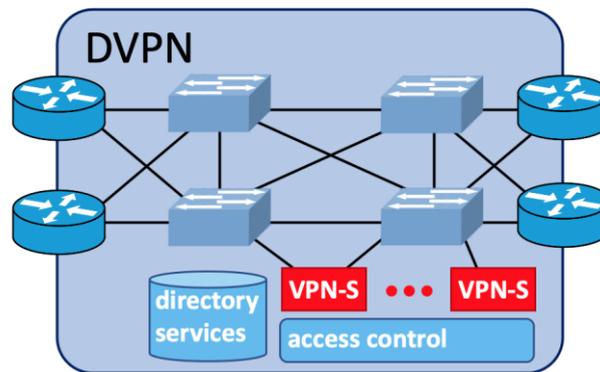


Figure 22: DVPN – inner life.

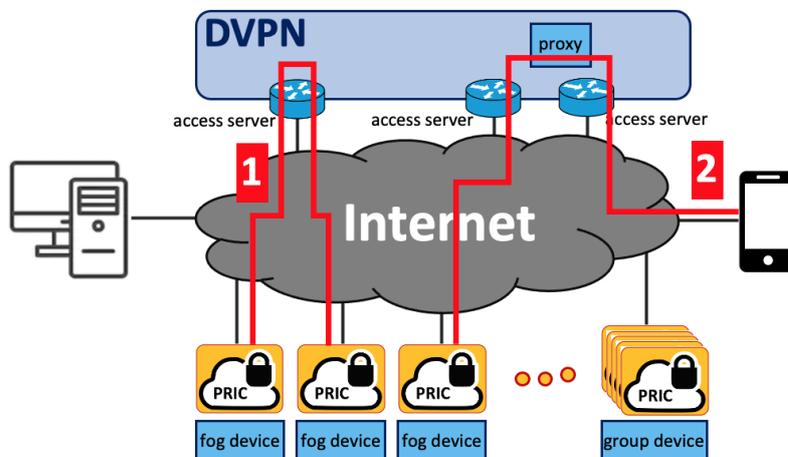


Figure 23: DVPN operation.



3.6.5 Security aspects – Example DVPN proxy

Security in IoT is a big concern, when it comes to control mission critical processes like smart grids. SIOT has been consequently built to address these security concerns, and some of them have already been illustrated above, such as:

- Security, privacy, and data protection enforced by the PRIC
- Service availability enforced by the swarm distribution of PRICs
- Redundancy regarding the inter-connection through the DVPN
- Data structuring and object orientation of the entire system architecture

We now want to illustrate on the example/use case of the proxy server, that SIOT really allows for highest level of security.

The DVPN hosts several servers referred to as proxy servers, and these servers are the intelligent entry points for external devices, e.g., a smartphone of the user.

When such a device connects, the access control mechanism checks whether the user of the smartphone has access, and if so, to which PRIC such access is granted. With the help of the directory service, the related PRIC is found and can be accessed by the user.

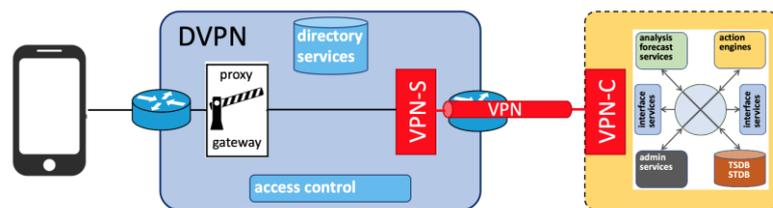


Figure 24: DVPN access.

However, the proxy server still does not allow for open access to such a PRIC; it rather restricts the access with the help of protocol specific gateways.

An even more stringent security can be introduced based on pattern analysis of the traffic between the PRIC and the external device. The gateway will learn, how such a pattern looks like.

With these measures, it is, for instance, also possible to protect the PRIC and the entire SIOT system from malware, since the infiltration of malware would either not use the one of the permitted protocols or would violate the known traffic patterns.

3.6.6 Conclusion

SIOT is a novel approach to build IoT systems. It works much different from traditional cloud-based models, even if such models might be “performance enhanced” by means of edge computing.

SIOT has been purposely built considering the high security, availability, redundancy, and real-time requirements of mission-critical systems like smart grids.



4 Results and discussion

4.1 Installation of real-time data infrastructure

Author: Till Richter, aliunid

4.1.1 Measuring concept

aliunid is currently recording real-time data from 527 devices. The devices which deliver real-time data consist of a mix between meter-behind-the-meter installations, smart meter readings, interfacing third party devices and measurements from PV inverters. The detailed overview is shown in Table 5.

For every utility that is contributing to the field test, a detailed plan has been created. An example is shown in Figure 25. Within the scope of this study, there were two different approaches:

- In two distribution grids, the focus was to engage as many prosumers as possible. This helps to test user interaction and understand various production and load patterns. The case studies described in Section 4.3 and 4.4 are based on this measuring concept.
- In the remaining distribution grids, the focus was on installing as many real-time data points as possible underneath a transformer station. This approach enables detailed load flow calculations in a subpart of the distribution grid in real-time. At the time of writing, these calculations are not completed, and results are subject to more in-depth analysis.

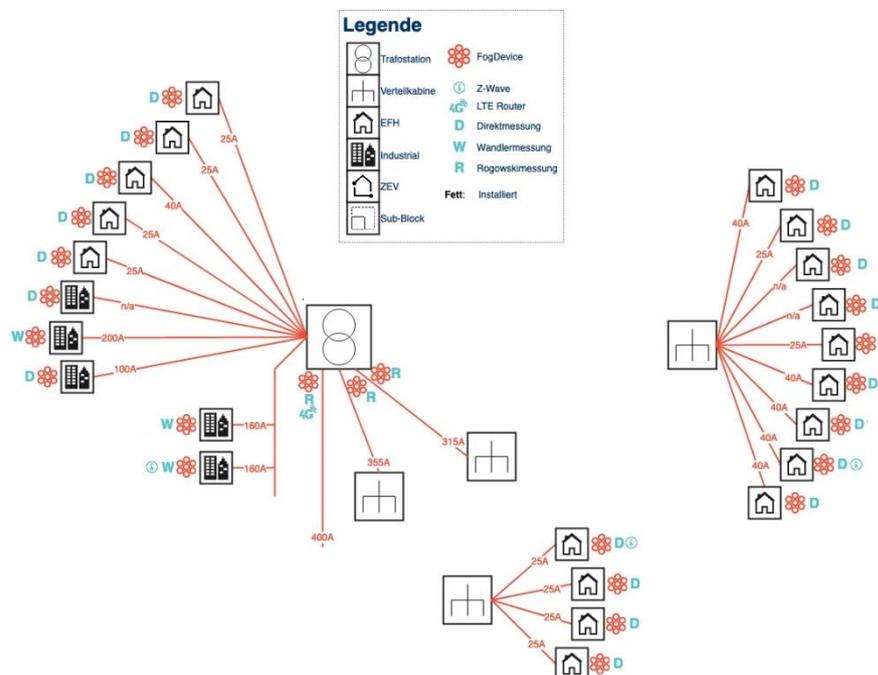


Figure 25: Example of a measurement plan for one utility participating in the aliunid field test.



Table 5: Overview of real-time data infrastructure of aliunid (status 29.10.20).

	Emxxx		Inverter		MeterSensor		BlinkSensor		3rd party		Total	
	inst.	active	inst.	active	inst.	active	inst.	active	inst.	active	inst.	active
<i>AET</i>	3	3									3	3
<i>EnBAG</i>	10	7									10	7
<i>Uster</i>	27	26									27	26
<i>esr</i>	10	7							196	196	206	203
<i>Elektra</i>	21	20									21	20
<i>ibi</i>	16	14									16	14
<i>ibm</i>	8	5									8	5
<i>ibw</i>	22	22	5	5	4	3					31	30
<i>lwa</i>	23	23			4	4	3	2			30	30
<i>localnet</i>	25	25							2	2	27	27
<i>SIG</i>	0	0									0	0
<i>seic</i>	13	13							122	122	135	135
<i>SWG</i>	0	0									0	0
<i>ylbe</i>	4	4									4	4
<i>aliunid</i>	4	4	2	2	3	4	4	3	9	9	22	23
Total	186	173	7	7	11	11	7	5	329	329	540	527

Goal per 31.12.2020	600
Difference	73

4.1.2 Propagation of load flows within the grid

In the figures below, the propagation of load flows in the distribution grid is visualized. Figure 26 shows the load profile of a single household. The load from the heating can be clearly seen in the night during winter. Figure 27 shows the heat map of the distribution cabinet one level higher. What becomes clearly visible here are the scheduled loads through ripple control, which result in vertical lines in the heat map. The load patterns measured at the transformer station one level up show how the statistical effects from the bigger network make the overall consumption patterns smoother and the effect of the scheduled loads more distinct (Figure 28). White space in the figure indicates missing data. The scheduled loads are an indication of flexibility, which can be used to optimize loads according to a chosen “tariff”. Section 10.2 elaborates more on this topic.

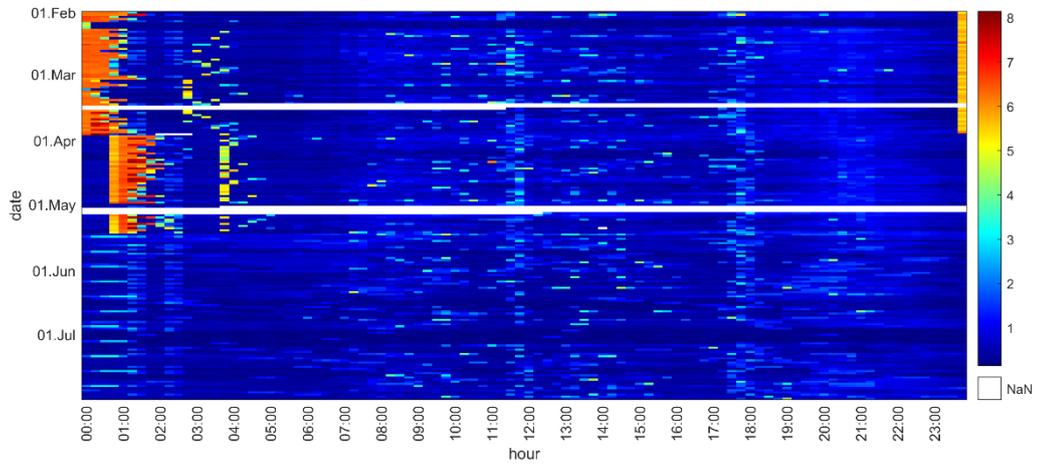


Figure 26: Recorded data single house.

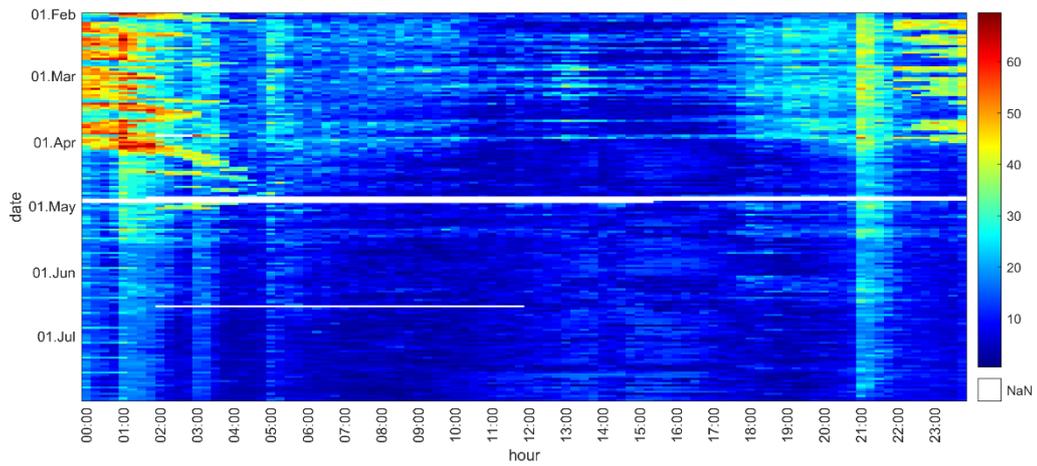


Figure 27: Recorded data substation.

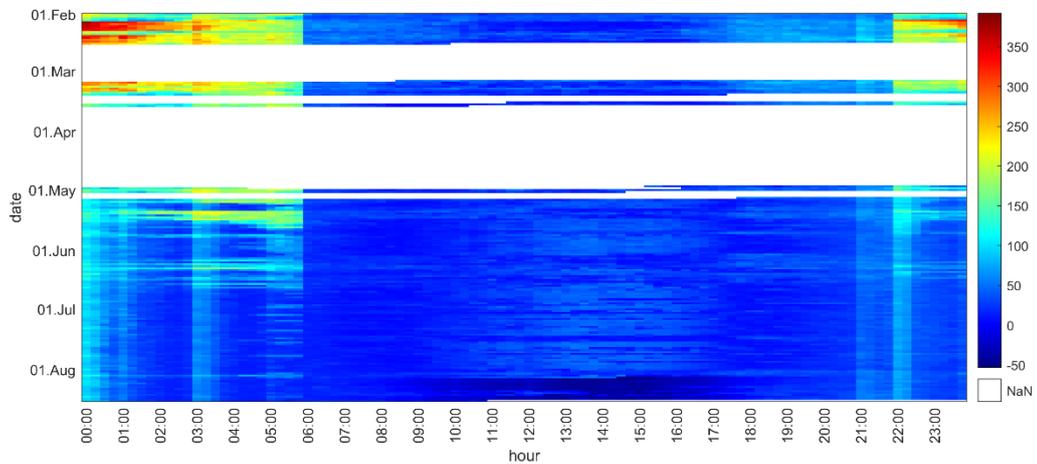


Figure 28: Recorded data transformer station.



4.1.3 Customer Interface and visualization of the four-value model

For every GB, aliunid's customer interface visualizes the four calculated values in the aliunid home app. As described in Section 3.3, inhaling and exhaling refer to the loads that can be shifted while the consumption is defined as all the loads that cannot be influenced.

Engaging end users on the topic of energy efficiency and their current consumption will help to make this topic relevant for a bigger audience. The aliunid home app will be the interface to the end customer. While the current field test app demonstrates the technical capabilities of providing a real-time data visualization to the customer using SIOT, the app that aliunid plans to bring to the market puts more emphasis on telling a story, better explaining the concept of a breathing energy supply, and engaging the customer.

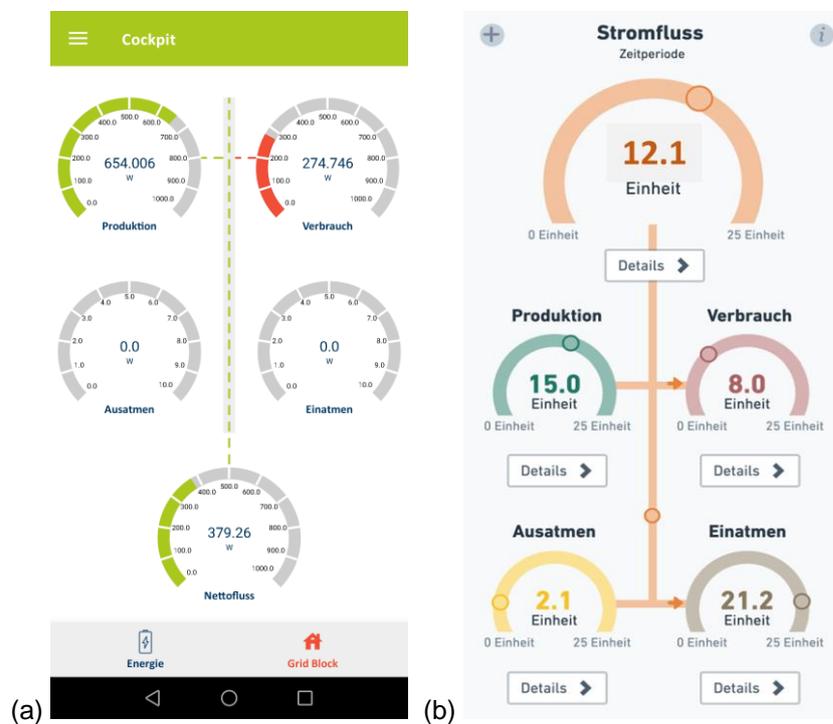


Figure 29: visualization of the four-value model (a) in the field test app and (b) in the planned release.



4.2 Breathing in one building using MPC to minimize CO₂

Authors: Dimitris Kouzoupis, aliunid

4.2.1 Experimental setup

In this section, the abstract description of an MPC controller in Section 3.4 becomes specific to the hardware setup described in Section 2.2. In addition to the high-level MPC controller, a low-level feedforward controller is implemented to compensate for the mismatch between actual consumption/production and forecast. The block diagram of the implementation is shown in Figure 30, where the orange arrows depict the flow of electricity and the blue ones the control loop.

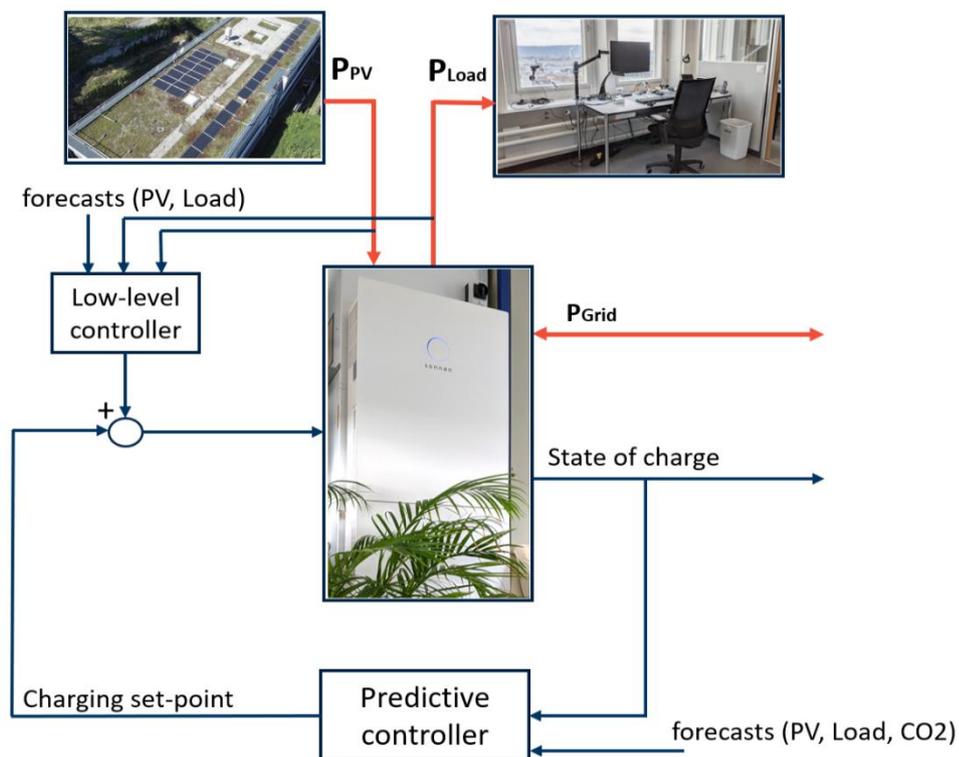


Figure 30: Block diagram of breathing experiment.

4.2.1.1. High-level controller

The predictive controller is called every 15 minutes to calculate the set-point for the power flowing in to or out of the battery. First, all necessary information is collected:

- Solar irradiation forecast for the next 24 hours in 15-minute intervals, converted to PV production forecast (see Section 10.7.1). The forecast is updated once an hour.
- Non-shiftable loads forecast for the next 24 hours in 15-minute intervals (average consumption over the past two weeks, during the same quarter of the day).
- CO₂ forecast for the next 24 hours in 15-minute intervals, calculated as explained in Section 3.2. The forecast is updated once a day between 18:00 and 19:00 and contains values until the next day at 23:00. When the available forecast becomes shorter than 24 hours, the remaining part is padded with the maximum value of the current trajectory.
- State of charge measurement from the battery interface.



Based on this information, an optimization problem is solved that calculates the optimal charging profile of the battery for the next 24 hours. The objective of the optimization problem is to minimize CO₂ emissions. Therefore, when the controller foresees that the available energy from the battery and the PV will not be enough to cover the future consumption, it purchases electricity from the grid in time periods with low CO₂ footprint.

To compensate for uncertainties in the simulation model of the battery and the weather forecast, only the first value of the charging profile is sent to the battery and the procedure is repeated at the next sampling time (see Section 3.4 where the concept of MPC is introduced).

However, due to the high degree of uncertainty in forecasting the solar irradiation and especially the electricity consumption within a building, a control scheme as such will hardly be able to interact with the grid as planned. For example, if one turns on a device that consumes a lot of electricity during the night, the high-level controller has no way to detect this event and all the required energy will flow from the grid to the device, even if the battery is fully charged. This fact highlights the need for a low-level controller, that considers the current consumption and PV production, and corrects the charging set-point accordingly.

4.2.1.2. Low-level controller

The principle of the low-level feedforward controller is quite simple, yet essential for the overall performance. The error between the current measurements of consumption and PV production and their forecasted values affects the charging set-point, such that energy flow prioritizes the battery instead of the grid. If more energy is produced than anticipated, the battery is charged with a higher power and vice versa. If more energy is consumed than anticipated, the battery is discharging more than planned instead of using electricity from the grid. A saturation logic ensures that the calculated value stays between the limits that the battery expects. Moreover, the controller blocks the energy flow to the battery when the state of charge reaches its maximum value and the energy flow from the battery when it reaches its minimum value. The low-level controller corrects the charging value of the battery once every second.

4.2.2 Results

The experiment described in the previous section was running on the demonstrator in Biel for a week. Aim of this section is to analyze the results and quantify how the designed controller improves the CO₂ impact of the demonstrator.

Section 4.2.2.1 compares the predicted behavior of the controller (i.e., the computed optimal trajectories) for a given point in time against a simple control strategy. The term ‘open-loop’ in the section title means that the effect of uncertainties in the forecasts and models is not taken into account yet, i.e., the results are presented from the point of view of the predictive controller. The more realistic, ‘closed-loop’ behavior is the topic of Section 4.2.2.2.

4.2.2.1. Open-loop benchmark

In this section, we compare the predicted behavior of the MPC controller to a simple control strategy that is further referred to as ‘baseline controller’. The baseline controller stores the surplus of energy ($P_{PV} - P_{Load}$) to the battery until it is full. If the consumption of the demonstrator is higher than the PV production, energy is drawn from the battery until the minimum specified state of charge is reached (same limit as the one set in the MPC controller). Analyzing the historical data of the battery showed that this simple control strategy is actually very close to the behavior of the proprietary controller of the battery (when in automatic operating mode).

To illustrate the comparison, two 24-hour prediction windows are analyzed. One for a sunny and one for a cloudy day. The calculated trajectories are shown in Figure 31 and Figure 32 respectively.

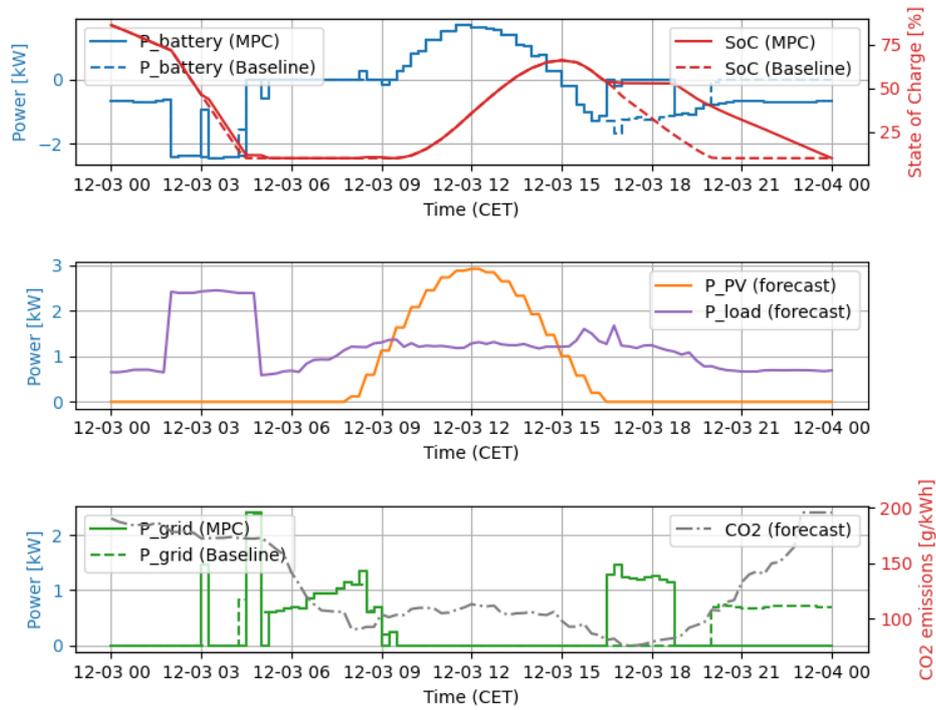


Figure 31: MPC versus baseline controller, sunny day.

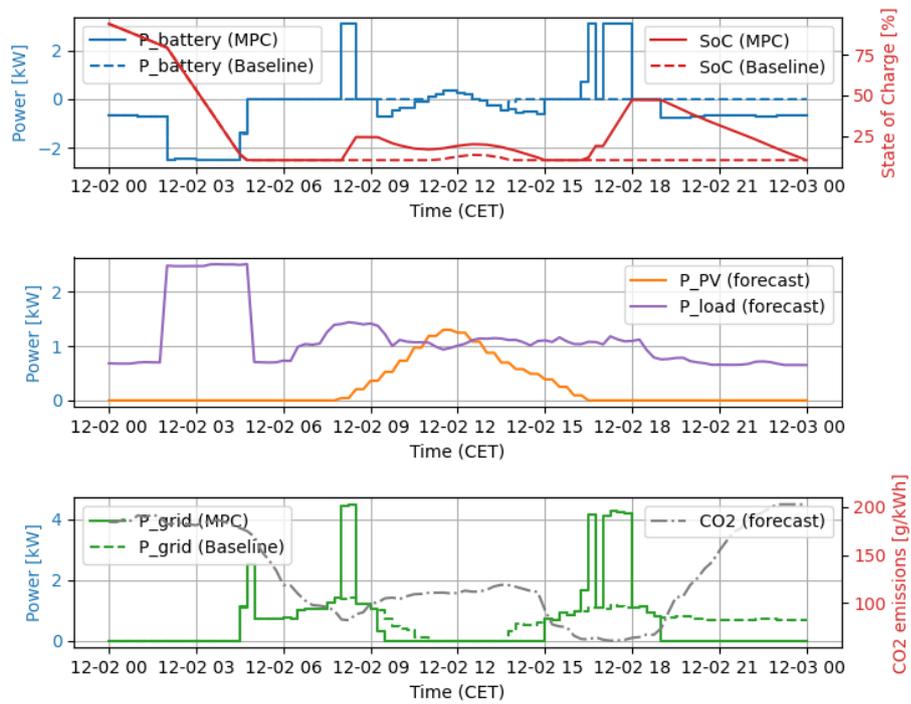


Figure 32: MPC versus baseline controller, cloudy day.



During the sunny day (based always on the available weather forecast), both controllers require about 8 kWh of energy from the grid to cover the demand. However, thanks to the available CO₂ forecast the MPC controller imports this energy when there is a notable decrease in the CO₂ while the baseline controller happens to charge the battery when the emissions are high. Note also that the controller cannot exclusively purchase energy during periods with the lowest CO₂ since it must also cover the forecasted consumption.

Another important point to understand the behavior of the MPC controller is that the underlying optimization problem has always a fixed 24-hour horizon. Therefore, it will always be optimal to empty the battery at the end of this horizon in order to minimize the CO₂ emissions. This is apparent in both Figure 31 and Figure 32. However, due to the receding horizon strategy described in Section 3.4, this end of the horizon is always shifted forward and will never be reached. Common remedies to reduce this effect is to penalize the difference of the last value of the state of charge from a desired set-point, but as illustrated by the results of Section 4.2.2.2, this was not necessary for the conducted experiments.

During the cloudy day, when the PV production is considerably lower, the points in time when energy is imported from the grid play an even more crucial role. As it can be observed in Figure 32, the baseline controller does not charge the battery and as a result the imported energy has a significantly worse CO₂ impact.

A comparison of the two controllers for the two cases considered is summarized in Table 6. The values refer to the 24-hour forecast at the point in time when the optimization problem was solved.

Table 6: CO₂ emissions and energy consumption for two optimal control problem instances versus baseline controller.

	MPC		Baseline		Difference	
	<i>Sunny</i>	<i>Cloudy</i>	<i>Sunny</i>	<i>Cloudy</i>	<i>Sunny</i>	<i>Cloudy</i>
Consumption [kWh]	7.63	14.84	7.63	14.16	+ 0%	+ 5%
CO₂ emissions [g]	865.52	1263.39	1066.21	1586.90	- 19%	- 20%
Average emissions [g/kWh]	113.48	85.16	139.78	112.09	- 19%	- 24%

It can be observed that the optimal controller's plan yields up to a 20% reduction in total CO₂ emissions with, however, some increase in energy consumption. This is due to the fact that the objective of the optimization problem is solely based on the minimization of total CO₂ emissions. To alleviate this effect, one can introduce a trade-off between these two costs (total consumption versus total emissions) by an appropriate reformulation of the objective. Table 7 shows how the optimum shifts by changing the weight of the two objectives. Scaling factor 100% means only emissions are minimized while scaling factor 0% means only energy consumption from the grid is minimized. For simplicity, only the cloudy day is considered.

Table 7: Pareto front values for simultaneous minimization of CO₂ emissions and energy consumption from grid.

Scaling factor [%]	Consumption [kWh]	CO₂ emissions [g]
100	14.84	1263.39
40	14.81	1263.82
10	14.68	1273.80
1	14.14	1551.59
0	14.14	2131.06

As one can notice, taking both objectives into consideration can help to reduce the consumption, but this reduction comes at the cost of higher CO₂ emissions. Note that with a minor consideration of the CO₂ emissions (only 1% weight) there is still a 2.2% reduction in total emissions with respect to the



baseline. It is clear than when one optimizes only for total consumption, there are infinitely many solutions that yield the same result and the negative impact on the CO₂ emissions is arbitrary (here 34% worse than the baseline). For the remain of this section, the scaling factor is set to 100%, i.e., total CO₂ emissions are minimized.

4.2.2.2. Closed-loop benchmark

The predictive controller was controlling the battery in the BFH demonstrator for a period of about a week (28th of November to 7th of December). The behavior of the battery together with the production, consumption and interaction with the grid are shown in Figure 33.

From the illustrated trajectories it is evident that the controller mainly purchases energy from the grid at favorable points in time with respect to forecasted CO₂ emissions. More precisely, the controller imports a total of 144 kWh that correspond to 17.7 kg of CO₂ (based on the actual values, not the forecast) and an average of 123.42 g/kWh. The average emissions of the energy in the grid are about 149.6 g/kWh during the same period. Compared to a simulation using the baseline controller,⁹ with a total of 21.56 kg of CO₂, the MPC controller results in 17.6% less emissions.

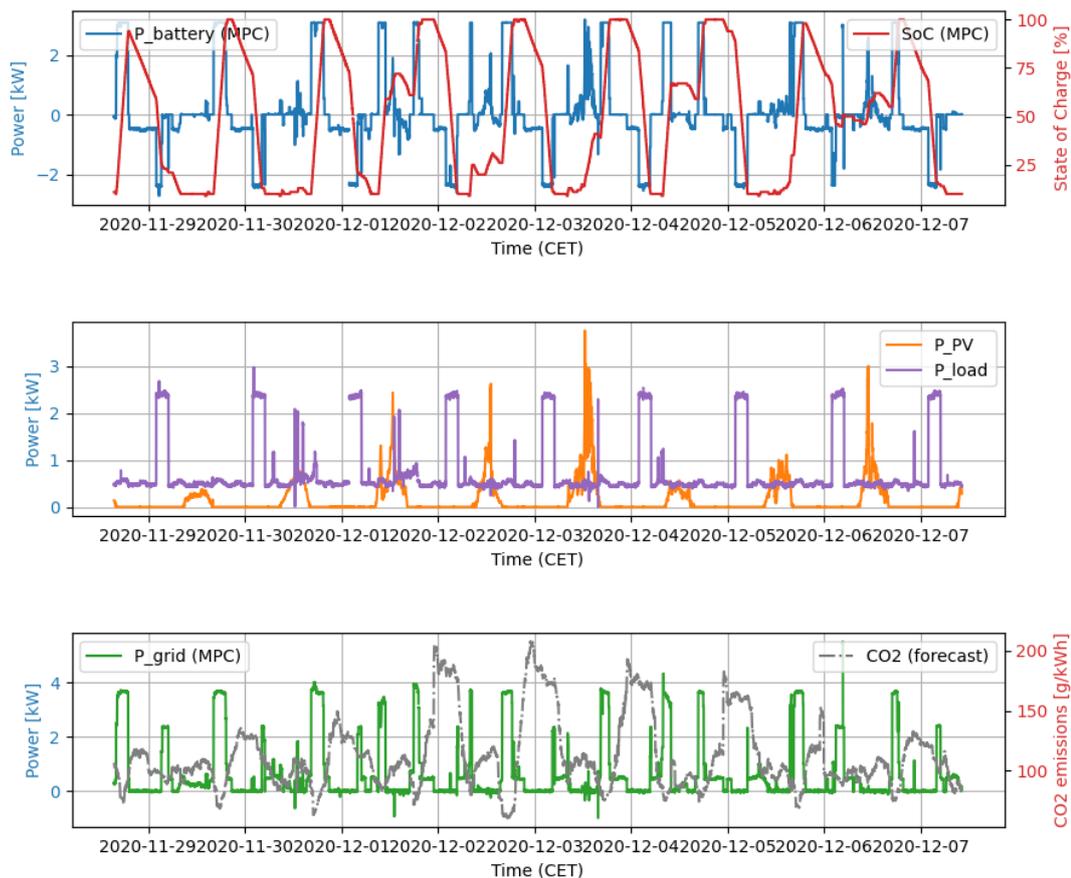


Figure 33: Measurements of breathing experiment in BFH demonstrator for a period of one week.

⁹ The model used to simulate the battery that is coupled to the baseline controller is identified using a least-squares fit on the data of the experiment. Consumption and PV data are readily available.



To get a better picture of the CO₂ savings, the behavior of the MPC scheme is compared to the baseline controller described in Section 4.2.2.1 for a period of 60 hours (a sunny day followed by a cloudy day).

The behavior of the two approaches is significantly different as shown in Figure 34. Due to lack of solar energy during the second day, the baseline controller keeps the battery empty, and the needs of the lab are covered directly by the grid. On the contrary, the MPC controller discharges the battery when CO₂ emissions are expected to be high and vice versa.

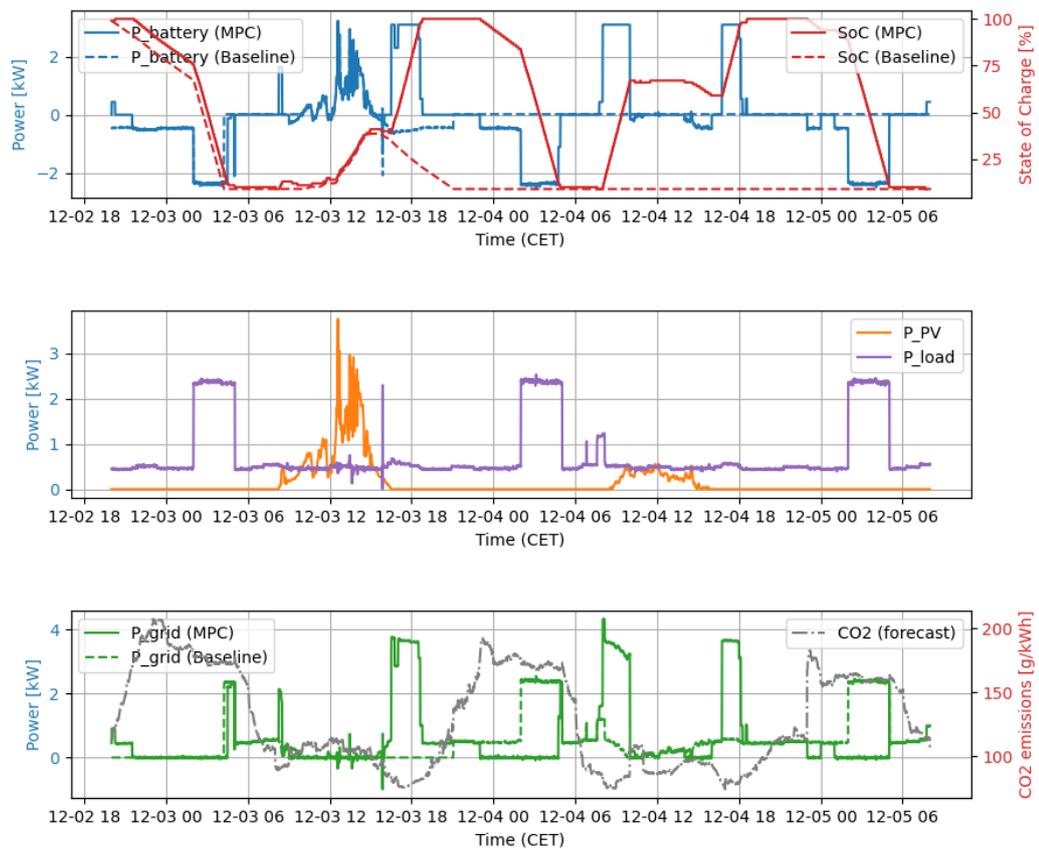


Figure 34: MPC versus baseline controller for a 60-hour window.



The consumption of the two schemes is summarized in the table below.

Table 8: CO₂ emissions and energy consumption in closed-loop.

	<i>MPC</i>	<i>Baseline</i>	<i>Difference</i>
<i>Consumption [kWh]</i>	33.90	31.22	+ 8.6%
<i>CO₂ emissions (forecast) [g]</i>	3313.17	4540.59	- 27.0%
<i>Av. emissions (forecast) [g/kWh]</i>	97.74	145.45	- 32.8%
<i>CO₂ emissions (actual) [g]</i>	4489.72	5045.55	- 11.0%
<i>Av. emissions (actual) [g/kWh]</i>	132.45	161.63	- 18.1%

The 8.6% increase in consumption is associated on the one hand to inaccuracies of PV and load forecasts, and on the other on energy losses due to the additional battery cycles. This behavior is expected since the objective of the controller is to minimize CO₂ emissions and not energy consumption, as discussed in Section 4.2.2.1. Note also that the reduction in actual emissions is considerably lower than forecasted - due to deviation of the CO₂ forecast from the actual emissions - but remains significant.



4.3 Breathing in multiple houses using DHW to minimize CO₂

Authors: Matthias Meier, aliunid

4.3.1 Setup

At present, domestic hot water (DHW) heating operation is commonly constrained to specific hours during the night, and the corresponding period is specified by the local utility. The infrastructure of aliunid however allows to exploit the demand side flexibility of domestic hot water tanks (DHWT) much further. To demonstrate this, DHWT consumption is optimized in two households (“house (a)” and “house (b)”) with the aim to minimize the impact on CO₂ emissions.

The general procedure is the following (see also Figure 35):

- CO₂ intensity is forecasted according to Section 3.2 every day shortly after 18:00.
- Optimum DHW heating window is determined.
- Commands are sent to the relay of the DHWT at the corresponding times to switch it on/off.

To ensure that the interval between two hot water heating cycles cannot get too large, i.e., the comfort of the residents is not compromised, the load is only allowed to be shifted within a predefined range of ten hours. The interval is chosen in a way to cover the range where the daily minimum in CO₂ intensity is typically occurring. Specifically, given the CO₂ metric used for this report, the range is chosen from 13:00 (“hard constraint 1”) to 23:00 (“hard constraint 2”). To further ensure that the DHWT is heated up completely, the “on” window of the relay lasts two full hours (green period in Figure 35). Since most DHWT’s usually only require between 30-90 minutes to reach the target temperature and turn off afterwards (regardless of the relay state), the window to seek for the minimum CO₂ is constrained to only one hour (orange dashed period in Figure 35).

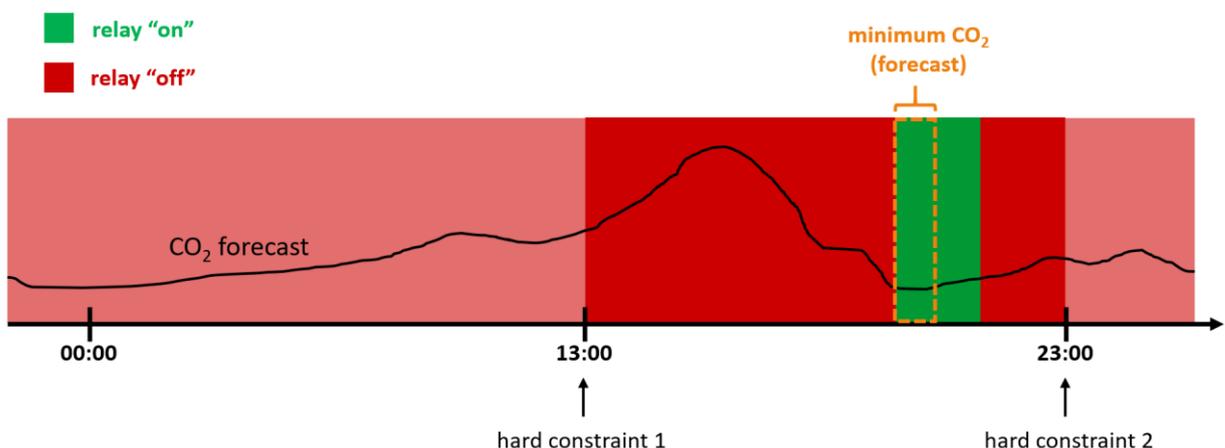


Figure 35: Illustration of the DHWT switching algorithm.

4.3.2 Simulation

To get an estimate of the potential in load shifting according to CO₂ with the above-described setup, the load shifting is first simulated. The simulation is based on historical CO₂ data (calculated according to Equation (3.1)) and historical CO₂ forecast according to Section 3.2.3. Furthermore, it is based on smart meter data of the two households involved in the experiment, which is disaggregated in order to extract the DHWT consumption and of which a typical daily consumption profile is extracted. The extracted typical daily DHW consumption is then shifted according to the CO₂ forecast and the resulting CO₂ emissions are compared with the standard consumption profile.



Power Consumption Disaggregation and Standard Profile

The smart meter data available provides aggregated¹⁰ active power and aggregated reactive power for all three phases. Depending on the type of DHWT (resistive or combined with heat pump), the DHWT power consumption may however show a distinct pattern in the reactive power phases or the active power phases. To disaggregate the DHWT power consumption from the aggregated power load, a median filter is applied on either the reactive or the active power of the three phases, followed by a k-means clustering algorithm. This procedure reliably detects the times when the DHWT is running (Figure 36).

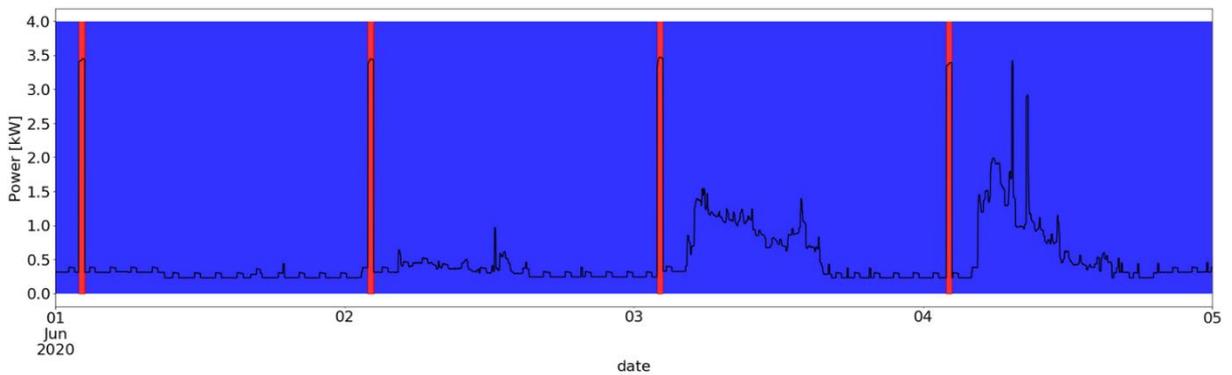


Figure 36: DHWT running times (red) of “house (a)” as extracted through k-means clustering of reactive power.

In the next step, the background is subtracted to retrieve the power consumption of the DHWT alone. Thereto, with a rolling window of the size of two hours, the minimum power value is calculated, and the resulting array is subtracted from the total active power. The standard DHWT power consumption value (dashed black line in Figure 37) is then defined to be the median of the background subtracted power consumption (orange line in Figure 37) during times when the DHWT was running (red parts in Figure 36).

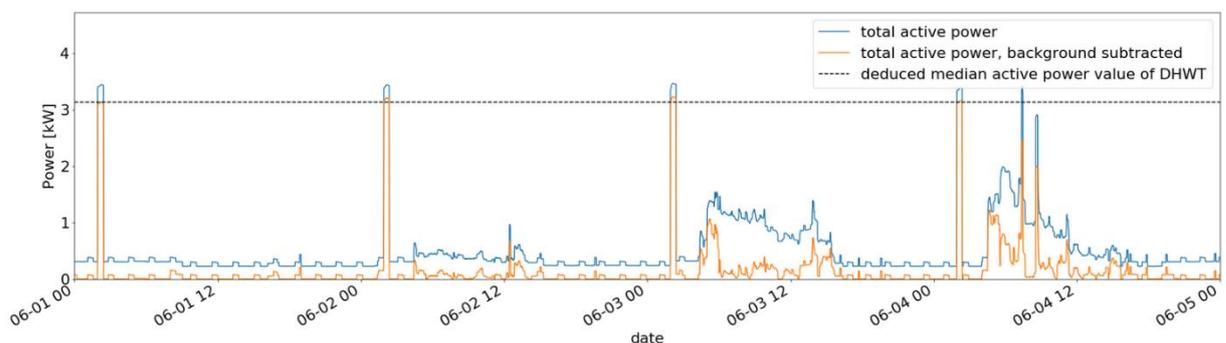


Figure 37: Total active power, background subtracted total active power and derived median active power value of DWHT of “house (a)”.

Having the DHWT’s power consumption disaggregated, its average daily run-time is calculated and its typical schedule is derived by inspecting the data. Figure 38 shows the actual, aggregated power data of “house (a)” and Figure 39 shows the resulting standard DWHT power consumption profile for the same period.

¹⁰ Aggregated in the sense that it represents the total power consumption of that household, i.e., the consumption of the DHWT is not sub-metered separately.

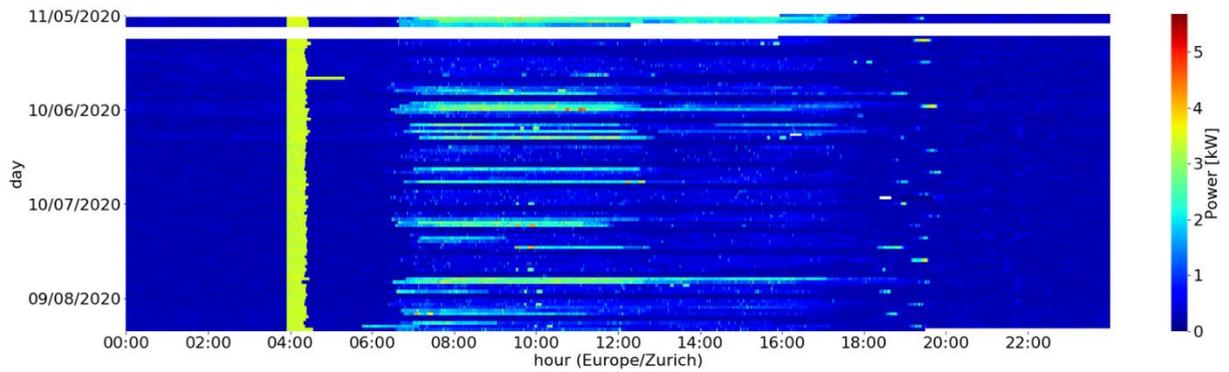


Figure 38: Aggregated power consumption of “house (a)”. White parts correspond to missing data.

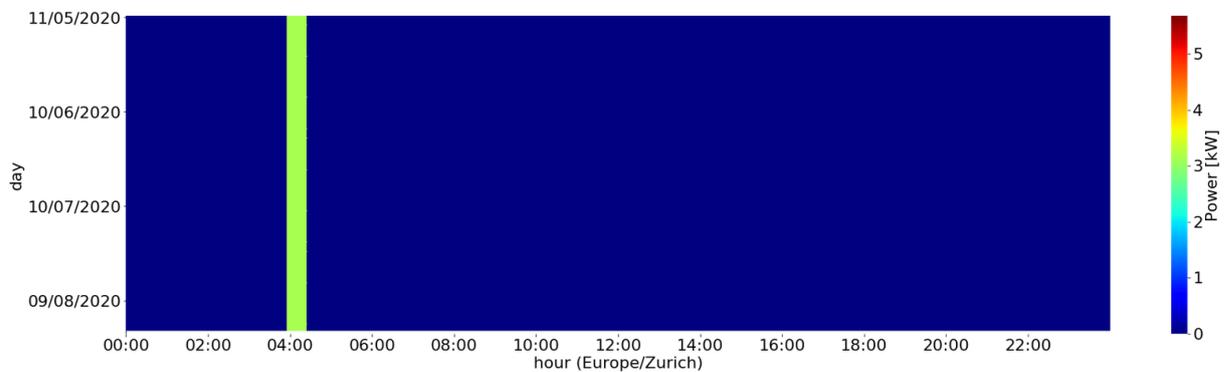


Figure 39: Derived standard DHWT power consumption profile for the same period.

Simulation Results

To evaluate the load shifting potential for a complete year, the standard profiles of the two houses are extended to the period from July 21st 2019 to July 20th 2020¹¹, and the load is allowed to be shifted according to the setup described in 4.3.1 either according to the CO₂ forecast or the actual CO₂ (Equation (3.1)). The results are shown in Table 9 and Figure 40 shows the resulting shifted DHWT power consumptions for “house (a)”.

¹¹ The period is chosen so that it lies almost exclusively in the validation set period of the CO₂ forecast model.

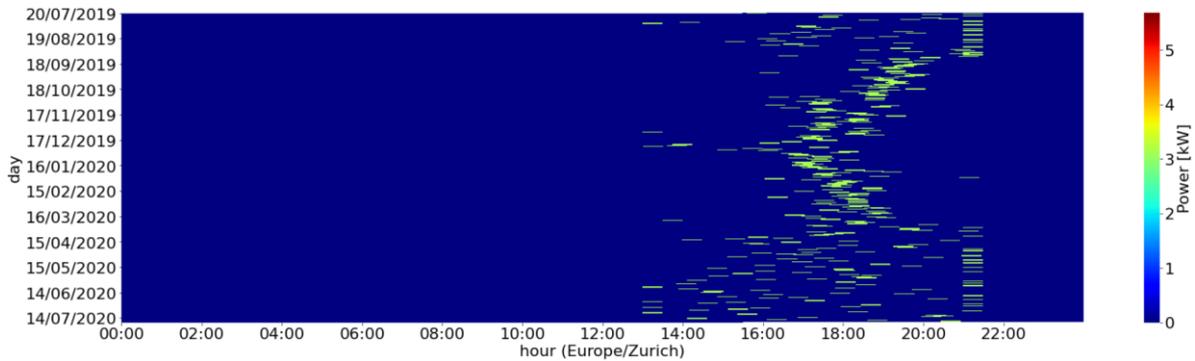


Figure 40: Synthetic DHWT consumption of “house (a)” shifted according to the experimental setup described in Section 4.3.1.

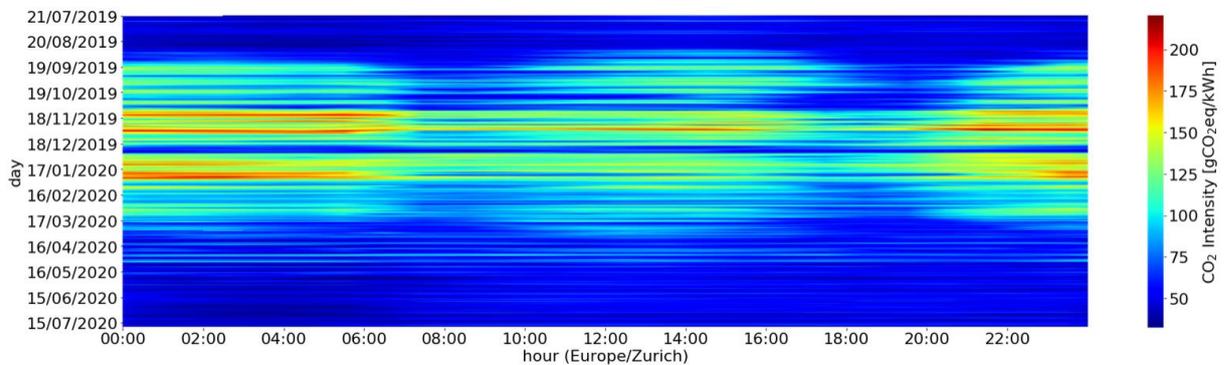


Figure 41: CO₂ forecast in the period of the simulation.

Shifting the DHWT power consumption with the given setup and the used metric for CO₂ may lead to a reduction of up to 25% of CO₂ emissions in the optimum case (“perfect CO₂ forecast”). With the CO₂ forecast used by aliunid, the savings may amount to roughly 20%.

Table 9: Results of the DHWT load shifting simulation.

	House (a)	House (b)
Total electricity consumption of DHWT July 2019-July 2020	576 kWh	2190 kWh
CO₂ emissions for DHW with <i>standard profile</i>	47.9 kg	184.7 kg
CO₂ emissions for DHW if <i>optimized according to CO₂ forecast</i>	38.9 kg (-19%)	149.3 kg (-19%)
CO₂ emissions for DHW if <i>optimized according to “real” CO₂ (i.e., “perfect CO₂ forecast”)</i>	35.5 kg (-26%)	135.6 kg (-27%)



4.3.3 Experiment

The actual experiment is performed during a period of 12 days. Figure 42 and Figure 43 show the resulting consumption profiles of the two households for a subperiod of the 12 days. Shown in grey are the periods during which DHWT tank operation is allowed and within which CO₂ is optimized according to the forecast - as described in Section 4.3.1. Shown in blue is the total power consumption of the corresponding household.

The disaggregation of the DHWT power consumption is performed as explained in the previous section, leading to the DHWT consumption profiles in red. To compare with the case of no load shifting, the corresponding DHWT consumption of any grey period is shifted to its default starting time to create the green default DHWT power profile. Furthermore, the actual CO₂ emissions and the forecasted CO₂ emissions are plotted in black and orange respectively.

Inspection of the two figures shows that shifting the loads according to the CO₂ forecast was successful. The DHWT was reliably running during times of minimum forecasted CO₂ within the predefined periods and without compromising the comfort of the residents.

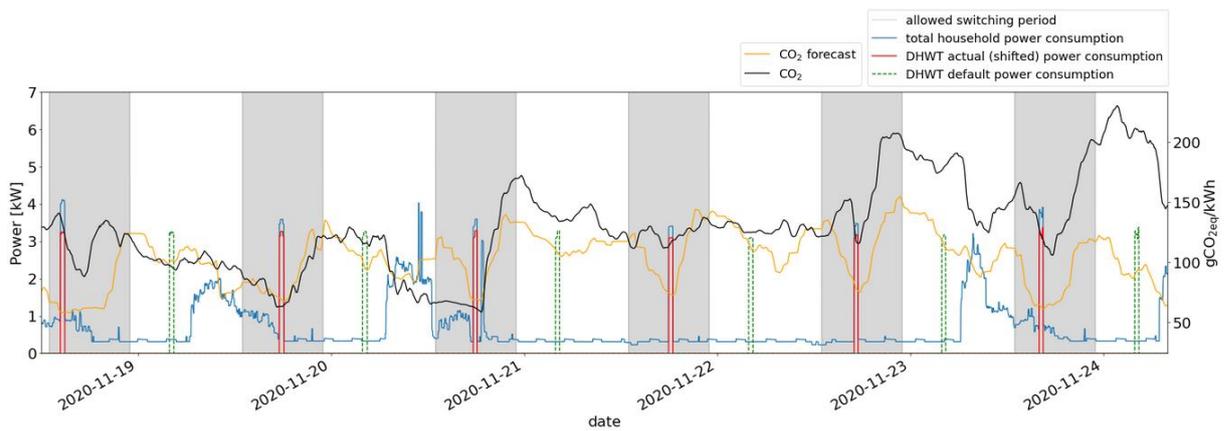


Figure 42: DHWT shifting experiment in "house (a)". The grey regions represent the periods within which boiler operation was allowed.

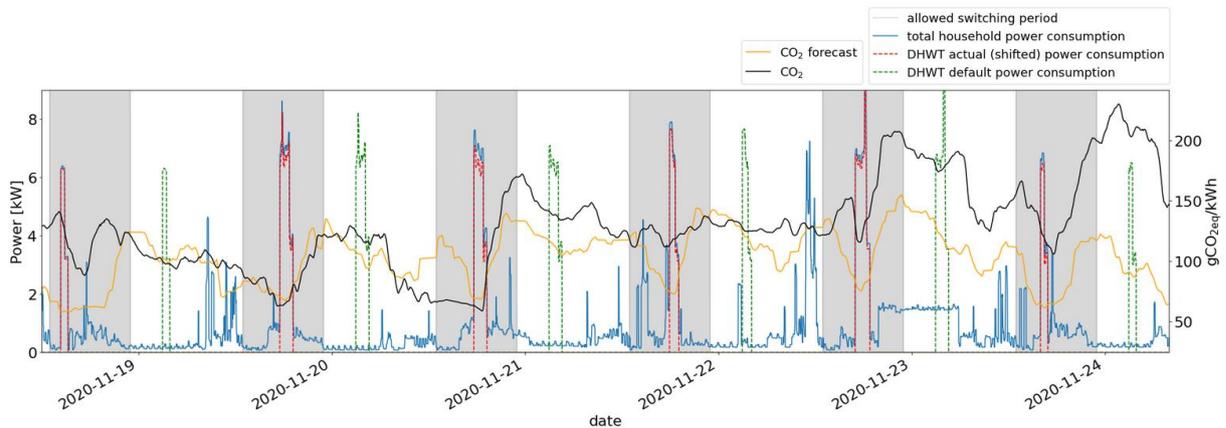


Figure 43: DHWT shifting experiment in "house (b)". The grey regions represent the periods within which boiler operation was allowed.



Furthermore, in most of the cases the DHWT was indeed shifted to times where the CO₂ was lower than during the default time. The overall results of the 12-day period are summarized in Table 10. It must be emphasized here that, as explained in more detail in Section 3.2.5, the reduction of CO₂ emissions shown in the table is only a number arising from the applied CO₂ metric. It does not directly reflect the real impact on carbon emissions.

The reason that the reduction is greater in the experiment than in the simulations is most probably due to the fact that the fluctuations in grid carbon intensity, and therefore the potential to reduce carbon emissions by load shifting, are much larger during winter times than during summer times (Figure 10).

Table 10: Results of the DHWT load shifting experiment.

	House (a)	House (b)
Total electricity consumption during experiment period	202.9 kWh	253.5 kWh
Electricity consumption for warm water	20.8 kWh (10%)	89.9 kWh (35%)
Non-shiftable electricity consumption	182.1 kWh (90%)	163.6 kWh (65%)
CO₂ emissions for DHW if <i>not shifted</i> (<i>default behavior</i>)	2.74 kg	12.11 kg
Actual CO₂ emissions for DHW (<i>forecast optimized</i>)	1.98 kg (-28%)	8.90 kg (-27%)



4.4 Breathing in multiple houses using HP to minimize CO₂

Author: Till Richter, aliunid

4.4.1 Setup

While Section 4.3 showed how to use flexibility from electric DHWT, this section shows how to use the flexibility of heat pumps (HP). The most common setup in Switzerland is that the heat pump is used both for SH and DHW. Since most heat pumps do not have an interface (e.g., Smart Grid Ready) that allows a more precise control, this use case demonstrates how flexibility can be used by disconnecting the device from the power supply for a limited amount of time.

The boundaries for the purposes of this experiment are shown in Figure 44. The utility allowed aliunid to turn off the HPs during high CO₂ once per day for two hours. However, for safety reasons, it was required to never block a HP between 4am and 6am to guarantee that there is enough hot water in the morning. In line with the DHW use case, the algorithm forecasts the CO₂ emissions at 6pm as soon as the forecast from third-party data sources is available. Based on this forecast the two-hour window with the highest forecasted emissions and which is outside of the “required on” window is identified. The controller sends then the “on” and “off” commands to the relay at the specified time.

The experiment was running for 7 days (between Oct 19 and Oct 26). During this time, heat pumps in 10 different houses were controlled and their power consumption (“inhaling”) was submetered.

4.4.2 Breathing to optimize CO₂

The aggregated time series of the measurements are shown in Figure 45. It can be clearly seen that the control algorithm worked as intended and the power of the HPs is reduced to zero during the time of the highest CO₂ emissions. As discussed in the previous section, the forecast is not always correct and therefore HPs are not guaranteed to switch off during the time with the highest emissions (e.g., on October 25). Table 11 shows the consumption for every individual house.

The results show that the principle works and that load of HPs can be shifted dynamically to minimize a previously defined tariff. However the quantification of CO₂ savings depends on the definition of a baseline, which is not a trivial task because of the nature of the problem (one cannot anticipate the possible behavior of the heat pump during the “cut-off” window reliably). The option to build data-driven models to derive a baseline case was evaluated, but to get meaningful results, more historical data are required. The current validation, however, did show an important quantification of flexibility, which is discussed in the next section.

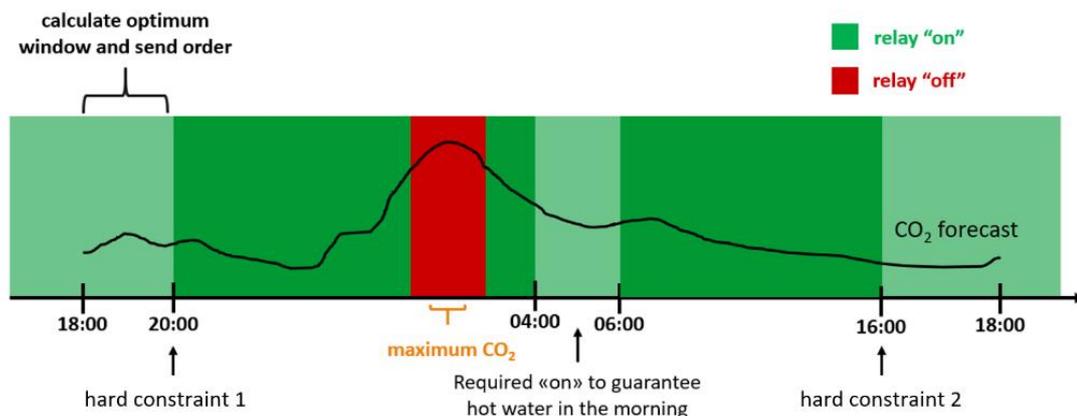
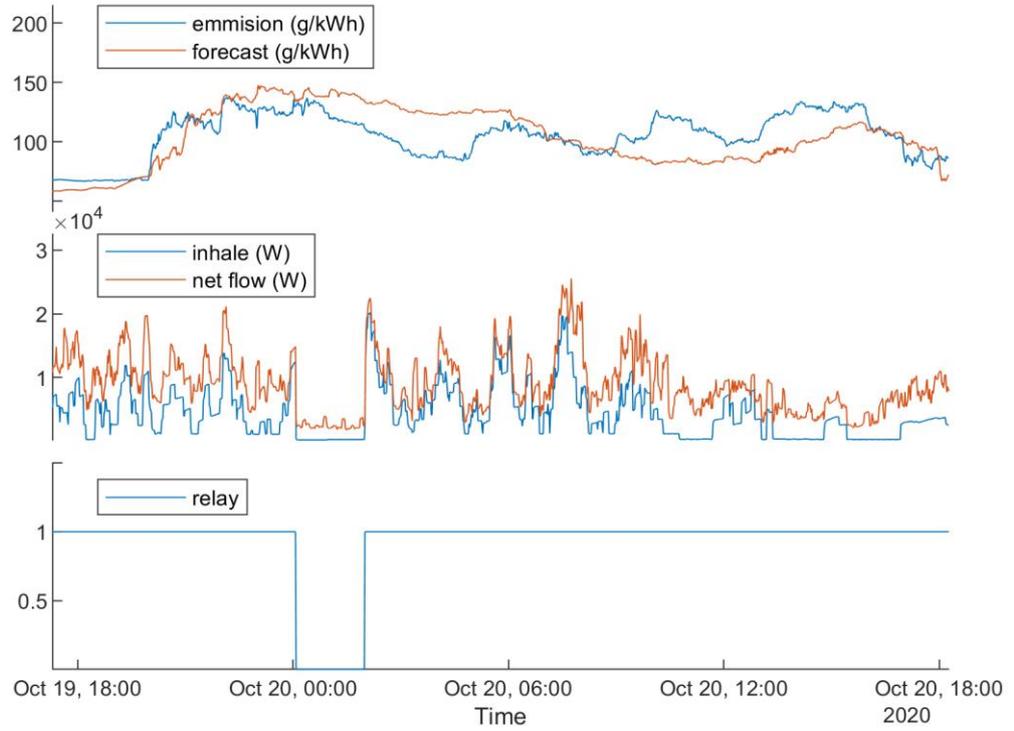
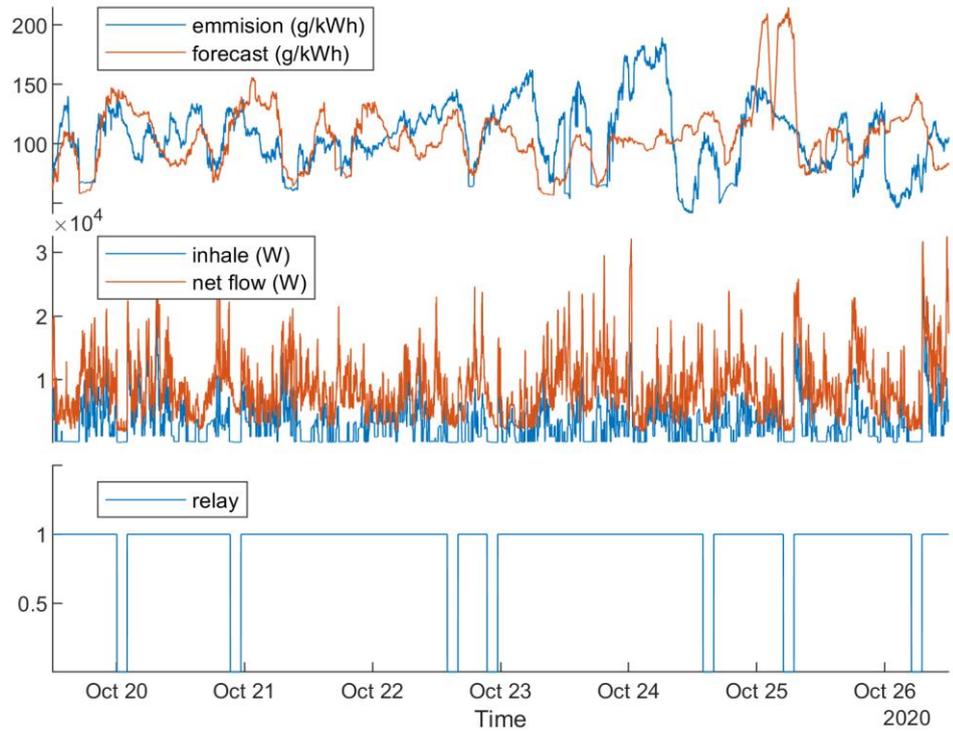


Figure 44: Illustration of the HP switching algorithm.



(a)



(b)

Figure 45: Load curves of 10 houses during "breathing with HP" for (a) one day and (b) 7 days.



Table 11: Summary of the measured CO₂ emissions during the experiment.

House ID	8	105	26	36	40	46	69	78	80	94	Sum
Total electricity consumption [kWh]	138	208	97	102	198	53	190	142	184	152	1466
Electricity consumption for heating (inhaling) [kWh]	48	80	45	31	71	10	63	39	94	70	551
Non shiftable Electricity consumption [kWh]	90	128	52	71	128	43	127	103	90	83	914
CO ₂ Emissions for HP [kg]	4.94	8.27	4.71	3.29	7.53	1.14	6.64	4.02	9.62	7.36	57.51

4.4.3 Impact of breathing on loads within a distribution grid

By switching off many HPs at the same time, the load in the grid can be significantly reduced. Figure 46 shows the measurements. When the relay was turned off, the accumulated load of the houses was reduced between 4 and 11 kW. When the relay was turned on again, the load was increased up to 7 kW. With such a small sample, the load increase/decrease varies heavily depending on the time of the day, the outside temperature and on coincidence. The more houses are included in such a load shifting, the more prominent and reliable these load decreases/increases become [9]. Compared to traditional ripple control, flexible shifting technology from aliunid makes such a demand management solution relevant for utilities soon, when there is a higher share of decentral renewable energy sources.

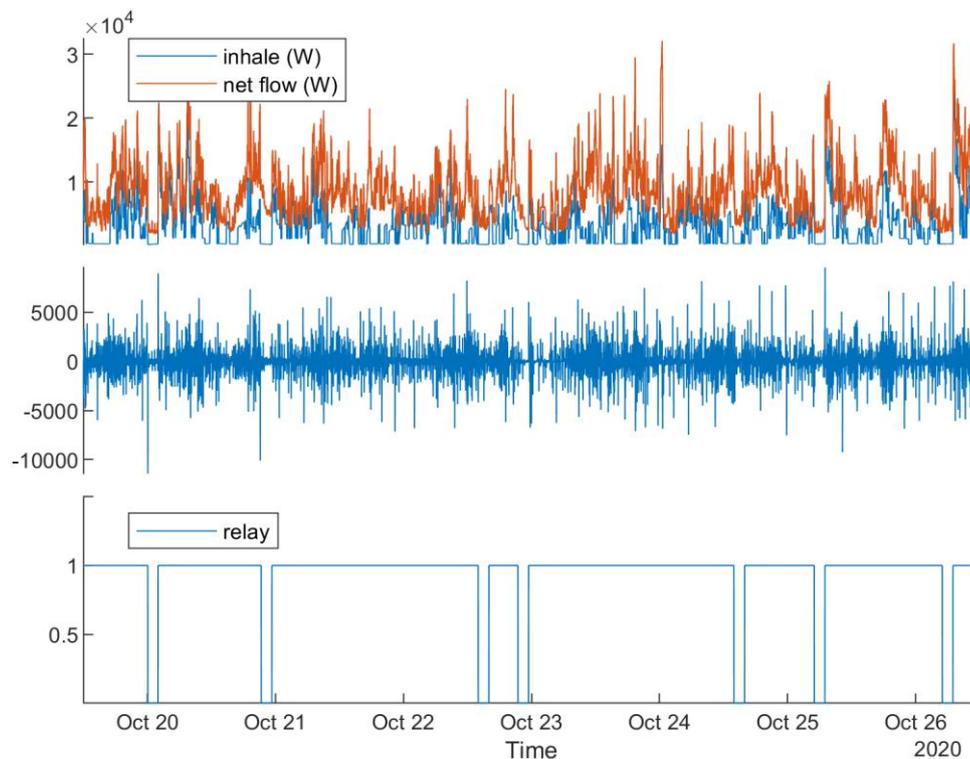


Figure 46: Changes in the accumulated power compared to the previous minute.



5 Conclusions

Authors: Till Richter, David Thiel, aliunid

Despite multiple challenges, a proof-of-concept implementation of an energy system that breathes was completed and its potential was investigated. A community of strong players in the Swiss energy market was formed, which is further pursuing the vision of a “new energy supply” with an integrated mindset. The competence to collect real-time data from different types of smart meters at low cost and store them securely in a private, distributed cloud was built and the hope is that commercial products based on the outcome of the project will impact and transform the energy market.

With the four-value model, an intuitive way on how to represent the flexibility of a household is proposed. On the one hand, the analogy of energy that breathes is simple enough to be understood by people that are not familiar with electricity. On the other hand, the four-value model and the grid-block structure allow to systematically aggregate the loads of individual households to guarantee data privacy without losing information on flexibility. These four values are the minimal set of information needed for energy management and optimization. The proposed model will be used within the community to identify and systematically exploit flexibility potential in energy systems, in a transparent way and in real-time.

During the project, it was demonstrated how to breathe based on an objective function that relates to carbon emissions of electricity consumption. However, besides these emissions, the developed controllers can also be used to optimize different cost functions, such as cost, self-consumption, or grid stability. The competence of calculating and forecasting CO₂ emissions of electricity in real-time is novel in Switzerland. However, research during the project indicated that this is by no means a trivial task, given the low degree of reliability of external data sources and the existence of different approaches on how to calculate emissions in a realistic way. Despite these difficulties, the proposed approach may open new opportunities in carbon emission trade, offering an alternative to the current certificate-based approaches. Being able to show end customers the real-time carbon emissions related to their individual electricity consumption may raise the necessary awareness to increase the interest in renewable energy and create market opportunities for new business models that target climate neutrality.

In Swiss households, there is a big variety of different electricity meters installed. Although Switzerland is currently rolling out smart meters, there is still an ambiguity in the open customer interface and load profiles are only stored with a resolution of 15 minutes. Within the field test, different approaches have been described on how to read real-time data, both from old legacy meters as well as new smart meters, and over 500 real-time data systems have been installed, integrated, and tested. The experience during the field test showed that installing or integrating different smart meter systems is not an easy endeavor, due to the lack of common standards. Moreover, establishing reliable data connection to the basement of houses, where most of the electricity meters are located, is not straightforward and has caused a lot of labor and cost-intensive troubleshooting. The continuous work on the distributed IoT system did ultimately lead to a robust, affordable, and secure data infrastructure, with a fully automated provisioning process that is necessary for the scaling of such a distributed system. Furthermore, it was demonstrated how to use relatively low cost IoT hardware (commodity) to retrofit existing HP and DHWT systems with the capabilities of demand side management. Nevertheless, once the smart grid ready interface of heat pumps will become more widely available, the control of heat pumps will be easier, and the flexibility will increase. The data infrastructure at customers' home will be the foundation of a breathing system and will allow to monitor, analyze, and control energy flows in real-time.

During the field test it became evident that the best place to store the “intelligence” (energy data, etc.) is not the smart meter. Electric infrastructure has a long lifecycle (up to 20 years). To avoid running into legacy issues with today's meter systems, the data processing unit needs to be decoupled from the data acquisition unit. Therefore, clear standards for the data interfaces are key to take full



advantage of the opportunities of digitalization. A recommendation of such a standard has been provided to SFOE as a separate work package. This report describes how to realize energy data access in an effective and efficient manner and the main findings are summarized in (Section 10.5.)

Considering that the energy grid is critical infrastructure, the IT solution developed during this project is based on a novel swarm cloud. This decentral architecture does not rely on central cloud computing services, which are susceptible to cyber-attacks. It instead creates a network of cells that are all able to function autonomously. The swarm cloud is perfectly suited to the four-value model and is not only resilient against cyber-attacks but also guarantees data privacy: The sensitive load information is only stored at the end user's device. At the next grid level, only aggregated values are available, which do not allow to draw any conclusions about the end user.

During the project, the following elements and methods were used to lay the basis for transforming the existing system into a "new energy supply" and the proposed breathing technology was demonstrated in three different use cases:

- The validation of **breathing in a single house** to minimize CO₂ emissions using a model predictive controller that interacts with a battery showed promising results, demonstrating a 20% improvement in the objective over the baseline. However, the performance could be improved further if real-time data from electricity production and imports were available, such that at least the first part of the calculated control trajectory is based on actual rather than forecasted information on carbon emissions.
- **Breathing in multiple houses** using DHWT was demonstrated successfully. Until now, many utilities use ripple control to control these loads based on a fixed schedule. By scheduling boilers during the most optimal time between 1pm and 11pm instead of 3am every day, the CO₂ footprint of warm water production was reduced by 25% according to the adopted metric. The CO₂ savings could be further improved if more information on the DHWT was available, e.g., the boiler temperature. This would however result in more complicated algorithms and require more hardware modifications in the houses.
- The third use case did demonstrate **breathing with HPs** in ten different houses within a distribution grid to optimize CO₂. It was shown that the basic principle works, and the operation of the heat pump was modified as planned. However, for a realistic quantification of the saved emissions aliunid would need to compare the measured load profile against a baseline profile, which is not trivial and subject to further work. Furthermore, the experiment demonstrated that the load of the grid could be reduced by up to 10kW, which is equivalent to a 30% reduction of the loads of the considered houses during periods of high carbon intensity.

In summary it can be concluded that compared to HP and DHWT, which can only be controlled by turning them off/on, the flexibility of batteries provides more freedom to breathe. The flexibility potential of HP and DHWT could increase if the thermal behavior of the house were be better modeled, which however would require additional instrumentation and limit fast rollout.

In all three cases, the controller did work as expected and considerable carbon savings have been demonstrated for the first and second use case, at least based on the chosen metric. It remains to be seen whether shifting strategies according to this metric will eventually improve the carbon intensity of the consumed electricity. Moreover, the demonstrated savings are also strongly depending on season and weather and it is expected that savings in a similar experiment during the summer would be lower. Furthermore, it must be stated that shifting loads from night to daytime is currently causing higher costs, since most utilities are using inflexible high and low tariff systems which are not supporting load demand management.

In addition to the three use cases, the effects of breathing in a bigger scale have been studied (Section 10.2). Based on a data driven approach the available flexibility options in distribution grids has been investigated and the needs for flexibility in grids with increasing PV production have been shown. This analysis highlights when critical times within the day/week/year occur within the network. With the capability of storing or shifting energy within a 12-hour window, the self-consumption (at EVU



level) or reduction of import/export (on Swiss level) could be significantly improved. Intelligent demand side management solutions of residential storage capacities (heating, warm water, stationary battery systems) are one possibility to achieve this in the future. For further improvement of the self-consumption however, other technologies like sector coupling are needed to store/shift energy up to 6 months.

This analysis shows the importance of a flexible system and justifies the efforts of this field test. Doing energy management holistically rather than individually helps to better achieve the targets of the Energy Strategy 2050. These first promising results are encouraging to continue this important work and scale a breathing system up to a size where it become relevant for the Swiss electricity supply.



6 Outlook and next steps

Authors: Till Richter, David Thiel, aliunid

The Swiss-wide field test “when energy breathes” was an exciting journey and helped the participants to learn a lot about flexible energy systems. However, the journey has not ended with the formal closing of the field test.

The field test demonstrated breathing in a basic setup to investigate its potential. So far, only few data points of breathing houses are available, and more work is needed to demonstrate breathing on a bigger scale. Important further technical questions are:

- How to solve large-scale distributed optimization problems to schedule loads in a network of interconnected systems? This aspect becomes relevant when the size of the managed loads is large enough to affect grid stability.
- How to deal with situations of critical grid stability and how to ensure that a breathing system is compliant with grid codes and does not interfere with primary and secondary control? A preliminary approach could be to implement central directives from the grid operator.
- Compared to the simple “open-loop” approach that was used so far to control HP and DHWT, real-time data would be essential in the implementation of a feedback controller, which would be better able to compensate for forecasting errors.

It became obvious in many cases that the availability of consistent real-time data is the beginning of any digitalization effort. Based on the learnings of this project it is recommended that utilities start recording data on their grid as soon and as broadly as possible. Experience from other industries show that data will be an asset that helps them to conduct studies and discover new applications of digitalization. In the future, it is expected that many new use cases will be discovered where data and digitalization is involved. Having data available will be the precondition to learn from data, find business opportunities, build data driven technology and create value, for example by increasing energy efficiency.

Besides the technical challenges, also economical and legislative questions are open.

- According to Swiss law, flexibility in residential houses is owned by the consumer. If these appliances are used for breathing, the customer must be motivated to let the utility manage the flexibility. Currently the profit opportunities of providing flexibility are limited, for demand side management as well as for supply side management. The question remains open how to make it more interesting for power producers and end customers to offer flexibility. Demand side management or supply side management solutions alone will not be profitable: The technology must be embedded in a viable business case and tailored to smoothly fit into the existing ecosystem while changes in the regulatory framework (e.g., grid tariffs that promote flexibility and encourage sector coupling) and market rules will become essential.
- In a similar direction goes the question on how to further optimize hardware, data connectivity and installation cost. Demand side management in residential houses is not yet a viable business case and the share of hardware and labor cost is relatively large compared to the margins. It seems currently most promising to include demand side management capabilities during installation of the appliances, which would avoid an expensive retro fit.

aliunid’s engagement on a breathing energy system will continue not only in further validation of the demonstrated use cases but also in implementing solutions that will be able to perform a continuous breathing. An overview of what was achieved and what are the immediate next steps in the different disciplines of aliunid is shown in Table 12 and Table 13. After closing this field test aliunid will focus on launching first commercial products for end customers. If Swiss people see how products positively impact their daily life, it is expected that their support in making the energy strategy 2050 happen will be bigger.



Breathing is not limited to use end user's flexibility. aliunid will further investigate how to include power plants in the flexible energy system, developing for example a solution to shift hydro production to times when there is no solar production.

aliunid is grateful for the early and steady support of Swiss Confederation represented by SFOE during the last 3 years. We learnt, that successfully bringing a vision like "when energy breathes" to earth is not possible without failures. Together we are strong!

Table 12: Next Steps regarding aliunid field test objectives

Objective field test	Results and relevance		Next Steps
Representation of the flexibility of a household in the four-value model in real-time	Sections 3.3 and 4.1.3	✓	Implement more broadly and scale from individual GB to GB structure
Calculation of a tariff (CO2 emission) in real-time for which the flexibility of a household or a distribution grid shall be optimized	Section 3.2	✓	Improve emission forecasting model, include additional data sources, and introduce new tariff functions
Robust, affordable, and secure acquisition of real-time data	Section 3.5	✓	Continue development and industrialize. Focus of integration of existing (smart) meters and start rolling out actors
Organization of data flows in a modern decentralized swarm cloud	Section 3.6	✓	
Controlling flexibility of a single-family household (SFH)	Section 4.2	✓	Continue developing optimal control
Controlling flexibility of several SFHs in a distribution network	Sections 4.3 and 4.4	✓	Continue learning by using real infrastructure, also after SFOE funded field test is over. aliunid plans to establish a new prototyping vehicle (aliunid InnoLab)

Table 13: Next steps working packages

Working packages	Results	Next Steps
#1: Regulation of a flexible energy supply: Currant legal framework hinders implementation of energy strategy 2050	Section 10.1	Focus on grid model and tariffs, data protection and customer interface. Addressing unbundling to avoid redundant infrastructure
#2: Technical flexibility: Demand side and supply side management	Sections 10.2, 10.3, 10.4	Flexibility in Grid: implement dynamic load flow calculation
#3: Digital infrastructure: Data protection and cyber security	Sections 10.5 and 10.6	Down-selecting R&D portfolio



7 National cooperation

Many thanks to our community, whose support and cooperation helped to make this project a success.

Table 14: National field test partners.

Field test partner	Abbreviation
Utilities	
Licht- und Wasserwerk Adelboden AG	lwa
Yverdon-les-Bains Énergies	ylbe
Energie Uster	Uster
Genossenschaft Elektra	elektra
Industrielle Betriebe Interlaken	ibi
Industrielle Betriebe Wohlen	ibw
Localnet AG Burgdorf	localnet
Industrielle Betriebe Murten	ibm
Société Electrique Intercommunale de la Côte SA	seic
Oiken (formally Énergies de Sion-Région SA)	
Energie Brig-Aletsch-Goms	EnBAG
Services Industriels de Genève	SIG
Städtische Werke Grenchen	SWG
Producers	
Azienda Elettrica Ticinese	AET
FMV	
Alpiq	
Universities	
ETH Zurich	
EMPA	
BFH	
HSG	
HES-SO	
IoT Developer	
Appmodule AG	



8 Publications

Table 15: List of publications

Title	Topic/Headline	Date	Authors	Medium
energate messenger.ch	Startup-Serie, im Portrait: aliunid	18.09.20	Mario Graf	Webportal, Online-Magazin, Newsletter
Hes-so Vaiais Vallis - Haute École d'Ingénierie Hochschule für Ingenieurwissenschaften	Senkung des Stromverbrauchs dank intelligenter Gebäude	17.09.20		Web, Soziale Medien
higgs	Energiewende: wenn in Haus, ein Quartier oder eine Stadt Energie ein- und ausatmet	19.06.20	BFH	Webportal, Online-Magazin
Swiss Environment & Energy, Innovation Monitor	Get the 2020 Energy Startup Map	06.05.20	Christina Marchand	Webportal
Geschäftsbereich 2019, Energie Uster AG	Ein atmendes System für die Zukunft	17.07.20		Geschäftsbericht
rundum Ausgabe 05, Energie Uster AG	Ein Netz, das atmet	01.12.19		Kundenmagazin
startupticker.ch	aliunid startet Feldversuch mit "atmender Stromversorgung"	10.05.19		Online-Magazin
Aargauer Zeitung	IB Wohlen teste Stromversorgung der Zukunft	08.05.19	Fabio Vonarburg	Print, Webportal
SRF	Erneuerbare Energien: Was tun, wenn der Wind nicht weht?	07.05.19	Klaus Bonanomi	Radio, Webportal
energate messenger.ch	Startup will neues Versorgungssystem schaffen	07.05.19	Mario Graf	Webportal, Online-Magazin, Newsletter
blue News	Testlauf für künftige Energieversorgung der Schweiz gestartet	07.05.19	SDA	Webportal, Online-Magazin, Newsletter



nau.ch	Testlauf für künftige Eenergierversorgung der Schweiz wird gestartet	07.05.19	SDA	Webportal, Online-Magazin
Le Nouvelliste	Energie intelligente: quand les lave-linges parlent aux barrages	07.05.19	SDA	Print, Webportal
Le Nouvelliste	aliunid, du test national à l'épreuve cantonale. Le commentaire de Pascal Fauchère	07.05.19	Pascal Fauchère	Print, Webportal
La Liberté	Une start-up teste l'alimentation électrique de l'avenir en Suisse	07.05.19	ATS	Print, Webportal
AGEFI	Une start-up teste l'alimentation électrique de l'avenir en Suisse	07.05.19	AWP	Print, Webportal
Inside IT	Schweizer Schwarm-IoT für flexible Energie-Verteilung	07.05.19	Thomas Schwendner	Webportal, Online-Magazin
NZZ	Mit Schweizer Cloud gegen den Flatterstrom	21.03.19	Giorgio v. Müller	Print, Webportal



9 References

- [1] A. o. t. P. A. F. C. o. C. Change, «Adoption of the Paris Agreement. Framework Convention on Climate Change,» Adoption of the Paris Agreement. Framework Convention on Climate Change, Paris, France, 2015.
- [2] Bundesamt für Energie BFE, «Energy Strategy 2050,» [Online]. Available: <https://www.uvek.admin.ch/uvek/en/home/energy/energy-strategy-2050.html>.
- [3] J. Gasser, «Demand side flexibilities on a building level,» ETH, Zurich, 2020.
- [4] M. Rüdüsüli, S. L. Teske und U. Elber, «Impacts of an Increased Substitution of Fossil Energy Carriers with Electricity-Based Technologies on the Swiss Electricity System,» 2019.
- [5] S. Karagiannopoulos, «Data-driven and Optimization-based Methods for Active Distribution Grids,» ETH Zürich, 2019.
- [6] Y. Sidqi, P. Ferrez, D. Gabioud und P. Roduit, «GOFLEX: Towards a New Integrated Solution for a Cost-Effective Use of Demand Response in Distribution Grids,» Berlin, 2020.
- [7] G. Meier, S. Liechti und M. Timm, «Energiemanagementsysteme: Digitales Werkzeug der Energieversorgung,» Energie Zukunft Schweiz, 2020.
- [8] Bundesamt für Energie BFE, «Schweizerische Gesamtenergiestatistik 2018,» 2018.
- [9] Y. Sidqi, P. Ferrez, D. Gabioud und P. Roduit, «Flexibility Quantification in Households- A Swiss Case Study,» Sierre, 2020.
- [10] G. Wernet, C. Bauer, B. Steubing, J. Reinhard, E. Moreno-Ruiz und B. Weidema, «The ecoinvent database version 3 (part 1): overview and methodology,» *The International Journal of Life Cycle Assessment* 21, pp. 1218-1230, September 2016.
- [11] Bundesamt für Energie BFE, «Schweizerische Elektrizitätsstatistik 2019,» 2020. [Online].
- [12] «Swissgrid,» [Online]. Available: <https://www.swissgrid.ch>.
- [13] «entsoe - Transparency Platform,» [Online]. Available: <https://transparency.entsoe.eu>.
- [14] M. Jakob, K. Volkart und D. Widmer, «CO₂-Intensität des Stromabsatzes an Schweizer Endkunden,» TEP Energy, Zürich, 2009.
- [15] R. Elliot, P. Hollmuller und M. Patel, «Émissions horaires de gaz à effet de serre liées à la consommation d'électricité - une approche incrémentale pour une économie ouverte: Le cas de la Suisse,» archive ouverte UNIGE, Genève, 2018.
- [16] A. Hawkes, «Estimating CO₂ emission rates for national electricity systems,» *Energy Policy*, Nr. 38, pp. 5977-5987, 2010.
- [17] «Computer implementiertes Verfahren zur dezentralen Steuerung eines Energienetzes». Patent PMURM001WO-2019104585.
- [18] J. B. Rawlings, D. Q. Mayne und M. M. Diehl, Model Predictive Control: Theory, Computation and Design, 2017.
- [19] Bundesamt für Energie BFE, «Smart Grid Roadmap,» 2015. [Online]. Available: <https://www.admin.ch/gov/de/start/dokumentation/medienmitteilungen.msg-id-56706.html>.
- [20] R. D'hulst, W. Labeeuw, B. Beusen, S. Claessen, G. Deconinck und K. Vanthournout, «Demand response flexibility and flexibility potential of residential smart appliances: Experiences from large pilot test in Belgium,» *Applied Energy*, pp. 79-90, 1 Oktober 2015.



- [21] K. Heussen, C. Steinbrink, I. Abdulhadi, V. Nguyen, M. Degefa, J. Merino, T. Jensen, H. Guo, O. Gehrke, D. Bondy, D. Babazadeh, F. Pröstl Andrén und T. Strasser, «ERIGrid Holistic Test Description for Validating Cyber-Physical Energy Systems,» *Energies*, Nr. 12, 2019.
- [22] Federal Department of Foreign Affairs FDFA, «Energie - Fakten und Zahlen,» 2019. [Online]. Available: <https://www.eda.admin.ch/aboutswitzerland/en/home/wirtschaft/energie/energie—fakten-und- zahlen.html>. [Zugriff am 26 April 2020].
- [23] B. f. S. BFS, «Road vehicles - stock, level of motorisation.,» 2020. [Online]. Available: <https://www.bfs.admin.ch/bfs/en/home/statistics/mobility-transport/transport-infrastructure-vehicles/vehicles/road-vehicles-stock-level-motorisation.html>. [Zugriff am 24 Juli 2020].
- [24] Federal Bureau of Statistics, «Energy field,» 2020. [Online]. Available: <https://www.bfs.admin.ch/bfs/en/home/statistics/construction-housing/buildings/energy-field.html>. [Zugriff am 26 April 2020].
- [25] F. Bünning, B. Huber, P. Heer, A. Aboudoni und J. Lygeros, «Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings,» *Energy and Buildings*, 2020.
- [26] Bundesamt für Energie, BFE, «Grundlagen der Ausgestaltung einer Einführung intelligenter Messsysteme beim Endverbraucher in der Schweiz,» 2014.
- [27] «IEC 62056,» [Online]. Available: en.wikipedia.org/wiki/IEC_62056.
- [28] Netbeheer Nederland –WG DSM, «Dutch Smart Meter Requirements - P1 Companion Standard,» 2016.
- [29] K. Garifi und K. Baker, «Non-Simultaneous Charging and Discharging Guarantees in Energy Storage System Models for Home Energy Management Systems,» 2018.
- [30] «electricityMap: Visualisierung der klimarelevanten Emissionen der Stromerzeugung,» [Online]. Available: <https://www.electricitymap.org/map>. [Zugriff am 28 Juli 2020].
- [31] Bundesamt für Energie BFE, «Stand der Wasserkraftnutzung in der Schweiz am 31. Dezember 2019,» 2020. [Online].
- [32] B. Tranberg, O. Corradid, B. Lajoied und T. Gibonelain, «Real-time carbon accounting method for the EUropean electricity markets (and Supplementary Material),» *Energy Strategy Reviews*, Bd. 26, Nr. 100367, 2019.
- [33] «Global Energy Observatory,» [Online]. Available: <http://globalenergyobservatory.org/geoid/45622>. [Zugriff am 28 Juli 2020].



10 Appendix

10.1 Summary: Governing decentral energy systems

Author: Bianca von Preussen, HSG

Die Energiestrategie 2050 führt zu einem Ausbau der **dezentralen erneuerbaren Elektrizitätserzeugung** mit dargebotsabhängiger Produktion. Sie erfordert Anpassungen in der Topologie und Bewirtschaftung der Netze, ein Management der Energienachfrage sowie einen Ausbau und systemdienlichen Einsatz von Speichern (Batterien, Power-to-X, etc.). Herz eines solchen Systems ist ein intelligentes Netz («Smart Grid»), das eine effiziente Steuerung bzw. den Ausgleich von Verbrauch und Angebot mittels innovativer Informations- und Kommunikationstechnologien ermöglicht. Grundlage dafür sind Smart-Metering-Systeme, die die relevanten Erzeugungs-, Verbrauchs- und Speicherdaten der Nutzer in Echtzeit erheben und für den systemdienlichen Einsatz zur Verfügung stellen.

In diesem Umfeld entstehen Möglichkeiten für **neue Geschäftsmodelle und Energiedienstleistungen**, wie sie von «aliunid» angeboten werden. Die vorliegende Studie identifiziert Hürden für die Verwirklichung solcher Geschäftsmodelle, die noch durch die Regulierung der «alten» Energiewelt begründet sind und daher potenziell wohlfahrtshemmend und anzupassen sind. Dabei wurden gemäss Förderauftrag des BFE folgende Bereiche einer näheren Untersuchung unterzogen: Datenschutz und Datensicherheit; Smart Meter – Echtzeitdaten und Nutzung; Restenergiebezug – Produkte und Bepreisung; Effiziente Netzkapazitäten; Atmung - Kompetenzen und Anreize.

Im Kern des Geschäftsmodells von «aliunid» steht die (systemdienliche) **Nutzung und monetäre Verwertung von Flexibilitäten**, die heute dem Erzeuger und Verbraucher obliegt. Auch wenn explizite Regelungen betreffend Flexibilitäten noch ausstehen (VE StromVG), so ist die vertragliche Nutzung und gebündelte Verwertung der Flexibilitäten durch Dritte (EVU, aliunid, andere Energiedienstleister) dennoch möglich, wenn auch aus regulatorischen Gründen mit Einschränkungen. Ein weiterer Pfeiler des Geschäftsmodells von «aliunid» bildet das **datengestützte Angebot von dezentralen Flexibilitätsmärkten**. Auch ein solches Angebot ist heute gestützt auf vertragliche Vereinbarungen verwirklichtbar. Dennoch ergeben sich mit Bezug auf die untersuchten Bereiche folgende Anregungen zur Weiterentwicklung des regulatorischen Rahmens und zur Erleichterung des Übergangs in die «neue» Energiewelt:

Datenschutz: Eine intelligente Steuerung des Energieverbrauchs und der Energieerzeugung ist im Einklang mit den Datenschutzbestimmungen durch die Beschränkung der verwendeten Smart-Meter-Daten im erhobenen Haushalt und durch die Aggregation der Daten durch den lokalen Versorger möglich. Eine weitergehende Nutzung der Smart-Meter-Daten, z.B. für energiefremde Dienstleistungen ist nur möglich, wenn dies für die Erfüllung eines Vertrages (z.B. mit dem EVU) notwendig ist oder eine Einwilligung gegeben wurde. Die Konformität der Datenverarbeitungen mit dem Datenschutzrecht kann also vertraglich sichergestellt werden.

Datensicherheit: Vertrauen in die Sicherheit der generierten Daten sind ein Eckpfeiler der Akzeptanz eines intelligenten Energieversorgungssystems. Hierzu bestehen erste regulatorische Ansätze, Richtlinien und Branchenstandards. Eine detaillierte gesetzliche Regelung gestaltet sich schwierig, da technische Sicherheitsmöglichkeiten und -standards laufenden Weiterentwicklungen unterworfen sind. Branchendokumente sind eher geeignet, den technischen Fortschritt abzubilden und enthalten hierzu bereits heute Regelungen. Anreize zur Übertreffung des Mindeststandards bestehen weder in monetärer noch in regulatorischer Sicht, insbesondere mit Blick auf die Gewährleistung der systemweiten Versorgungssicherheit. Unternehmen, welche die Datensicherheit durch besondere Massnahmen erhöhen, können diesen Umstand gegenüber ihren Kunden immerhin anpreisen, so auch «aliunid». Weitergehende Ergebnisse werden im Rahmen des Arbeitspaketes 3 dieses Projektes «Digitale Infrastruktur und Sicherheit» erwartet.



Smart Meter – Echtzeitdaten und Nutzung: Auch wenn die regulatorischen Mindestvorgaben an intelligenten Mess-, Steuer- und Regelsystemen heute keine echtzeitnahe Lastgangmessung vorsehen, sollte es den Messstellenbetreibern freistehen, Geräte zu installieren, die eine höhere Granularität der Daten gewährleisten. Die Betriebs- und Kapitalkosten von Steuer- und Regelsystemen können via Netznutzungsentgelte der Gesamtheit der Endverbraucher angelastet werden. Die Kapital- und Betriebskosten von Dienstleistungen, die auch von Dritten (z.B. von «aliunid») angeboten werden können, sind nicht anrechenbar und deshalb von den Netzkosten klar und nachvollziehbar abzugrenzen.

In diesem Zusammenhang sind die Entflechtungsvorgaben gemäss Art. 10 Abs. 2 StromVG zu beachten, deren Verletzung zur Strafbarkeit gemäss Art. 29 StromVG führt. Sofern es sich bei den Smart-Meter-Daten um Netzdaten handelt, ist deren Verwendung für netzfremde Dienstleistungen ausgeschlossen. Diese Beschränkung der Nutzbarkeit von Smart-Meter-Daten auf reine Netzanwendungen erscheint angesichts der Investitionskosten für Smart-Meter-Daten als nicht sachgerecht. Mindestens müsste der Endkunde oder Produzent in die Verwendung der Smart-Meter-Daten für aussernetzliche Anwendungen einwilligen können. Immerhin zeigt das von «aliunid» angestrebte Modell, dass in einem gewissen Rahmen auch unter bestehendem Recht Dienstleistungen auf Basis der Smart-Meter-Infrastruktur möglich sind. «Aliunid» sieht vor, über eine (diskriminierungsfrei vom Netzbetreiber angebotene) offene Schnittstelle des Smart-Meters Daten selbst zu generieren. Ein solches Vorgehen verstösst nicht gegen Art. 10 Abs. 2 StromVG.

Restenergiebezug – Produkte und Bepreisung: Die Abhängigkeit der Prosumers vom Verteilnetz nimmt stetig ab. Dennoch wurde die Solidarität unter den Eigenverbrauchern bereits gestärkt. Für kleine Endverbraucher mit einer Anschlussleistung von bis zu 15 kVA gibt es nur eine Kundengruppe, egal ob diese über eine Produktionsanlage oder einen Speicher verfügen, wodurch die Netzkosten nach verteilungspolitischen Gesichtspunkten angelastet werden.

Bezüglich der Netznutzungstarife gilt im bestehenden Netznutzungsmodell das Ausspeiseprinzip, d.h. dass das Netznutzungsentgelt von den Endverbrauchern je Ausspeisepunkt zu entrichten ist. Bei den Endverbrauchern werden Kraftwerke und namentlich PSKW für ihren Eigenbedarf explizit ausgenommen, d.h. sie haben kein Netznutzungsentgelt zu entrichten. Das heutige Netzentgeltsystem ist statisch. Unterjährige Anpassungen sind nicht vorgesehen. Grundsätzlich zulässig wären auch Wahltarife und dynamische Tarife. Auch bei den Energieliefertarifen wären dynamische Tarife unter Einhaltung der gesetzlichen Vorgaben zulässig.

Effiziente Netzkapazitäten – Kompetenzen und Anreize: Auch wenn für die Nutzung von Flexibilitäten kaum regulatorische Vorgaben bestehen, können Dritte (Endverbraucher, Erzeuger oder Speicher, die beide Eigenschaften in sich vereinigen) bereits heute mit den Netzbetreibern den Einsatz eines intelligenten Steuer- und Regelsystems vertraglich regeln. Den technischen Möglichkeiten des Marktes (z.B. Anbietern wie «aliunid») ist es überlassen, Mehrwertdienste und Lösungen zu entwickeln, die auf dieser vertraglichen Nutzung des Netzbetreibers aufsetzen. Entsprechende Investitionen sind dann profitabel, wenn durch verschiedene Dienstleistungen mehrere Erlösströme erzeugt werden können.

Allerdings ist das auf der Grundlage von Art. 17b Abs. 2 und 3 StromVG i.V.m. 8c StromVV geschaffene Flexibilitätsmodell nicht mit den bestehenden Entflechtungsvorgaben gemäss Art. 10 Abs. 2 StromVG abgestimmt. Auch wenn das Geschäftsmodell von «aliunid» davon nicht unmittelbar betroffen ist, soll an dieser Stelle nochmals darauf hingewiesen werden. Smart-Meter-Daten müssen aus dem Netzbetrieb stammen und dürfen im Sinne des Unbundling nicht für andere Tätigkeiten des Netzbetreibers als Netznutzung und Netzanschluss genutzt werden. Unter andere Tätigkeiten fällt die wettbewerbliche Tätigkeit eines integrierten Netzbetreibers. Darunter fällt namentlich die Nutzung für den Verkauf von Strom an freie Kunden.

Atmung – Kompetenzen und Anreize: Die Versorgungssicherheit ist von übergeordneter Bedeutung. Es besteht ein verstärkter Bedarf an Zwischenspeicherung von Energie, da die Produktion aus dargebotsabhängigen Quellen zeitlich nicht notwendigerweise mit dem Verbrauch zusammenfällt. Speicher werden zwar zunehmend in den gesetzlichen



Bestimmungen genannt, beinhalten aber keine Konkretisierungen. Vor allem die Unsicherheit bezüglich der anrechenbaren Netzkosten von Speichern und einer allfälligen Vergütungsfähigkeit als Netzverstärkung bilden ein wesentliches Hemmnis für die Einführung innovativer Lösungen. Bezüglich dem Netznutzungsentgelt werden Stromspeicher für die aus dem Netz bezogene Elektrizität derzeit als Endverbraucher betrachtet, verbunden mit der Pflicht zur Entrichtung eines Netznutzungsentgeltes. Soweit die Speicherung von Elektrizität nicht für den Eigenverbrauch erfolgt, ist fraglich, ob eine solche Betrachtung korrekt ist oder für Speicher die PSKW-spezifischen Ausnahmen gelten sollten.

Eine ganzheitliche Betrachtung von Energieversorgungssystemen unter Verbindung der Sektoren Strom, Gas und Wärme/Kälte wird heute durch die bestehenden, sektorspezifischen Regelungen im Energiebereich erschwert. In den anderen gekoppelten Märkten (Gas, Fernwärme, etc.) bestehen heute keine Regelungen diesbezüglich. Die Sektorkopplung ist kein Grund, in diesen Märkten ebenfalls eine umfangreiche Regulierung einzuführen. Sie würde sicherlich innovationshemmend wirken. Vielmehr sollte die Stromgesetzgebung für die Sektorkopplung geöffnet werden.

Die Nutzung von Flexibilitäten, wie im Geschäftsmodell von «aliunid» vorgesehen, ist heute regulatorisch möglich. Flexibilitäten können auf vertraglicher Basis gesichert und vermarktet werden. Die datengestützte Systemoptimierung ist datenschutzkonform durchführbar. Ob sich Systemoptimierer wie «aliunid» zukünftig allein über solche Energiedienstleistungen finanzieren können, hängt allerdings in grossen Teilen von den regulatorischen Rahmenbedingungen ab: Haushalte und Prosumers werden von Knappheitssignalen des Marktes abgeschirmt; die Netzkostenregulierung setzt wenig Anreize zur lokalen oder regionalen Netzoptimierung; für zentrale Pumpspeicher gelten heute noch PSKW-spezifische Ausnahmen; lokale Optimierungen durch Sektorkopplung stossen im Gegensatz zum blossen Kapazitätszubau auf keine besonderen Anreizstrukturen.

Der derzeitige Wandel stellt für die Unternehmen der Energiebranche **Herausforderung und Chance** zugleich dar. Die traditionelle Kernaufgabe der Energieversorgung wird durch datengestützte Dienste ergänzt. Diese Entwicklung verändert den Wettbewerb und bietet den Energieunternehmen die Chance, sich gegenüber Wettbewerbern deutlicher abzugrenzen und mit attraktiven Zusatzdiensten bestehende Kunden zu binden und neue Kunden zu gewinnen. Obwohl diese «neue» Energiewirtschaft im Grundsatz auf generelle Zustimmung stösst, bestehen hinsichtlich des regulatorischen Rahmens zwar keine unumstösslichen Hindernisse, doch auch keine besonderen Anreizstrukturen zur Förderung dieses Wandels.



10.2 Summary: Identifying critical times and corresponding flexibility options

Authors: Natasa Vulic, Martin Rüdüsüli, Empa

10.2.1 Introduction

With the increasing trend in PV production levels, each DSO is experiencing unique challenges in balancing production and consumption volumes within their network areas. Estimating the ability of individual DSOs to effectively self-consume PV production surpluses – both current and future – relies on the careful analysis of their consumption and production patterns. This analysis helps us identify critical times within the day/week/year that may benefit from a more flexible system, as well as options and existing potential for making the respective networks more flexible.

The work presented here focuses on developing a generalized methodology to investigate flexibility needs and options of DSOs. Rather than using models, we adopt a data-driven approach for analyzing production and consumption patterns that make it possible to pinpoint critical times and investigate existing potentials for shifting production and consumption volumes. Using measured profiles, we developed methodologies for addressing the following questions:

1. How can we determine flexibility needs of a network area? Do these needs arise due to temporal or spatial mismatch between consumption and PV production?
2. How do flexibility needs change throughout the day/week/year? When are they highest/lowest? How can we identify them?
3. What are the benefits of daily load-shifting for maximizing regional PV self-consumption? What are its limits?
4. What are the available options for making the network more flexible? What volumes can be expected and already exist in parts of the grid?

These methodologies are briefly outlined below. To demonstrate them, we also present the preliminary analysis of the two participating utilities (*EVUs*).

10.2.2 Data

Network area inflows and outflows, including production and consumption profiles, are based on EVU data (supplied at 15-min resolution for the year 2019) provided by aliunid, and pre-processed by Empa.

10.2.3 Flexibility need indicators (spatial vs. temporal)

The flexibility needs of individual EVUs depend on the mismatch between consumption and PV production. To evaluate both the degree and type of mismatch experienced by the EVUs, we develop and compare certain indicators of self-consumed PV production. Specifically, we compare the *net* surplus PV production (i.e., after concurrent inflows and outflows at the network boundary are balanced out for all the connection points to the neighboring network areas) to the *actual* surplus PV production (i.e., sum of all the individual outflows). To enable the comparison between different EVUs, these values are expressed as a ratio (percentage) to the total consumption. The former (*net*) indicates the degree of temporal mismatch in the network area, whereas the latter (*actual*) indicates the combined temporal and spatial mismatch. By comparing the values of these two indicators, the following observations can be made:

- If *net* surplus PV production = *actual* surplus PV production, the mismatch between consumption and production within the network area is mainly temporal.
- If *net* surplus PV production < *actual* surplus PV production, there is an additional spatial mismatch, with the degree of that mismatch indicated by the difference between these values.



These indicators can then be used to provide tailored recommendations to the EVUs. In the first use case, we compared these indicators for two EVUs (A and B) and observed significant differences between them. These two EVUs have comparable ratios of total PV production to total consumption of 8.7 and 11.9%, respectively. Nevertheless, after PV production surpluses are taken into account, the resulting (consumed) PV production slightly decreased for EVU A (to 8.5%), but quite dramatically for EVU B (to 7.6%), when compared to their total PV production. While their *net* production surpluses revealed negligible temporal discrepancies for their respective network areas (non-existent for EVU A, negligible for EVU B), their *actual* production surpluses revealed stark differences: 1.7% vs. 36%, respectively. Based on these observations, the recommended approach for EVU A is therefore more generalized focusing on the anticipated flexibility needs (for the entire network area), while the one for EVU B, requires a more immediate localized approach. Nevertheless, both can benefit from increased digitization at subordinate network levels (transformer station, distribution cabin, end-consumer level) to better understand where the mismatch arises (or may arise in the future) and how to address it.

While these indicators give us an overall summary of flexibility needs for the entire year, a closer look at production and consumption patterns (daily, weekly, seasonal) is necessary to better understand *when* they occur (for both current and future production levels).

10.2.4 Patterns of production and consumption

The daily profiles of consumption and PV production over the course of the year are clustered by season and weekday-weekend. For each cluster, average daily production and consumption profiles are determined, together with the corresponding min-max ranges. Within each cluster, the production and consumption profiles are compared to each other, as well as across different clusters with respect to profile shape, magnitude, and variability.

In Figure 47, the average daily consumption and production patterns (in blue and orange, respectively) are provided for EVU A. The profiles are shown for four months representative of each season, with consumption profiles also separated between weekdays and weekends. The corresponding shaded areas show the ranges for each of these loads at each time-step, an indication of their variability.

From the comparison of the PV production and consumption profiles, the following observations can be made:

- Consumption has a distinct weekday/weekend drop due to reduced productivity (rows).
- The production and consumption ranges for the full year (top row, shaded areas) capture the seasonal variation for weekdays and weekends separately.
- Maximum daily PV production, as well as the number of hours available, is seasonally-dependent (increasing from January to July, and decreasing again in October).
- Maximum daily consumption is seasonally-dependent in the opposite direction (decreasing from January to July, increasing again in October for both weekdays and weekends).
- A peak following sunset (indicated by dashed vertical lines) is seasonally dependent (most prominent in January, reappearing in October) likely due to increased demand for public and private lighting; consequently, the peak increase also coincides with the end-of-day drop in PV production.
- Largest mismatch between production and consumption is observed in January during weekdays; smallest mismatch is observed in July during weekends (red arrows).
- PV production range, an indication of variability, appears highest in Summer and Spring months.

Localized analysis of EVU B, with a focus on a residential area with prosumers, will be included in the full report.

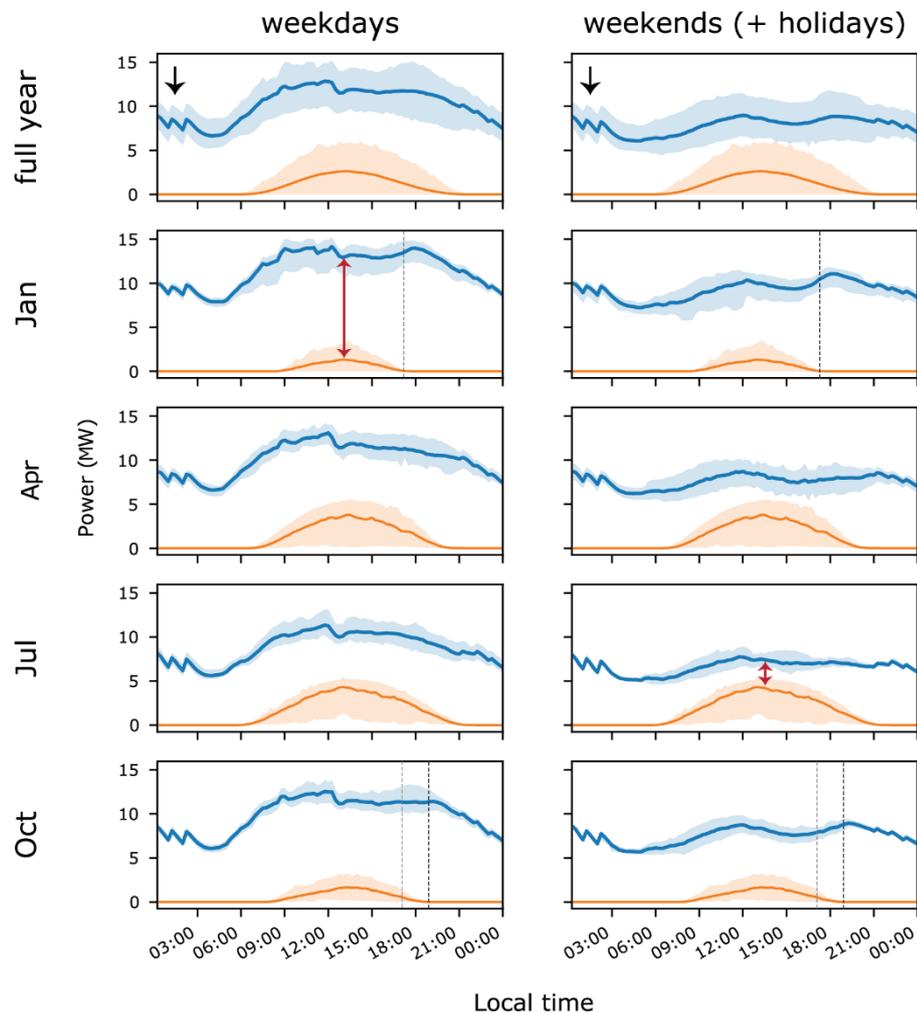


Figure 47: Average daily consumption (blue lines) and PV production (orange lines) profiles by weekday- weekend (columns) and months representative of the four seasons (rows) with minimum and maximum ranges (shaded areas). The largest and smallest load mismatch between PV production and consumption occurs in January during weekdays and in July during weekends, respectively (red arrows). Regularly scheduled loads are also observed (black arrows). When day-lighting hours are short, an additional evening peak follows the sunset times (shown for January and October, vertical dashed lines).

The potential flexibility options – namely electricity-based domestic hot water (DHW) boilers and space heating – are identified in the consumption profiles as scheduled and seasonal loads, respectively. The presence of peaks that occur at the same time-of-day throughout the year are checked against the (fixed) load-scheduling hours of the EVU (i.e., switch-on and switch off times of particular loads using ripple control) to confirm their source; we have used these to identify the peaks corresponding to the domestic hot water boiler charging (direct electric or heat pump) that occur during nighttime hours (black arrows indicate the more prominent of those peaks). In addition, the seasonal changes in electric consumption, attributed to the combination of heating demand (throughout the day), lighting demand (starting at sundown), and lifestyle changes (in the evening hours) are captured in the consumption profile range for the full year (top row). Using this information, we develop methods to decompose the profile based on separate load contributions – namely DHW and heating demand – and estimate their potential as a source of flexibility within the network. The methodology of profile decomposition, and the corresponding results, will be discussed in more detail in the full report.



10.2.5 Benefit of load shifting with increased PV production (“best case scenario”)

The benefits of daily load-shifting are analyzed for scaled PV production. First the PV production profile of the network area of EVU A is scaled until a certain annual production (as percentage of annual consumption) is reached. At each of these levels, the effective PV production-to-consumption is computed for the base (no load-shifting) case. Then, load-shifting windows of various duration between 3 and 24 hours are applied over the course of the year, and the maximum benefit of this load shift is computed. Consumption profiles remain the same for all PV expansion levels, and no losses associated with load-shifting or transmission are considered.

In Figure 48, effective PV consumption (y-axis) is shown for each PV production level (x-axis) considering several load-shifting windows indicated in the legend (colored bar sections). Within these load-shifting windows, the loads are assumed ideally shiftable, representing the best-case scenario for the given consumption/production profiles.

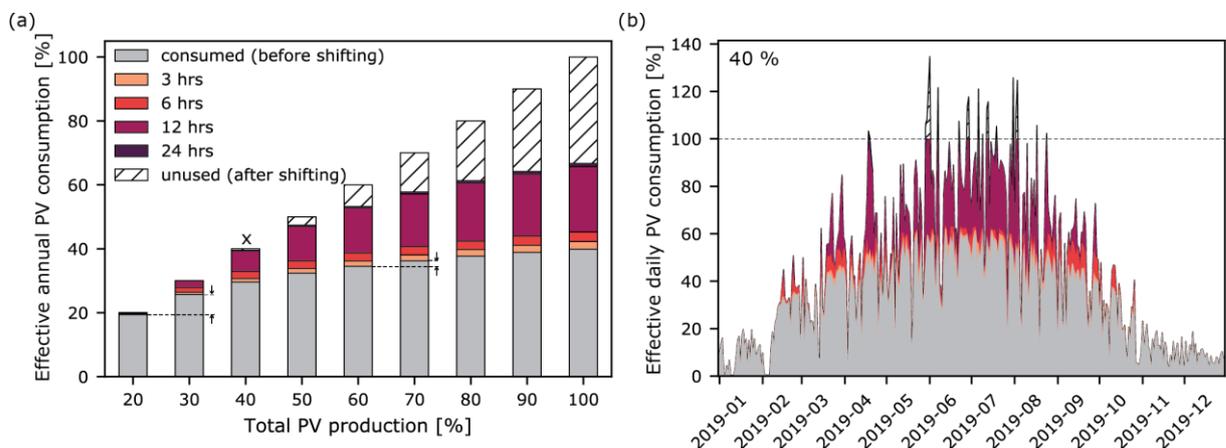


Figure 48: (a) Effective PV consumption (y-axis) at each PV production level (x-axis) without load shifting (grey) and load-shifting windows of specified duration (colors). Remaining unused PV production (hatched), following perfect load-shifting (up to 24 hours) is also shown. (b) Daily effective PV production- to-consumption throughout the year for the 40% annual production case. Minor differences are observed between 12 and 24-hr load shifting windows since the diurnal mismatch is already addressed.

We make the following observations:

- Additional increase in effective PV production will diminish gradually from 6-1% with each additional 10% increase in total PV production; in the figure, we highlight the effective increase for the 20→30% and 60→70% transitions for the base case scenario (approx. 6 and 2%, respectively).
- Additional increase in effective PV production is improved with load-shifting.
- With ideal daily shifting, approx. 40% of consumption could be met with PV production (with negligible outflows during weekends in summer shown in b).
- 24-hr shifting window brings negligible benefit over 12-hr because the day-night imbalance is already addressed.



10.2.6 Preliminary concluding remarks

From the analyses conducted thus far, the following steps are proposed for identifying flexibility needs, critical times, and available options for making the network more flexible:

- Calculation of the flexibility need indicators for the potential EVU network area to determine the main causes of inflexibility - temporal, spatial, or a combination of both.
- If the mismatch in production and consumption is mainly spatial, following analysis is localized to the relevant regions (transformer station, distribution cabin, end-consumer).
- If the mismatch in production and consumption (for the network area as a whole) is mainly temporal, following analysis is generalized to the entire network.
- Observations of when these discrepancies arise (season → weekday-workday → time of day).
- Identification of flexibility options based on scheduled and seasonal patterns (i.e., domestic hot water boiler charging and electricity-based space heating demand, respectively).
- Estimation of volumes from these contributions through profile decomposition (peak separation, seasonal contributions).
- Best case scenario with increased PV production for maximized self-consumption (what time shifts and volumes are necessary to address them).

The unique characteristics of each DSO – in terms of their consumption patterns, existing flexibility potential, and network topology – determine their readiness to adapt to current (and planned) PV expansion levels. Developing approaches to analyze their characteristics using minimal existing information (prior to extensive digitalization) is essential to ensuring the implementation of effective monitoring and control strategies that address their individual needs. The methodology developed here is a step in that direction.

The full report will include an expanded finalized analysis of the presented methodologies that will also be applied to the Swiss national grid profiles.

The necessary demand-response (DR) schema at the end-consumer level activating flexibility to respond to daily variations (which were investigated here for the network area) is addressed in Section 10.3. In addition, technical limitations of transmitting surplus PV production (i.e., critical times due to over-voltage) are addressed in Section 10.4.



10.3 Summary: Minimal information for collaborative consumption control

Authors: Hanmin Cai, Philipp Heer, Empa

10.3.1 Introduction

Recently, there has been a growing penetration of distributed renewable energy sources (RES) in the distribution system due to favorable financial policies, growing climate awareness, and promising economic assessments. While this supports the transition towards a more sustainable energy system, system operators cannot continue to operate their assets as in the past, when distribution grids were treated as pure sinks of energy. Power quality issues such as reverse power flows or increased voltage levels are possible due to rooftop PV production at periods of high solar irradiance. These issues may reach levels at which expensive network expansion becomes necessary based on traditional distribution system planning. Simultaneously, the digitalization of electrical loads in residential buildings has also gained momentum due to increasing business potential linked to digitalization. In the near future, the sizable storage in stationary batteries/electric vehicles (EVs) and high inertia of thermal systems will enable building energy management systems to modify how buildings traditionally interact with the power system without compromising end-user comfort. This report exploits local building energy flexibility as an alternative to the traditional prohibitively expensive network expansions. More specifically, the research questions addressed in this report are as follows:

1. What is the necessary information to be exchanged between buildings and upper-level controllers and how could we embed flexibility potentials that are dependent on time and initial condition in the exchanged information?
2. To what extent can we design an optimal controller to reduce equivalent carbon emission while respecting operating limits and users' preferences and how will this optimal controller further extract flexibility metrics to send them to the next super-ordinate cell for collaborative energy scheduling?

This rest of this summary briefly outlines our approaches to answer these two main questions and is structured as follows: Section 10.3.2 discusses proposed flexibility metrics that address research question 1. Section 10.3.3 outlines the methodology that supports the building energy management system design, the implementation at NEST, and a systematic evaluation of experimental results, addressing research question 2. Section 10.3.4 summarizes preliminary results and outlines the anticipated results in the full report and outlook for future investigation. In addition to the research and development tasks, Section 10.3.5 outlines the exchange of information and knowledge transfer between project partners.

10.3.2 Proposed flexibility metrics

There are two flexibility metrics investigated, namely, flexibility envelope and four-value time series. The flexibility envelope provides extensive details of flexible power levels and available duration. The four-value time series concept presents a simplified alternative of communicating flexibility that has a low communication cost.

10.3.2.1. Flexibility envelope

The flexibility envelope concept originally proposed in [20] is adapted to describe the time and initial condition dependent flexibility while taking into account anticipated boundary conditions, such as weather and user preference. The flexibility envelope extensively captures the thermal inertia of the building and the domestic hot water tank, the bi-directional charging of a stationary electric battery and electric vehicles, and curtailable PV. It is envisioned that buildings self-report this flexibility to the local distribution system operator (DSO) for network level dispatch (see Section 10.4). For example, with the flexibility envelope, a DSO also knows the duration that heat pumps (HPs) can be deferred without



compromising end-users' comfort. Currently, we extract flexibility for the following controllable devices: EV, HPs for space heating and domestic hot water, PV, and stationary battery.

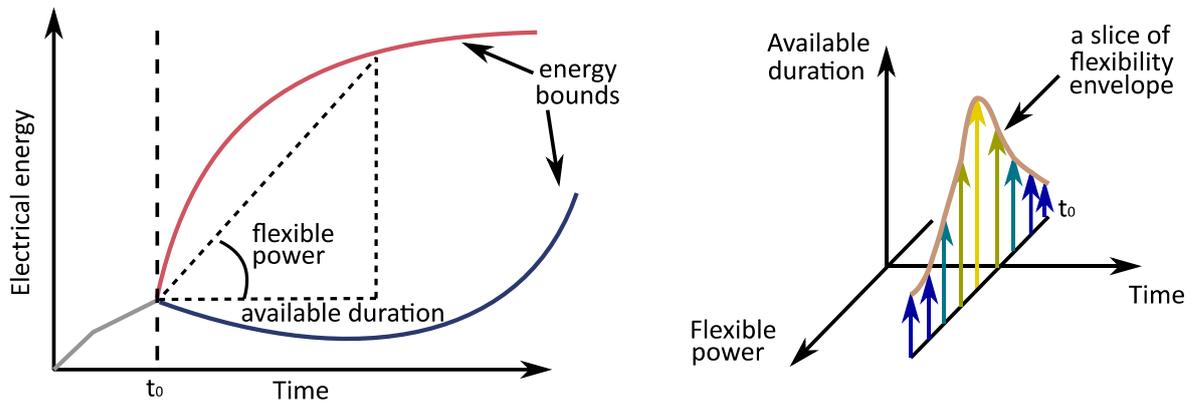


Figure 49: Illustration of flexibility envelope adapted from [20] where t_0 denotes next time step. The left figure depicts upper and lower energy bounds that allow derivation of flexible power and corresponding available duration. The right figure maps the time-varying power and duration into 3-dimensional space and the resulting red curve represents a slice of flexibility envelope at t_0 .

The interactions between buildings and DSO are further illustrated in Figure 50.

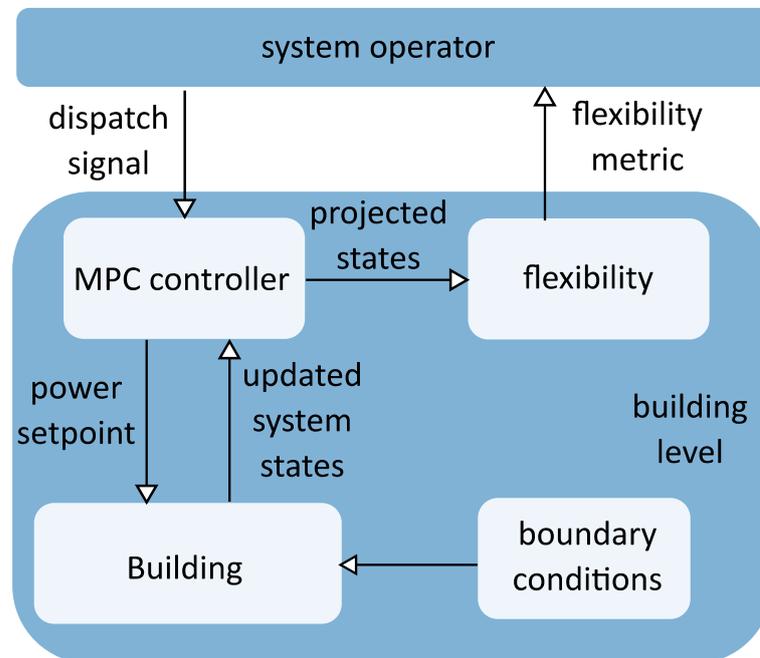


Figure 50: Structure of overall experimental implementation. Interactions with system operator are not implemented and are depicted to illustrate the integration into system management.



10.3.2.2. Four-value time series

In addition to flexibility envelope described above, four-value time series were proposed by aliunid. Please refer to Section 3.3 for more information.

10.3.3 Methodology

Research question 2 introduced in Section 10.3.1 is investigated by the following actionable steps:

1. Propose a system configuration that would emulate a typical Swiss household and is as representative for field deployment as possible, depending on available facilities at NEST.
2. Analyze historical data collected at NEST and develop models for the devices.
3. Develop proof-of-concept controller in simulation with the developed models of the proposed system configuration.
4. Design and plan experiments that support systematic verification of the deployment.
5. Execute optimized experiments and design end-user satisfaction survey to collect end-users' feedback.
6. Verify experimental results step by step and validate controller design according to expected outcomes and criteria specified in the experiment plan.

The Empa demonstrator NEST [21] is a highly modular facility that allows researchers to combine different equipment to emulate a desired experimental setup in a district. We leveraged the infrastructure at NEST to emulate a fully equipped single-family building. Specifically, we designed an emulator that includes a rooftop PV installation, a space heating HP, a domestic hot water HP, a stationary electrical battery and a bi-directional charging station for an EV. Although there was no bi-directional charging EV present during the project, we modelled EV driving patterns and changing schemes to emulate the respective loads. This set of equipment represents the promising flexibility sources of future Swiss residential buildings. Notably, the EV and the HPs have been identified as promising flexible resources because energy consumption in the building and transport sectors constitutes 65% of total energy consumption in Switzerland [22]. Besides, HPs have been widely adopted since 2000 [23] because of their low environmental impacts and operating costs. During the same period, the amount of registered EVs has increased by a factor of 38 [24]. In the controller design, an aggregated power profile of uncontrolled loads is also included in the optimization. Hence, the remainder of the summary assumes these components in the system configuration.

10.3.3.1. Modelling and proof-of-concept controller

We first analyzed the historical data collected at NEST. Notably, the thermal characteristics of building envelope and domestic hot water heating are non-linear. In practice, a first principle model would require a large amount of building construction details. We acknowledged that first principle modelling presents a barrier to scale up the adoption of the model predictive control strategy. Hence, we built upon recent work of data-driven modelling tools from Empa. We applied a data-driven approach to extract the thermal models for space heating and the domestic hot water tank, which are then used in the predictive controller development and flexibility envelope analysis. We exploited the rich data collection throughout the year to build a data-driven model that can capture sufficient dynamics of an underlying system, focusing on using as little physical knowledge as possible.



Table 16: List of components included in the emulator of a residential building.

Name	Location	Number of Units
PV	NEST	1
Battery	NEST	1
SH	NEST	2 bedrooms: 272, 274 and 1 living room: 273
DHW	NEST	1
Fixed load	NEST	1
EV	Computer script	1

We also acknowledge the situation that temperature sensors are often missing in domestic hot water tanks. Without such information, a controller will not be able to exploit the energetic flexibility within the temperature variation of the water tank. We attempted to address this problem by using a recurrent neural network based deep learning method. We investigated the possibility of recovering the energy status of the water tank with only a memory of past power inputs into the tank and time index. The RMSE of estimation can be as low as 10% depending on user water draw pattern. Before rolling out the control strategy at NEST, simulations were performed to validate the strategy as a proof-of-concept.

10.3.3.2. Implementation and experiment design

The demand response (DR) strategy has been implemented in NEST, with the experimental results being systematically quantified and verified. The components that we used to emulate a residential building are listed in Table 16. This work package deals with increasingly complex cyber-physical systems with interactions among electrical appliances, controllers, communication and data storage infrastructure, as well as external stakeholders. To incorporate all of these aspects a holistic and systematic approach is crucial in our development. We leveraged international experience on smart energy system development and testing, by choosing an approach based on the methodology described in [21]. We have designed a comprehensive description of test cases with a list of experiment plans to verify and assess the results systematically. In addition, we designed intermediate experiments to assess the sufficiency of models. One such intermediate experiment is presented in Section 10.3.4 with preliminary results. The full plan will be detailed in the full report.

10.3.3.3. Close the loop: user feedback

One critical boundary condition of the controller is the end-users' comfort. In the project, we have adopted users' preference metrics, such as indoor temperature comfort zone, commonly suggested in the existing literature and past projects [25]. However, to fully benefit from the NEST as a living lab, a user feedback survey was created for the residents to fill out regularly. The questions include the assessment of thermal comfort for both domestic hot water and space heating. With this survey, we also validate our controller and experiment design in practice. End users were not informed about the purpose of running these experiments, and we tried to minimize information revealed on the survey to avoid subconscious bias. Their qualitative ratings were collected and will be converted to quantitative metrics for further analysis and interpretation. Although the number of residents is currently limited, we expect the methodology to be generic enough to be used in the future large-scale roll-out of the technology developed in this work package. Outcomes of such a survey can be 1) verification of whether the current control strategy is sufficient; 2) and revealing the potential of relaxing the existing thermal comfort zone. Positive feedback would validate the controller development, and an extended thermal comfort zone can be used to increase flexibility potential.



10.3.4 Preliminary results and outlook

Intermediate experiments were designed to verify the accuracy of modelling in implementation. Preliminary results of these experiments are partly presented below. The final report will include a carbon emission-oriented building energy management that reacts to electricity carbon footprint profile and minimize total equivalent carbon emission. Experimental results have indicated that flexibility in occupied residential buildings can be utilized to serve customized objectives while respecting end-users' preference and comfort, as shown in Figure 51 and Figure 52.

The Space Heating (SH) demands in room 272 and 274 are minimized to save energy and reduce peak power while respecting customized comfort zone shown as dashed lines in Figure 51. For room 273, there was a large temperature drop on 2020-10-19 due to long window opening. The large glass wall in room 273 is also the reason for temperature overshooting during the day. In the experiment, only heating demand is controlled. Thus, the overshoot outside upper temperature limit cannot be avoided because heating was not activated. Overall, this validates the accuracy of developed SH model that captures the room temperature dynamics.

The Domestic Hot Water (DHW) heating demand was also actively controlled with a comfort zone of [45 °C, 60 °C]. In addition, the temperature is boosted to [59 °C, 60 °C] once a week to avoid Legionella contamination. The results on average tank temperature and corresponding power inputs are summarized in Figure 4. It shows developed controller sufficiently respect the specified temperature range and can guarantee a sufficient amount of hot water. Both Figure 51 and Figure 52 validate the accuracy of modelling and control strategy.

In addition, the self-sufficiency for the 5-day experiment is found to be around 17% due to low PV production. Apart the energy scheduling shown in Figure 51 and Figure 52, the controller also actively communicates the flexibility envelope and four-value time series to the next super-ordinate cell. A more detailed analysis on the results will be summarized in the final report.

10.3.5 Knowledge and technology transfer

The current report is connected to Sections 10.2 and 10.4 as follows.

- Section 10.2 reveals that reverse power flow occurred at parts of the network and emphasizes the need for load shifting to maximize PV self-consumption. This necessitates this work package as the controller developed here is capable of shifting loads to maximize self-consumption and provide local flexibility resources to address local issues.
- While this report investigates energy scheduling and flexibility potential at the building level, the dispatch strategy at the district level is addressed by Section 10.4, which can assess network health status with machine learning techniques and send dispatch signals to the individual building to alleviate network constraint violation.

During the project, we designed a setup to sufficiently represent that in the field. The energy management system is developed by Empa. Validated controller configurations were handed over to aliunid for field deployment. This has been realized after promising results have been obtained and verified by the research team at Empa, which also supported aliunid to benchmark the controller implemented in NEST.

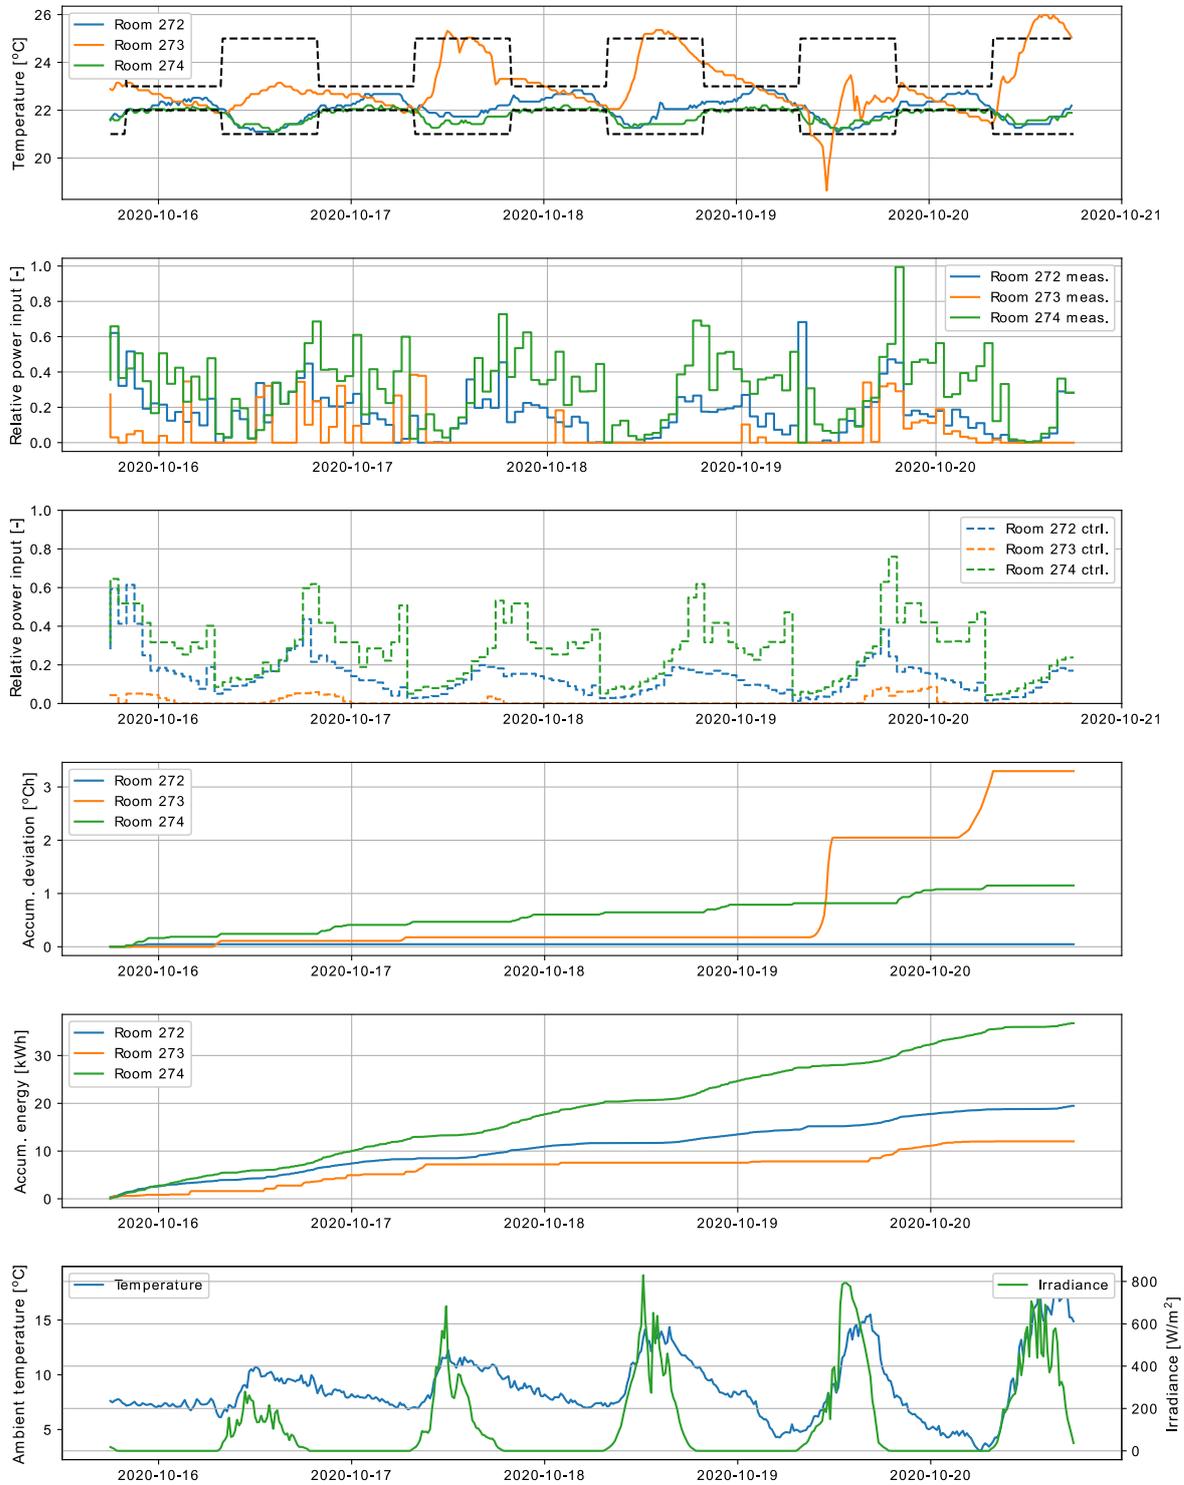


Figure 51: Experimental results with space heating control. The dashed black line indicates the thermal comfort zone.

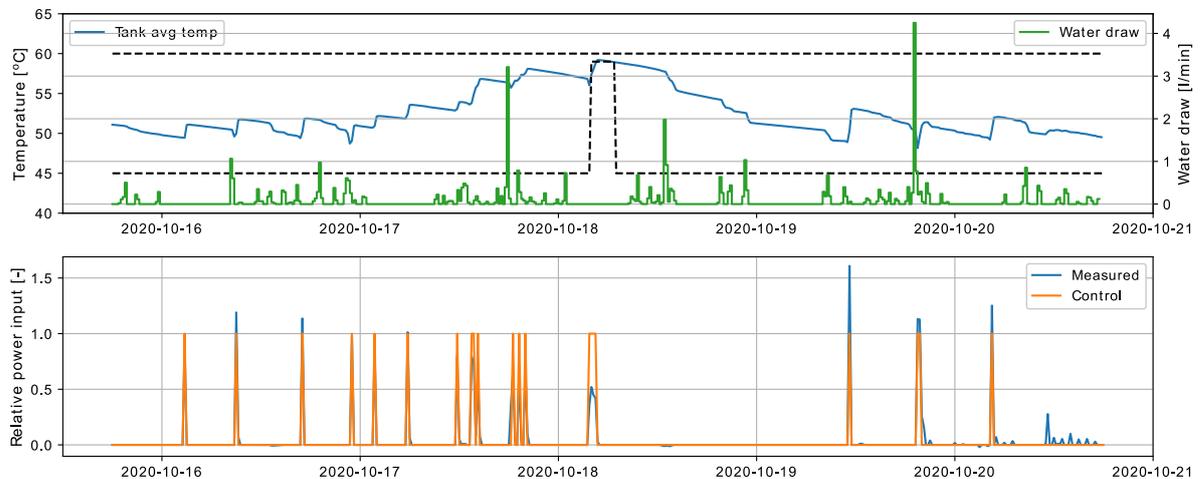


Figure 52: Experimental results with domestic hot water tank control. The dashed black line indicates the thermal comfort zone.

10.3.6 Take-home messages

- Model-based predictive controllers had been developed based on historical data with a minimal amount of physical knowledge. Preliminary results show these models have sufficient accuracy to optimize energy schedules while respecting end-users' preferences and comfort levels. Accumulated temperature deviation can be kept below 1 Kh in the 5-day experiment.
- HPs for the domestic hot water tank can be switched off for an extended duration. The observed flexibility from the hot water tank is larger than space heating. The reasons include the well-insulation building envelope and large tank size.
- Stationary batteries can maximize PV self-consumption and provide substantial bidirectional flexibility, which is significantly higher than those observed in the domestic hot water tank and space heating.
- When an EV capable of bidirectional charging is present and connected to the charging station at home, it can contribute to energy arbitrage and reduction of equivalent emission. Besides, it also complements the stationary battery and provides substantial bidirectional flexibility. Although an EV has a higher energy capacity than the battery, its energy and power flexibility are highly dependent on its usage pattern.
- PV curtailment can provide flexibility, although its level is low at the time of the experiment.



10.4 Summary: Safe flexibility utilization in active distribution grids

Author: Stavros Karagiannopoulos (ETH Zürich)

10.4.1 Background and project overview

The Power Systems Laboratory (PSL) of ETH Zurich conducts research in the area of electrical and integrated energy systems including their planning, design, and operation stages. In collaboration with aliunid, PSL has investigated some aspects of control and optimization of active distribution grids. Within the scope of this collaboration, the following four tasks were investigated:

- Consideration of the power flow physics for safe grid operation: Calculation and illustration of the flexibility amounts at the household levels that can be utilized safely at different points of the grid. The derived shrunk flexibility bounds consider the physics of the power flows and the network constraints in terms of thermal loading and voltage magnitudes.
- Data-driven tools to estimate grid health status: Development of data-driven tools to estimate the status of the distribution networks (DN) without requiring a full two-way communication link among all units. Within the scope of this collaboration, PSL has designed classifiers that are trained to forecast the status of the grid in terms of power quality constraint satisfaction based on historical data and expected operating conditions. Thus, depending on the amount of measurement data that will be available, a classifier can estimate and predict future unsafe operating points. This will be used as a signal to modify controllable units to avoid power quality issues.
- Interactions between GB1 and GB2: In case of an unsafe estimated grid status, a signal should be sent from GB2 to GB1 at a lower hierarchical level to ask for modifications that will not result in constraint violations.
- Demonstration of the proposed methods and evaluation of the control schemes via case studies: The benefits of the developed methodologies were highlighted by considering realistic grid configurations and data for various DNs in Switzerland.

10.4.2 Power flow calculations and data-driven control scheme

In the developed data-driven control scheme, a training stage (*offline stage*) is used to derive a representative classifier¹² based on the unique characteristics of each network and the solutions of the load flow problem on historical data. Then, during the real-time (*online*) stage, no power flow calculations need to be run since the classifier can take over this functionality.

Offline Stage: Figure 53 illustrates the offline stage in a block structure. The input data block contains the properties which are used as input both for the model creation as well as for the real-time operation. Some examples that can be used as input data can include solar radiation data, loading data, temperature information, active/reactive power measurements at specific nodes, local time and date, etc. The second block is related to grid parameters and contains information concerning the topology of the network, the impedances of the lines, the installed capacity of the DERs and the nodes where information is provided in detail, i.e., participants in the aliunid concept. The third block refers to power flow calculations that are needed to calculate the power quality and guarantee safe grid operation even when flexibility is used. The power flow analysis describes the operating state of the entire power system, by calculating some quantities such as the complex voltage at each bus. Numerical methods that rely on successive approximations, or iterations, are used to solve the power flow. An iterative scheme, namely the Backward-Forward Sweep (BFS) power flow was implemented

¹² By the term classifier we refer to the mathematical function that maps input data to a category, in this case using machine learning.



for the scope of this work package. The result of this block refers to the voltages (magnitude and phase) of all network nodes.

To create the classifiers, different machine learning methods were compared. In the end, Support Vector Machines (SVM) was chosen. A SVM is a supervised learning algorithm that can be used for binary classification or regression. Three categories were chosen: the 'good' grid status without network violations, 'alerted' state (Medium class) where some constraints are approaching the minimum or maximum acceptable values and the 'critical' case, where one or many of the network constraints are violated. Intuitively, high PV injections increase local voltages triggering over voltages or thermal overloads, low load conditions trigger undervoltage issues and high load conditions trigger under voltages and thermal overloads.

Online Stage: During real-time operation, only new input data are needed so that the classifier can be evaluated, and the health status estimated. These will be provided by the results of GB1 (see Section 10.3) and on globally available information, e.g., regarding solar radiation data. Figure 54 illustrates this approach.

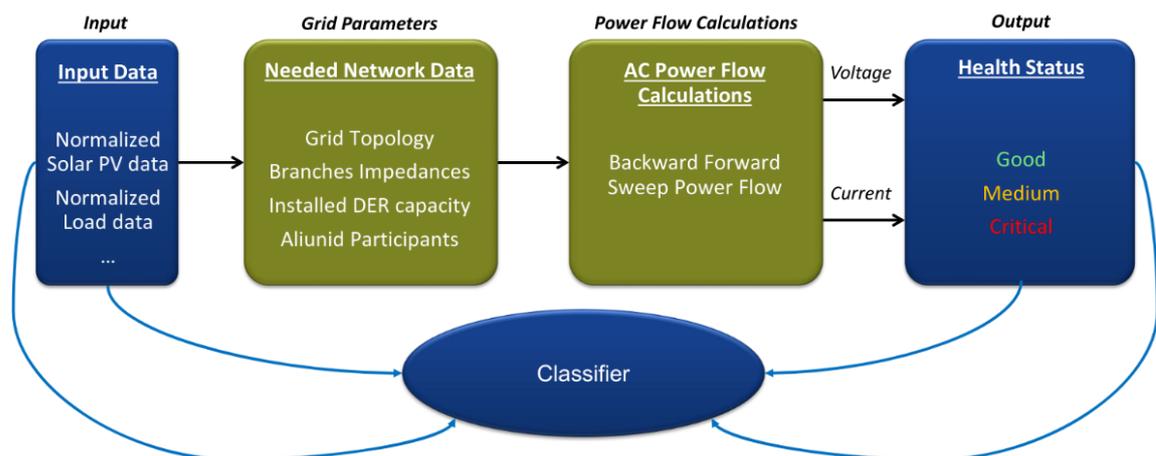


Figure 53: Proposed overall methodology for the offline training stage.

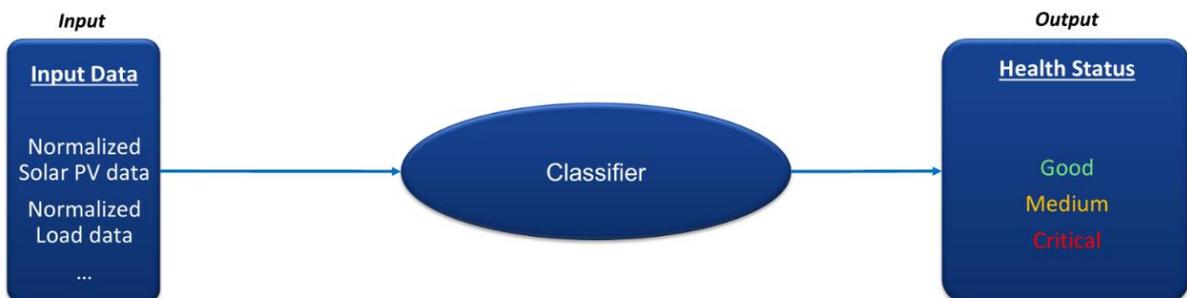


Figure 54: Methodology for the real-time, online-stage.



10.4.3 Interactions between GB1 and GB2

Based on the outcome of the classifier evaluation, a GB1-GB2 interaction scheme was developed. The working principle is as follows: When a problematic status is identified by the classifier (i.e., Medium or Critical), GB2 can send one common signal to all aliunid participants, such that they modify their scheduling optimization in the direction that will alleviate the power quality issue. This way is the simplest and most robust way to impose a certain behavior to aliunid's participants. However, it can be suboptimal, since also participants that do not create quality issues might need to modify their behavior.

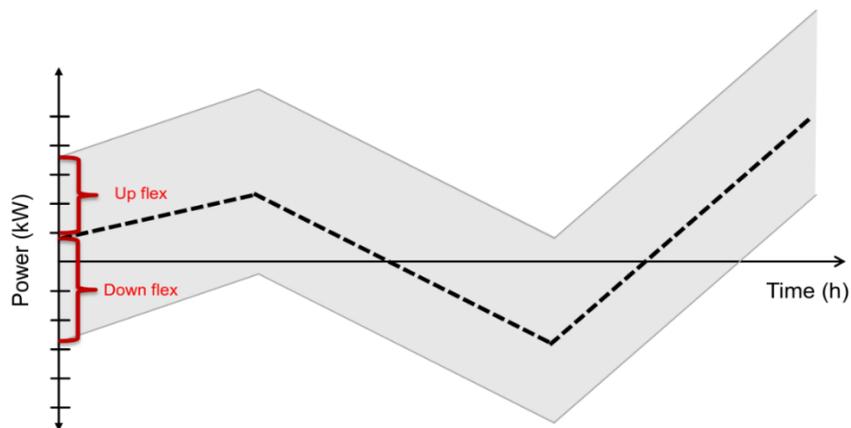


Figure 55: Baseline and flexibility bounds communicated to GB2.

Figure 55 illustrates the main principle of the directive signal proportional to the flexibilities offered by GB1 and communicated to GB2. The dotted black line corresponds to the baseline present and forecasted trajectory of one household-level GB1. The gray areas above and below the baseline illustrate the forecasted up and down flexibility of this user. GB2 will use the grid health estimation for the current time step to estimate and alleviate potential power quality issues. Thus, in case of problematic health status estimation, a directive signal can be sent to GB1. In case of overloading or overvoltage issues, the baseline needs to be reduced, by utilizing down flexibility. Similarly, in case of undervoltage issues, the baselines would need to increase requesting up-flexibility resources.

10.4.4 Case Studies

The presented approach was tested on a model of a typical European LV and MV grid. The input data were based on the solar radiation at the area of Switzerland and on the wind profiles based on the wind speed data from two areas in Germany. To assign different loading behavior at each node, a $\pm 20\%$ p.u. random deviation from the typical load profile was superimposed.

Furthermore, a real case referring to a grid taken from IBW was considered. Some of the units in this grid participate in aliunid's framework and provided detailed measurements. By subtracting the measured nodes from the total loading of the transformers, the loading of the remaining nodes could be calculated without measurements. Rooftop PV units were added at all nodes with loads to simulate future scenarios with high PV penetration.

Overall, the accuracy of the obtained classifier is 95.6% for the simulated grid and 95.2% for the IBW grid. However, since most of the cases considered lie in the normal operating condition without too much stress on the grid, i.e., at night periods, or when the intermittent generation is low, one can be interested in the accuracy levels for the alerted (medium), or critical cases. The case studies helped to understand the strengths and the weaknesses of the classifiers. Areas with poor accuracy levels indicate that the derived classifier model could not capture all classes satisfactorily. In such a case, possible mitigation actions would include more input data and understanding the reasons at the specific classes with low accuracy values.



10.5 Summary: Requirements customer interface

Author: Andreas Danuser, BFH

10.5.1 Einleitung

Mit dem Titel «SMv2 Spezifikation» hat aliunid AG im April 2020 ein Dokument bereitgestellt, in welchem die Anforderungen an ein intelligentes Messgerät (engl. smart meter) für die Zwecke einer atmenden Versorgung hergeleitet und beschrieben sind.

Zentraler Punkt der Forderungen von aliunid ist eine *standardisierte* und *diskriminierungsfrei nutzbare* Kundenschnittstelle (engl. CII – Customer Information Interface), über welche lokal (d.h. am Ort der Messung) Echtzeitdaten der erfassten el. Energieflüsse mit kurzer Latenz und einer zeitlichen Auflösung im Sekundenbereich bereitgestellt werden.

Diese Forderung der aliunid ist nicht neu. Tatsächlich wird eine solche offene Kundenschnittstelle in verschiedenen Dokumenten seitens des Bundesamtes für Energie (BfE) und des VSE gefordert und soll ab 1.1.2021 bei allen intelligenten Messgeräten verbindlich vorliegen.

Die Forderung liegt also klar vor. Aber die entsprechende Beschreibung ist zu wenig detailliert ausgeführt, was dazu führt, dass unterschiedliche Umsetzungsvarianten im Felde ausgerollt werden, welche zudem auch funktionale Defizite aufweisen. Aus diesem Grunde spricht aliunid oft von einem sog. «Mess-Tschungel».

aliunid unternimmt unter dem Titel «SmartMeter Kundenschnittstelle» einen Vorstoss, die Umsetzung der Forderung nach einer solchen Kundenschnittstelle detaillierter vorzuschlagen.

Dabei geht dieses Dokument wie folgt vor:

- Beschreibung und Begründung der Anforderungen an diese Kundenschnittstelle
- Vorstellung und Bewertung der heute vorhandenen ansatzweisen Umsetzung der Kundenschnittstelle
- Konkreter Vorschlag einer idealen Umsetzung
- Vorschlag einer Umsetzung auf Basis der heute bekannten Ansätze

Das vorliegende Dokument ist lediglich eine Zusammenfassung des umfassenderen Dokumentes «SmartMeter Kundenschnittstelle». Die detaillierteren Beschreibungen sind im vollumfänglichen Bericht dargestellt, auf welchen hier summarisch hingewiesen wird.

10.5.2 Systemumgebung

Die Dokumente von verschiedenen Quellen, insbesondere auch das Dokument [26] zeigen die Grundstruktur eines «Intelligenten Messsystems, was hier zusammengefasst ist:

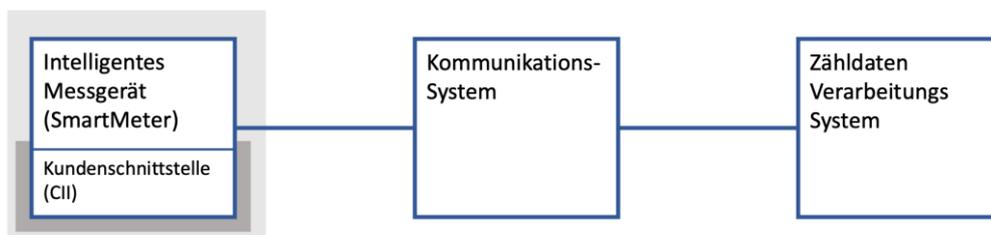


Figure 56: Intelligentes Messsystem.



Aktuell (d.h. im 2020) sind verschiedene solcher Systeme bereits installiert bzw. im Rollout. Diese Systeme sind proprietär, was zur Folge hat, dass Systeme verschiedener Hersteller nicht kombiniert werden können.

Der Fokus des vorliegenden Berichtes liegt jedoch auf einem kleinen, aber wichtigen Teilaspekt, nämlich der sog. Kundenschnittstelle des Intelligenten Messgerätes (CII des SmartMeters) – siehe grau hinterlegte Teile.

Diese offene, dokumentierte und allgemeinen Standards folgende Kundenschnittstelle wird in den Dokumenten gefordert und soll in Echtzeit (im Sekundenbereich) und diskriminierungsfrei Verbrauchs- und Produktionsdaten für beliebige Anwendungen (z.B. Energiedienstleistungen oder Gebäudeautomation) bereitstellen.

Allerdings wird diese Schnittstelle in den vorhandenen Dokumenten technisch zu wenig detailliert beschrieben, so dass wiederum unterschiedliche Standards angewendet werden – eine einheitliche Definition dieser Schnittstelle fehlt, was dazu führt, dass eine flächendeckende Anwendung unter Einbezug der Messgeräte verschiedener Hersteller unmöglich oder nur mit grossem Aufwand möglich ist.

10.5.3 Anforderungen

Auf Basis systematischer Überlegungen und den Erfahrungen aus dem Feldtest der aliunid werden nun die Anforderungen im Überblick dargestellt:

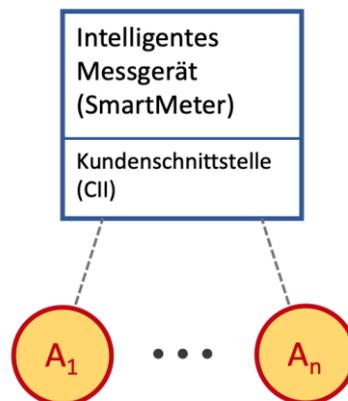


Figure 57: Kundenschnittstelle CII.

1. Die Schnittstelle (CII) soll Hersteller-übergreifend gleich sein.
2. Die Schnittstelle soll auf allgemein bekannten, dokumentierten und häufig verwendeten Standards beruhen.
3. Die angewendeten Standards sollen zudem über mind. 20 Jahre unterstützt werden können.
4. Die Schnittstelle soll direkt lokal vorliegen - also keine Datenlieferung über ein nachgeschaltetes System.
5. Die Schnittstelle muss ohne Umbau des Gerätes (mechanisch und/oder IT-Konfiguration) nutzbar sein.
6. Mehrere Anwendungen (also ≥ 1) sollen die Schnittstelle gleichzeitig nutzen können.
7. Diese gleichzeitige Nutzung soll zudem ohne Interferenzen (gegenseitige Blockade) erfolgen können.



8. Unterschiedliche Anwendungen haben unterschiedliche Bedürfnisse hinsichtlich dieser Daten: «wann» (Zeitpunkt(e), wie oft) und «was» (welche Daten) werden benötigt. Die Schnittstelle muss deshalb gleichzeitig unterschiedliche Bedürfnisse hinsichtlich Dateninhalte und Auslesezeitpunkt unterstützen.
9. Die Auslesung soll effizient erfolgen können.
10. Die ausgelesenen Daten sollen entweder direkt der Anwendung oder indirekt mittels eines Gateways zugeführt werden. (Gateways werden benötigt, falls am Ort der Messung (z.B. Keller) keine geeignete Datenverbindung vorliegt.)
11. Da am Ort der Messung oft keine frei zugängliche Energiequelle vorliegt, soll die Schnittstelle über eine Energieversorgung für zumindest ein Gateway verfügen.
12. Die Schnittstelle muss rückwirkungsfrei sein und eine negative Beeinflussung (z.B. Umkonfiguration des SmartMeters) verunmöglichen.
13. Die Schnittstelle muss die Daten optional verschlüsselt bereitstellen
14. Die Schnittstelle muss das Messgerät kryptologisch eindeutig authentisieren.

Eine Herleitung mit Begründung dieser Anforderungen wird im umfassenderen Dokument «SmartMeter Kundenschnittstelle» geliefert.

10.5.4 Bestehende Ansätze

Messdaten in Echtzeit werden schon heute bereitgestellt. Bereits die alten Ferraris-Zähler informierten den Kunden mit Hilfe der Drehscheiben über den aktuellen Energiefluss.

Seit der Einführung elektronischer Zähler und aktuell der Intelligenten Messgeräte wurden diese Informations-Schnittstellen stets verbessert. Nachfolgend sind die wichtigsten Technologien zusammengefasst (auch hier wird für eine tiefere Begründung auf das umfassendere Dokument verwiesen):

LED-Anzeige

Es wird eine gut sichtbare LED genutzt, welche nach Durchfluss einer bestimmten Energiemenge (typischerweise 1Wh) kurz blinkt. Die Zeit zwischen zwei Blink-Ereignissen dient also dazu, die (durchschnittliche) Wirkleistung zum aktuellen Zeitpunkt zu ermitteln.



Figure 58: LED, welche mittels Blink-Impulse Wirkleistungen anzeigt.

Vorteile:

- Einfache Messtechnik
- Sehr weit verbreitet und Hersteller-übergreifend nutzbar
- Reicht eigentlich in den meisten Fällen aus, um die aktuellen Leistungswerte zu ermitteln



Nachteile:

- Es wird lediglich summarisch (über alle 3 Phasen) die Wirkleistung gemessen
- Demnach: keine Messung von Spannung, Strom, Frequenz, $\cos\phi$, Blind- und Scheinleistung
- Eher ungenau und nicht zeitaktuell bei geringen Energieflüssen
- Wird generell nicht als Lösung für die Kundenschnittstelle betrachtet

Optische Schnittstelle nach IEC62056 [27]

Mit Hilfe von zwei LED, welche jedoch im Infrarot-Bereich arbeiten, kann eine serielle bidirektionale Datenverbindung zwischen dem Messgerät und einem Auslesegerät (Applikation) hergestellt werden. Der entsprechende IEC-Standard legt DLMS/COSEM als Protokoll sowie OBIS als Datenformatierung fest.



Figure 59: Mess-Kopf (links), welcher eine optische Schnittstelle nach IEC62056 (rechts) anbindet. Zudem gut sichtbar (rot umrandet): 2x LED zur Anzeige von Wirk- und Blindleistung.

Vorteile:

- Umfassende Datensätze für alle Phasen verfügbar, inklusive Spannung, Strom, Frequenz $\cos\phi$, Wirk-, Blind- und Scheinleistung
- Echtzeitdaten verfügbar, wenngleich nicht im Sekundenbereich
- Schnittstelle gut nutzbar
- Information kann verschlüsselt werden

Nachteile:

- DLMS/COSEM definiert 4+1 Variante → aufwändiges Protokoll, oft schwierig zu beherrschen
- Interferenzen bei Auslesung durch mehrere (\geq Anwendungen)
- Nicht rückwirkungsfrei
- Wird von vielen Versorgern nicht als Kundenschnittstelle betrachtet

M-Bus Push

Diverse SmartMeter-Systeme basieren auf einem Konzept, wonach der SmartMeter quasi als Messzentrale agiert und sich über M-Bus z.B. mit Gas- und Wasserzählern verbindet und die entsprechenden Verbrauchswerte ausliest und an ein Zählerten-Verarbeitungssystem (siehe Figure 56) übermittelt. Der SmartMeter, als M-Bus Master agierend, kann über den M-Bus auch lokal Messdaten an angeschlossene Geräte / Anwendungen übermitteln, was mit „M-Bus Push“ bezeichnet wird.

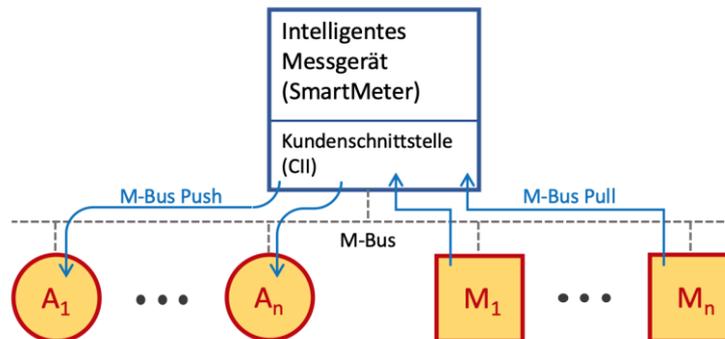


Figure 60: Kundenschnittstelle umgesetzt mit M-Bus.

Vorteile:

- Umfassende Datensätze für alle Phasen verfügbar, inklusive Spannung, Strom, Frequenz cosPHY, Wirk-, Blind- und Scheinleistung
- Information in Echtzeit (Auflösung je Sekunde)
- M-Bus ist ein standardisiertes Protokoll
- Speisung lokal vorhanden
- Lösung von Versorgungsunternehmen unterstützt
- Könnte für ≥ 1 Applikationen nutzbar gemacht werden, was aber neue SW für SmartMeter braucht

Nachteile:

- Benötigt in der Regel eine Umkonfiguration der Messgeräte vor Ort (Zeitaufwand ca. 1h)
- Wenig Dokumentation vorhanden
- Hersteller-abhängige Umsetzung
- Aktuell nur von einer Applikation nutzbar
-

DSMR/P1

Die Herausforderung „Kundenschnittstelle“ wurde in Holland bereits sehr gut adressiert mittels einer eigenen Spezifikation «Dutch Smart-Meter Reading» (siehe [28])

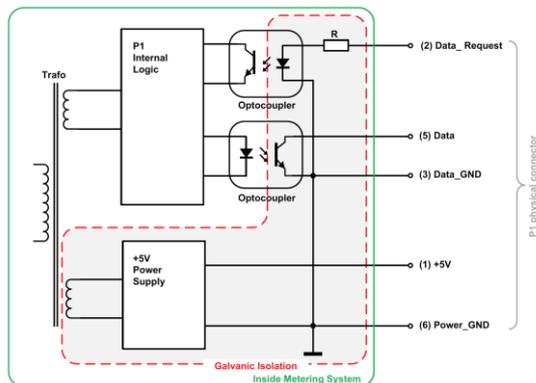


Figure 61: Kundenschnittstelle gemäss DSMR/P1 (Kopie aus Dokument [28]).

Vorteile:

- Sauberes, einfaches Konzept
- Gute Dokumentation
- Umfassende Datensätze für alle Phasen verfügbar, inklusive Spannung, Strom, Frequenz, cosPHY, Wirkleistung, Blind- und Scheinleistung
- Information in Echtzeit (Auflösung je Sekunde)
- Sauberes, einfaches Konzept
- Gute Dokumentation
- Speisung lokal vorhanden
- Lösung von Versorgungsunternehmen unterstützt

Nachteile:

- Aktuell nur von einer Applikation nutzbar

Ideale Umsetzung: CII «ideal»

Der Vorschlag für eine ideale Umsetzung des CII ist durch die oben genannten Anforderungen motiviert, insbesondere durch folgende Forderungen:

- ≥ 1 Applikation soll gleichzeitig auf die Informationen des SmartMeters zugreifen können.
- Die Zeitpunkt(e) der Auslesung sowie die Zusammensetzung der benötigten Daten ist Sache der Applikation.
- Die Lösung soll auf offenen und allgemein in der IT genutzten Standards basieren.

Untenstehende Figur links zeigt den Vorschlag der idealen Umsetzung.

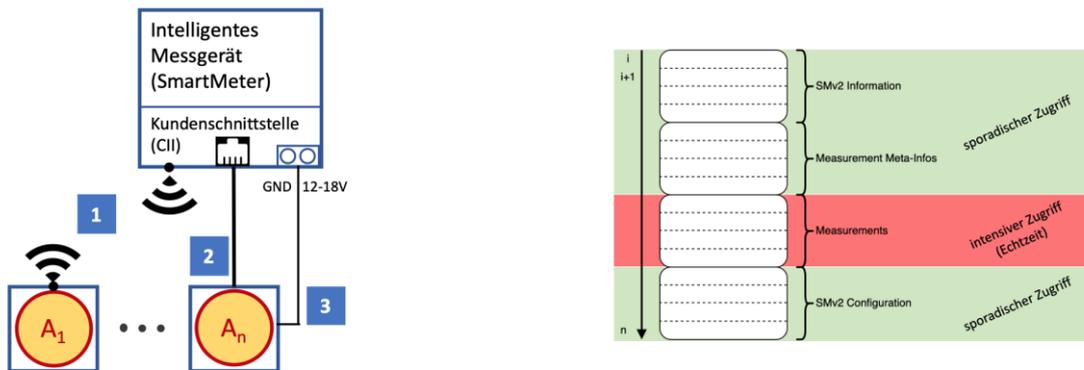


Figure 62: SmartMeter CII «ideal» auf Basis Modbus/TCP über WIFI und Ethernet.

Der SmartMeter verfügt hierbei sowohl ein über eine WIFI- wie auch über eine Ethernet-Schnittstelle und agiert als Server in einem privat adressierten IP-Netzwerk. Zudem liefert der SmartMeter für den Betrieb von nachgeschalteten Geräten eine (begrenzte) Energie von typ. max. 2Watt, was einer Energiemenge von knapp 1,5 kWh pro Monat entspricht.

Applikationen (z.B. A_1), welche als Client(s) arbeiten, können sich über WIFI mit dem SmartMeter verbinden (Weg 1); dies kann gleichzeitig durch mehrere Applikationen erfolgen.

Applikationen können sich alternativ über Ethernet mit dem SmartMeter verbinden (Weg 2). Die Ethernet-Schnittstelle kann ebenfalls durch verschiedene Applikationen gleichzeitig genutzt werden, was dann aber zusätzlich einen zwischengeschalteten Ethernet-Hub bedingt (zwischen Ethernet-Stecker und Applikationen).

In geringem Umfang wäre auch die Versorgung von Anwendungen mit Energie möglich (Weg 3).

Als Protokoll wird das Client/Server Protokoll Modbus/TCP verwendet, welches über beliebige IP-Netzwerke und somit über WIFI und Ethernet angewendet werden kann.

Das Modbus-Protokoll definiert, dass Informationen (also z.B. der aktuelle Wert der Wirkleistung) in sog. Registern abgelegt sind. Diese Register können - ähnlich wie Webseiten beim bekannten Web-Protokoll HTTP/HTTPS – abgefragt werden. Das Modbus-Protokoll arbeitet also nach dem Request→Response Verfahren.

Die Applikation (Client) stellt eine Anfrage (Request) an den SmartMeter (Server), welcher eine entsprechende Antwort (Response) gibt. Beispiel: Anfrage = «lese Register 1000, welches die gesamte Wirkleistung beinhaltet» und Antwort = «10,7kW».

Das Protokoll ist sog. «stateless» (mehr Info darüber im umfassenderen Dokument), und der Zugriff auf die einzelnen Register / Informationsinhalte kann kontrolliert werden (→ Sicherheit!).

Die Registerbelegung, d.h. die Regelung wo welche Information steht, muss herstellerübergreifend standardisiert sein, was in Figur rechts dargestellt wird!

Die Nutzung von Modbus/TCP über WIFI und/oder Ethernet kombiniert mit einer einheitlichen Registerbelegung führt also zu einem sinnvollen, offenen und einfach zu definieren Standard.

In diesem Zusammenhang ist auch der folgende Hinweis wichtig: Für Geräte wie Wechselrichter, Batteriespeicher, etc. wurde ebenfalls auf Basis Modbus eine einheitliche Spezifikation namens «Sunspec» erstellt, welche die Anbindung solcher Geräte Hersteller-übergreifend regelt. Deshalb erscheint es dem Autor sinnvoll, ein analoges Vorgehen für SmartMeter vorzuschlagen.

Mit diesem Vorgehen kann also einfach und kostengünstig eine Hersteller-übergreifende Lösung angeboten werden, welche somit flächendeckend Anwendungen unterstützt und auch ideal in bestehende Systeme der Gebäudeautomation integriert werden können.

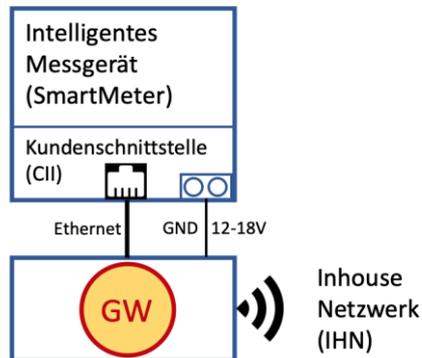


Figure 63: Anschluss an Inhouse Netzwerke (IHN).

In vielen Fällen sind an den Orten, wo SmartMeters zum Einsatz kommen (neben dem Fehlen von nutzbaren Energiequellen) keine geeigneten Netzwerke (WIFI oder Ethernet) verfügbar. Trotzdem sollen die SmartMeters in entsprechende Inhouse-Netze (IHN) eingebunden werden können, welche z.B. auf niederfrequenten Verfahren oder LoRa basieren, die lediglich Datenraten von 100 – 1000 bits/sec ermöglichen.

Mit Hilfe der vorgeschlagenen idealen Umsetzung kann dies mit einfachen und Kosten-günstigen Gateways realisiert werden, welche bezüglich Auslesehäufigkeit und Datenumfang den Applkationen und Netzwerken (IHN) angepasst sind, was in obenstehender Figure 63 gezeigt wird.

Umsetzung mit bestehender CII

Mit Hilfe eines kostengünstigen Gateway-Moduls können SmartMeter, welche das CII mit Hilfe M-Bus-Push oder DSMR/P1 umgesetzt haben, ebenfalls CII «ideal» unterstützen, was in untenstehender Figure 64 illustriert wird.

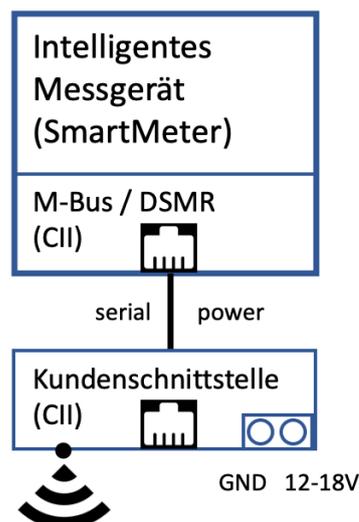


Figure 64: Integration bestehende SmartMeter an CII «ideal».

Das entsprechende Gateway Modul empfängt dauernd die vom SmartMeter über M-Bus oder DSMR gesendeten Zählerinformationen und stellt die Information so zusammen, dass die gewünschten Informationen zur gewünschten Zeit mittels Modbus/TCP über WIFI oder Ethernet abgerufen werden können.



Es ist zu untersuchen, in welcher Weise auch SmartMeter, welche nicht auf M-Bus Push oder DSMR basieren, mit so einem Gateway nachgerüstet werden können. Es scheint jedoch wahrscheinlich, dass dies möglich ist.

Schlussfolgerung

Unter dem Begriff CII «ideal» wird ein Vorschlag präsentiert, welcher die identifizierten Anforderungen an ein Customer Information Interface (CII) eines SmartMeters allesamt erfüllt.

Da dieser Vorschlag doch erheblich von den heute bereitgestellten SmartMeter CII – sofern vorhanden – abweicht, wird aufgezeigt, wie diese bestehenden Lösungen mit geringen Kosten auf den vorgeschlagenen neuen, einheitlichen und Hersteller-unabhängigen Standard migriert werden kann.

Für eine solche Migration eignen sich die beiden oben erwähnten, heute verwendeten Verfahren M-Bus-Push und DSMR/P1 am besten, wobei DSMR/P1 favorisiert würde.

Für Netzbetreiber, welche heute auf Basis von bestehenden SmartMeters eine Beschaffung planen, empfohlen wird deshalb als Alternative zu CII «ideal» den in Holland durchgesetzten Standard DSMR/P1 anzufordern.



10.6 Summary: Requirements IoT security

Author: Andreas Danuser, Tobias Rothen, BFH

10.6.1 Einleitung

In seinem Dokument mit Titel «Smart Grid Roadmap Schweiz» ([BfE 03/15]) postuliert das Bundesamt für Energie (BfE), dass intelligente Stromnetze, sog. Smart Grids, der «Weg in die Zukunft der Schweizer Elektrizitätsnetze» sind. Weiter definiert das BfE: «Als ein Smart Grid wird ein elektrisches System verstanden, das unter Einbezug von Mess- sowie meist digitaler Informations- und Kommunikationstechnologien den Austausch elektrischer Energie aus verschiedenartigen Quellen mit Konsumenten verschiedener Bedarfscharakteristika intelligent sicherstellt.»

aliunid teilt diese Sichtweise: Smart Grids werden Basis für die Umsetzung der vom Schweizer Stimmvolk beschlossenen Energiestrategie 2050 sein. Dabei sollen Smart Grids nicht nur auf Elektrizitätsnetze sondern auch auf Gas- und Wärmenetze angewendet werden (Sektoren).

aliunid baut eine «atmende Versorgung»: Die Produktion, die Speicherung und der Verbrauch werden analysiert und nach verschiedenen Kriterien optimiert beeinflusst, unter Einbezug aller Energieformen.

Hierzu entwickelt aliunid eine digitale IT-Plattform, «dEVU» genannt (digitales Energieversorgungsunternehmen), welche konsequent auf einem IoT-Ansatz basiert. Das heisst: Eine sehr grosse Anzahl Sensoren und Aktoren werden verteilt über das gesamte Versorgungsgebiet eingesetzt, um ein möglichst genaues Gesamtbild der Energieflüsse zu erhalten («digitaler Zwilling»), und um zielgerichtet diese Energieflüsse zu beeinflussen. Dies wird mit den Attributen «ubiquitous» (allgegenwärtig) und «pervasive» (durchdringend) umschrieben.

Damit wird die Energieversorgung von der korrekten Funktion dieser IT-Plattform abhängig: die Systeme müssen verfügbar sein, korrekt funktionieren und die erfassten und übertragenen Informationen (Messdaten und Befehle) müssen korrekt sein. Zudem muss sichergestellt werden, dass die Privatsphäre der Nutzer und Kunden trotz massiver Datenerfassung und Analyse nicht verletzt wird! Die IT-Plattform, welche auf einem IoT-Ansatz basiert, muss also «sicher» sein.

Das Dokument «IoT-Security» beschreibt einen Ansatz, wie ein solches System realisiert werden kann.

Dabei umfasst der Begriff Security nicht nur die Aspekte der Datenverschlüsselung, sondern vielmehr auch die Verfügbarkeit der Systeme, die Korrektheit und die Anonymisierung der Daten, wie im Abschnitt oben illustriert.

Das Dokument verfolgt zudem einen systemischen Ansatz: Die Sicherheit wird nicht primär mit Kryptologie, sondern vielmehr über eine geeignete Strukturierung der Daten und mittels der Architektur des Gesamtsystems erzielt, natürlich an vielen Orten ergänzt mit den Verfahren der Kryptologie.

Der Bericht ist wie folgt gegliedert:

Er beginnt mit einer Erläuterung, welche Aspekte das Thema IT-Sicherheit beinhaltet, gefolgt von einer Beschreibung Herausforderung «atmende Versorgung».

Danach wird aufgezeigt, dass erstens die Datenstrukturierung und zweitens die Systemarchitektur der IT-Plattform die Hauptpfeiler für die IT-Sicherheit sind (Abschnitte «4-Werte-Modell» und «SIOT»). Sodann wird aufgezeigt, wie die Systemverfügbarkeit, die Datensicherheit und der Datenschutz gewährleistet werden.

Im Abschnitt «Schlussfolgerung» wird Bilanz über die dargestellten Verfahren gezogen.

Das vorliegende Dokument ist lediglich eine Zusammenfassung des umfassenderen Dokumentes «IT-Security». Die detaillierteren Beschreibungen sind im vollumfänglichen Bericht dargestellt, auf welchen hier summarisch hingewiesen wird.



10.6.2 Was bedeutet IT-Sicherheit?

Untenstehende Figure 66 zeigt, wie ein Sender S eine Meldung M an einen oder auch mehrere Empfänger E sendet:



Figure 65: Grundmodell für die Aspekte der IT-Sicherheit.

Dabei stellen sich sechs wichtige Fragen, welche die sechs Aspekte der sog. IT-Sicherheit identifizieren:

1. «**Privacy**»: Ist sichergestellt, dass ausschliesslich berechtigte Empfänger die Meldung verstehen können?
2. «**Integrity**»: Ist sichergestellt, dass die Meldung unverändert übertragen wird?
3. «**Authenticity**»: Stammt die Meldung von einem zuverlässig bekannten Sender?
4. «**Anonymity**»: Ist sichergestellt, dass man nicht erkennt, von welchem Sender die Meldung stammt?
5. «**Authorization**»: Sind Sender S und Empfänger E berechtigt, die Meldung M zu senden resp. zu empfangen?
6. «**Availability**»: Sind Sender S und Empfänger E in der Lage, die Meldung S zu senden / zu empfangen?

Schnell wird klar, dass die Aspekte 3 und 4 in einem gewissen Konflikt zueinanderstehen; dies ist jedoch die Herausforderung des Datenschutzes, denn wir müssen einerseits wissen, ob eine Information von einem berechtigten Sender stammt, andererseits auch gewähren, dass dieser Sender anonym bleibt.

In Kürze sei hier angedeutet, mit welchen Methoden man die sechs Aspekte adressiert.

«Privacy» kann mit Verschlüsselung oder mit Hilfe von Direktverbindungen ohne Aussenzugriff erreicht werden.

«Integrity» kann mit kryptologischen Methoden, z.B. Digitalen Signaturen, erreicht werden. Zudem muss sichergestellt werden, dass die Meldung E wahrheitsgetreu.

«Authenticity» wird ebenfalls mit kryptologischen Methoden (z.B. Digitalen Signaturen) erreicht.

«Anonymity» sicherzustellen, benötigt geeignete Daten- und System-Strukturen möglicherweise ergänzt mit kryptologischen Verfahren.

«Authorization» wird auf über Zugriffskontrolle und auf Basis von sicherer Identifikation von Sender und Empfänger erreicht.

«Availability» wird erzielt durch Sicherstellung der Qualität der Infrastrukturelemente, durch Redundanzen und – sehr wichtig im IT-Zeitalter – durch Aussperren von von Fremdeinflüssen (z.B. Malware).

Diese kurze Andeutung soll mit den nachfolgenden Abschnitten konkreter beleuchtet werden.



10.6.3 Atmende Energie-Versorgung und Smart Grids

Motivation und Ausgangspunkt für die vorliegenden Arbeiten im Themenbereich IoT-Sicherheit ist die «atmende Energie-Versorgung». Die Motivation für diese Art der Versorgung sowie die entsprechende Funktionsweise wird ausführlich in anderen Dokumenten dargelegt.

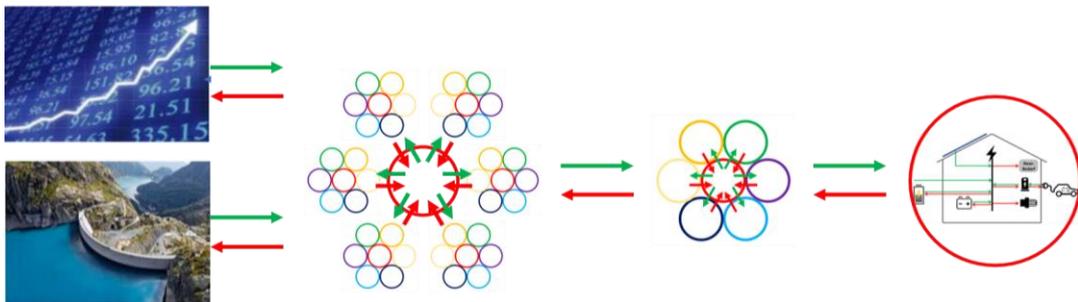


Figure 66: Übersicht atmende Versorgung ausgehend vom Prosumer (rechts im Bild).

Diese atmende Versorgung setzt beim dezentralen Verbraucher und/oder Erzeuger, Prosumer genannt, an; hier wird – falls vorhanden – die dezentrale Produktion (z.B. PV) und der Verbrauch gemessen, sowie ermittelt, wie viele Energie aufgenommen („eingatmet“) oder abgegeben („ausgeatmet“) werden kann (Flexibilität).

Mehrere solcher Endnutzer werden in Gruppen zusammengefasst (z.B. ein Quartier), welche wiederum zu grösseren Einheiten (z.B. Ortschaft, Region) gebündelt werden. Auf Basis der dezentral gemessenen Daten können auch für all diese Gruppen und Einheiten Produktion, Verbrauch und Flexibilität berechnet werden.

Damit wird eine Transparenz geschaffen, dank welcher die Energieflüsse optimiert und die Ressourcen effizient eingesetzt werden können.

Voraussetzung dafür ist: Die Energieflüsse müssen dezentral an vielen verschiedenen Orten gemessen und auch beeinflusst werden können, was ja die Aufgabe von Smart Grids ist (vgl. [19]).

Messungen und Beeinflussungen müssen in Echtzeit vollzogen werden, was konkret wie folgt zu verstehen ist:

Table 17: Bedeutung des Begriffs „Echtzeit“ in Smart Grids

Aktivität	Frequenz	Latenz
Messung	1" – 10"	0.1" - 1"
Analyse	10" – 30"	1" - 5"
Beeinflussung	ca. 60"	10" - 15"

Die „Frequenz“ sagt aus, wie oft die Handlung vollzogen wird, und die „Latenz“ ist die Zeitspanne, innerhalb welcher die entsprechenden Resultate vorliegen sollen.

Heute werden in modernen intelligenten Mess-Systemen (IMS, vgl. [26]) Daten alle 15' erhoben, also $4 \times 24 = 96$ Erhebungen pro Tag; die entsprechenden Resultate werden ein Tag später verfügbar gemacht werden (Latenz: 24h).

Eine atmende Versorgung ist jedoch auf Werte wie in Figure 66 dargestellt angewiesen!

Für die atmende Versorgung der aliunid wird daher eine Smart Grid Lösung auf Basis eines IoT-Ansatzes vorgeschlagen, was in Figure 67 schematisch dargestellt ist: An vielen (sehr viele!) verteilten Punkten werden mit Hilfe von verschiedenen Methoden die Daten der Energieflüsse erfasst



(Messpunkte); wiederum an vielen (sehr viele!) Punkten werden mit unterschiedlichen Verfahren Energieflüsse beeinflusst (Steuerpunkte). Mess- und Steuerpunkte werden zusammen als Datenpunkte bezeichnet.

Die Datenpunkte sind über Kommunikationskanäle mit Rechenwerken verbunden, welche einerseits Daten analysieren, darauf basierend Prognosen erstellen und geeignete Aktionen einleiten („Intelligenz“). Zudem werden die Daten geeignet „persistiert“ (Wikipedia: Persistenz ist in der Informatik der Begriff, der die Fähigkeit bezeichnet, Daten oder logische Verbindungen über lange Zeit bereitzuhalten.). Damit können Auswertungen und Bewertungen (u.a. Rechnungen) erstellt werden.

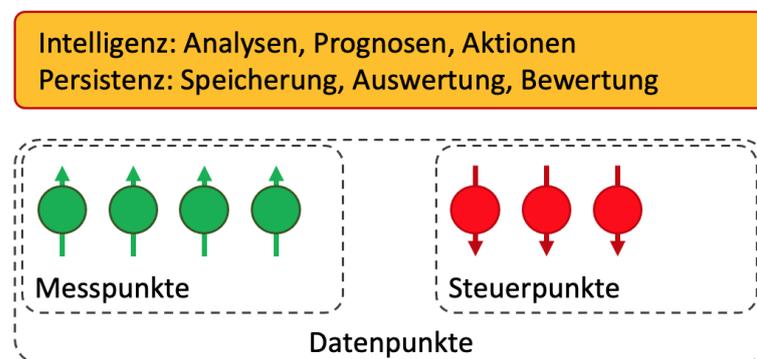


Figure 67: Schematische Darstellung der Funktionalität eines IoT-basierten Smart Grids.

Obschon in dieser schematischen Darstellung die Datenpunkte „uniform“ behandelt werden, ist doch die Bedeutung und vor allem die Sensitivität dieser Datenpunkte unterschiedlich. Diese Feststellung hat für das Thema Sicherheit in IoT-Systemen sehr zentral!

Messpunkte sind ungleich relevant, was mit folgendem Beispiel illustriert wird: Fehlen die Produktionsdaten eines Gross-Kraftwerks ist die Gesamtsicht auf das Energienetz stärker eingeschränkt als wenn die Verbrauchsdaten eines einzelnen Gerätes beim Endkunden nicht vorliegen. Entsprechend sind die Anforderungen, welche bezüglich Sicherheit der Datenerfassung und Übertragung bei solchen Mess- und Steuerpunkten wesentlich stringenter. Dasselbe gilt betreffend die Verfügbarkeit der Rechenwerke, welche die grössere Teile eines Netzes analysieren; Falschaussagen betreffend einen einzelnen Haushalt sind weniger einschneidend als Aussagen über das energetische Verhalten einer ganzen Region. Hingegen sind die aggregierten Daten von vielen kleinen Prosumers wiederum „System-relevant“.

Genau umgekehrt verhält es sich hinsichtlich des Datenschutzes, mit welchem die Privatsphäre des einzelnen Prosumers geschützt werden soll: Die Daten des einzelnen Prosumers sollen anonym bleiben, derweil aus den aggregierten Daten einer ganzen Gruppe von Prosumer keine Rückschlüsse auf die Privatsphäre einzelner gemacht werden können.

Daraus folgt, dass bei der Bemessung der Sicherheitsvorkehrungen in einem IoT-System die Relevanz der ansonsten uniformen Mess- und Steuerpunkte berücksichtigt werden muss.

Für die Beurteilung und Bemessung der Sicherheits-Massnahmen in einem IoT-System sollen die folgenden **Security Policies** dienen, welche grundlegende Eigenschaften beschreiben, die ein sicheres Smart Grid System erfüllen muss:

Policies:



- P1 Das System liefert **korrekte** (realitätsnahe) **Messdaten**
- P2 Die unterliegende **Grundfunktion** des Systems (Energieverteilung) ist **immer verfügbar**
- P3 Das System muss mit **mangelnden Informationen** umgehen können
- P4 **Kritische Informationen** zum Energiesystem sind **geschützt** vor Fremdzugriff
- P5 Die **Privatsphäre des Endverbrauchers** wird erhalten

In den folgenden Kapiteln wird jeweils Bezug auf besagte Policies genommen. Dabei wird mittels Massnahmen und Regeln ein entsprechendes **Security Model** formulieren welches sämtliche Policies abdecken soll.

In der Einleitung zu diesem Bericht wurde betont, dass die Sicherheit in einem IoT-basierten System nicht primär nur mit Kryptologie, sondern vielmehr über eine geeignete Strukturierung der Daten und mittels der Architektur des Gesamtsystems erzielt wird. Dies wird nun in den zwei folgenden Abschnitten eingeführt.

10.6.4 4-Werte-Modell

Die Smart Grid Plattform zur Steuerung eines Energienetzes nach dem Verfahren der atmenden Versorgung ist wesentlich geprägt von geeigneten Strukturen und (Daten-)Modellen. Dieses Modell wurde unter dem Titel «Computer implementiertes Verfahren zur dezentralen Steuerung eines Energienetzes» beschrieben und zur Patentierung angemeldet (siehe [17])

Das Modell basiert auf sog. Gridblocks: Gemäss dem Modell wird das gesamte Energienetz in einzelne Blöcke (Gridblocks) unterteilt. Mehrere kleinere Gridblocks können zu grösseren Gridblocks zusammengefasst werden, welche ihrerseits wieder mit anderen Gridblocks zusammengefasst werden können.

Dabei beginnt man beim Prosumer (eine Wohnung, ein Einfamilienhaus, ein Firmengebäude, ein Notstromaggregat, ein Blockheizkraftwerk, etc.), wo reale Energieflüsse *gemessen* und *ermittelt* (zusammen «erfasst») werden, und wo Geräte und Anlagen *gesteuert* werden. Ein Prosumer bildet ein sog. erster Gridblock, welcher aus Sicht der Energieflüsse mit den vier (4) charakteristischen Werten Produktion (P), Verbrauch (V), Einatmen (E) und Ausatmen (A) – kurz [P,V,E,A] – beschrieben wird (siehe rechte Seite der untenstehenden Figur).

Mehrere erste Gridblocks können zu einem zweiten Gridblock zusammengefasst werden, also z.B. mehrere Wohnungen zu einem Mehrfamilienhaus. Mehrere zweite Gridblocks können zu einem dritten Gridblock zusammengefasst werden, usw. Dies wird durch den linken Teil der untenstehenden Figur dargestellt.

Wie der erste Gridblock werden auch die zweiten, dritten, vierten usw. Gridblocks aus Sicht der Energieflüsse mit den vier Werten [P,V,E,A] beschrieben.

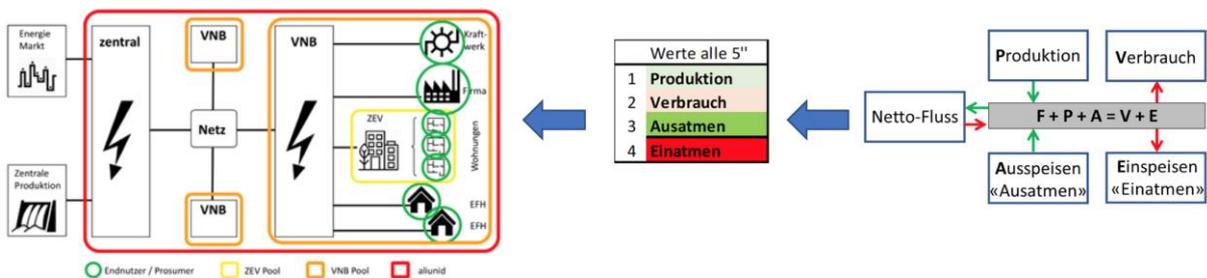


Figure 68: 4-Werte-Modell als Basis zur Steuerung eines Energienetzes.



Ein Gridblock ist also für die atmende Versorgung durch die vier Werte (i) Produktion, (ii) Verbrauch, (iii) Einatmung und (iv) Ausatmung vollständig beschrieben. Deshalb wird das Modell auch einfach «4-Werte Modell» (4WM) genannt und die Werte als Quadrupel [P,V,E,A] bezeichnet.

Mit Blick auf die Sicherheit von IoT-Systemen bietet diese Strukturierung eines Energienetzes in eine (grosse) Anzahl von in sich verschachtelten Gridblocks folgende Vorteile:

- Datenschutz 1: Die vielen bei einem Prosumer gemessenen Daten werden auf die vier Werte [P,V,E,A] abstrahiert und „nach aussen“ freigegeben. Interne Detailspekte bleiben verborgen.
- Datenschutz 2: Die vier Werte von verschiedenen Prosumer werden zu einen zweiten Gridblock aggregiert. Ab diesem Punkt kann kein Rückschluss auf einen einzelnen Gridblock mehr gemacht werden.
- Sicherheit: Jeder Gridblock stellt quasi eine autonome Einheit dar und wird in Realität tatsächlich durch ein eigenes Rechenwerk ausgewertet und gesteuert.

10.6.5 SIOT

IoT-Systeme zeichnen sich dadurch aus, dass eine (grosse) Menge von Sensoren Daten eines physischen Systems erfassen und dieses System mittels Aktoren beeinflusst wird. Hierzu werden traditionell die Messdaten an eine zentrale Cloud gesendet, welche die Daten speichert, analysiert und auswertet sowie Befehle an die Sensoren sendet, die das physische System beeinflussen. Dies wird im linken Teil der Figure 69 gezeigt.

Diese Architektur hat verschiedene nachteilige Eigenschaften hinsichtlich Ausfallsicherheit, Datenschutz, etc., was im umfassenderen Dokumente «IT-Security» detailliert wird.

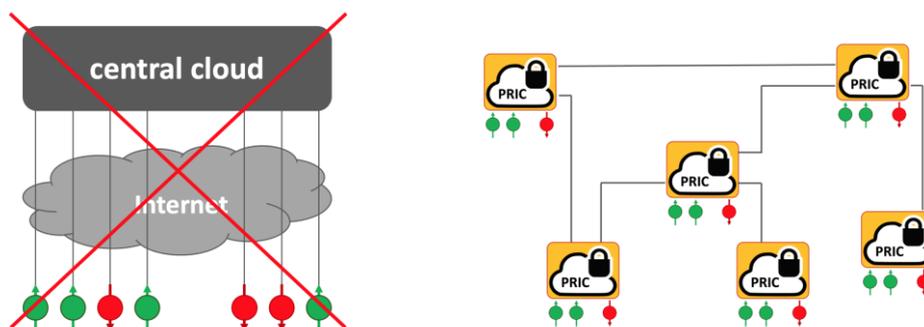


Figure 69: Klassischer Cloud-Ansatz (links) versus Schwarm-Ansatz (rechts).

Der von aliunid gewählte Ansatz «SIOT» unterscheidet sich deutlich von dem eben beschriebenen «klassischen Cloud-Ansatz»!

Der Begriff «SIOT» besteht aus «S» und «IOT», wobei der Buchstabe «S» die Kerneigenschaften von SIOT anspricht: sicher, simpel, smart, strukturiert, Swiss-made und – vor allem – Schwarm. Ein SIOT-System besteht aus einer Vielzahl von sog. «Private Clouds», abgekürzt mit dem Namen «PRIC». Ein SIOT-System ist ein Schwarm von PRICs. Jede dieser PRICs stellt eine autonome Einheit dar, welche mit Hilfe von Sensoren Messdaten erfasst und über Aktoren ein System beeinflussen kann. Die verschiedenen PRICs eines Schwarms tauschen gegenseitig Daten aus, d.h. PRIC-X sendet z.B. Daten zu PRIC-Y und erhält Daten von PRIC-Z, etc.

Ein erstes zentrales Element von SIOT ist also die Private Cloud (PRIC), deren Innenleben in Figure 70 dargestellt wird:

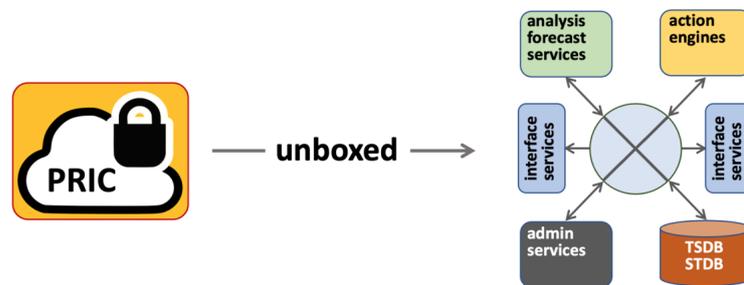


Figure 70: Struktur einer Private Cloud PRIC.

Eine PRIC arbeitet völlig autonom, erfasst Messwerte von Sensoren („interface services“), speichert Werte in einer sog. Zeitreihen-Datenbank („TSDB“), analysiert diese Werte und erstellt möglicherweise Prognosen („analysis forecast services“), trifft Entscheide für Aktionen („action engines“) sendet Befehle an Aktoren („interface services“). Weiter beinhaltet eine PRIC auch administrative Dienste („admin services“) wie z.B. ein Verzeichnis aller angeschlossenen Sensore und Aktoren wie auch die Verbindungen zu anderen PRICs. Diese Daten werden in einer Struktur-Datenbank („STDB“) abgelegt. Natürlich müssen all diese Einheiten intern geeignet miteinander kommunizieren können.

Die Figure 71 zeigt die *physische* Struktur eines gesamten SIOT-Systems. Die Private Clouds (PRIC) werden jeweils auf einem Server gehostet. Wird eine PRIC auf einem (kleinen) Server direkt „vor Ort“ beim physikalischen Prozess gehostet, spricht man von einem sog. «Fog Device». Private Clouds können aber auch als ganze Gruppe auf einem «Group Device» gehostet werden. All diese Server sind mit dem Internet verbunden.

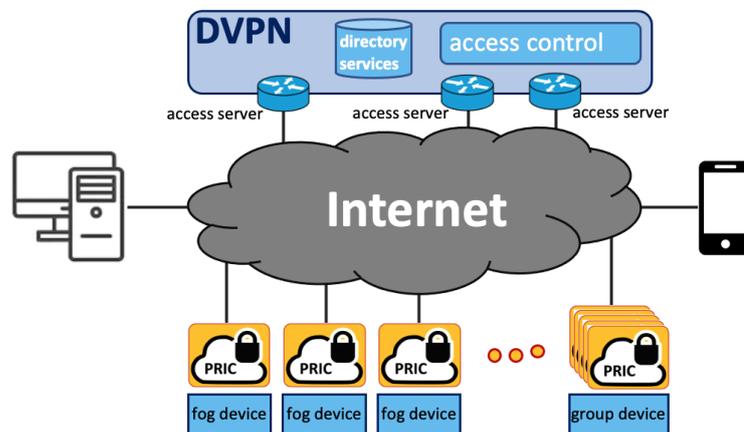


Figure 71: Physische Struktur von SIOT mit DVPN.

Ebenfalls mit dem Internet verbunden ist das sog. «DVPN» (Distributed Device Virtual Private Network), welches das zweite zentrale Element von SIOT ist.

Das DVPN ist ein verteiltes und redundantes Netzwerk, das die Verbindung zwischen den (weltweit) verteilten Private Clouds ermöglicht. Hierzu hat das DVPN die zentralen aber redundanten Funktionen «directory services», also Verzeichnis-Dienste, und Dienste, welche den Zugang zu den verschiedenen Ressourcen regeln.

Sodann dient das DVPN dazu, Nutzern den geeigneten und erlaubten Zugang zu den Daten und Funktionen der einzelnen Private Clouds zu ermöglichen. Das DVPN stellt also redundant und sicher die Verbindung unter den PRICs und und den Zugang zu Daten und Funktionen der einzelnen PRICs her.



Mit Blick auf die Sicherheit von IoT-Systemen bietet diese Schwarm-IoT Architektur folgende Vorteile:

- **Datenschutz:** Die Daten werden lokal durch die Private Cloud (PRIC) erfasst, verarbeitet und gespeichert. Nur ausgewählte Daten werden zielgerichtet mit anderen Private Clouds ausgetauscht.
- **Ausfallsicherheit:** Fallen Teile des Systems (z.B. ein Server) aus, so ist lediglich ein (kleiner) Teil des Gesamtsystems betroffen.
- **Robustheit:** Die dezentral gehostete und autonom arbeitende Private Cloud arbeitet auch dann weiter, wenn andere Teile des Gesamtsystems ausfallen sollten – möglicherweise mit reduzierter Funktion.

10.6.6 Datenfluss - Systemanalyse

Um die Sicherheit des Systems zu beschreiben folgen wir in diesem Kapitel exemplarisch dem Fluss der Daten durch das System. Dabei unterscheiden wir zwischen vier Teilsystemen /Datenflüsse:

- Der Erfassung von **Echtzeitdaten** und deren Propagation durch die Gridblock (GB) Struktur
- Die **Echtzeitsteuerung** via Direktiven und deren Propagation/Umsetzung durch die GB Struktur
- Die **Verwaltung der dezentralen Private Cloud** durch den Endbenutzer
- Das **Management der Infrastruktur** (Inbetriebnahme, Betrieb sowie Abbau der Teilsysteme)

Erfassung von Echtzeitdaten

Bei der Erfassung der Echtzeitdaten handelt es sich um einen «bottom-up» Prozess. Die Daten werden in einem ersten Schritt dezentral von Sensoren erfasst, und an die «Private Cloud», den physikalischen Gridblock weitergeleitet. In einem zweiten Schritt werden die Daten in die «4-Werte» Struktur umgewandelt und an den übergeordneten Gridblock weitergeleitet. Darauf wiederholt sich der Prozess mehrmals: Werte aller Subgridblocks werden aggregiert, die 4-Werte werden abgeleitet und an den übergeordneten Gridblock weitergeleitet.

Zu den Sicherheitsrisiken in diesem Teilsystem gehören der Ausfall von Messpunkten, Messfehler, die Manipulation von Daten bei Erfassung, Übermittlung oder Verarbeitung sowie Informationslecks. Gefährdet sind dabei primär die Security Policies korrekter Messdaten (p1), mangelnde Informationen (p3), Sicherheit kritischer Informationen (p4) sowie die Privatsphäre des Endverbrauchers (p5).

Die Massnahmen, um diesen Risiken entgegen zu wirken bestehen zu einem Grossteil aus der Absicherung der dezentralen Kommunikation. Wir zeigen dabei im Detail welche Protokolle und Technologien in drahtlosen und verkabelten Netzwerken beim Endkunden verwendet werden können und wo deren Limitationen liegen. Zentral ist zudem, dass sämtliche privaten Daten die «private cloud» aufgrund des 4-Werte Modells gar nicht erst verlassen. Eine wichtige Rolle spielt hierbei die dezentrale Identity and Access Management (IAM) Funktion der «private Cloud». Diese erlaubt die komplett dezentrale Konfiguration und Verbindung von Sensoren, was die Privatsphäre des Endbenutzers (p5) stützt. Zudem ist die Hard- und Software nach aktuellem Stand der Technik vor Fremdzugriff zu schützen (Hardening).

Um den zweiten Teil der Echtzeitdatenerfassung zu schützen, wird primär der Kanal zwischen den Gridblocks geschützt (z.B. durch Technologien wie VPN, TLS und end-to-end Verschlüsselung). Um der Privatsphäre der Endbenutzer (p5) gerecht zu werden, werden die Daten nicht nur aggregiert und damit anonymisiert; der Informationsfluss der Gridblocks wird zusätzlich mittels separater Subdomain komplett vom Datenfluss Endbenutzer und Datenfluss Infrastruktur Management getrennt.



Echtzeitsteuerung

Die Echtzeitsteuerung folgt im Gegensatz zur Datenerfassung dem «top-down» Prinzip. So genannte «Direktiven» werden zentral ausgelöst und an unterliegende Gridblocks versendet. Die lokale Gridblock Intelligenz («GB Controller») setzt dabei die Instruktionen in der Direktive zur Nutzung von Flexibilität («ausatmen» / «einatmen») in Form von neuen Direktiven um und versendet diese weiter an seine Subgridblocks. Sobald die Direktiven einen physikalischen Gridblock erreichen, werden sie vom entsprechenden GB Controller in explizite Steuerungsbefehle an die angeschlossenen Aktoren umgesetzt.

Die Sicherheitsrisiken in diesem Teil betreffen die Manipulation, Beeinflussung- oder Unterbindung besagter Direktiven bis hin zur Unterbindung von ausgelösten Steuerbefehlen im dezentralen System (physikalischer Gridblock). Zudem gibt es energietechnische Herausforderungen (z.B. Timing der Steuerung) welche bei der Implementation des GB Controller beachtet werden müssen. Diese Risiken haben das Potenzial zur Gefährdung der Grundfunktion des Systems (p2) sowie der Sicherheit kritischer Informationen (p4).

Die Massnahmen für die Steuerung decken sich grösstenteils mit den Massnahmen aus der Echtzeitdatenerfassung. Diese umfassen Technologien zur Absicherung der Verbindungen (TLS, VPN, end-to-end encryption) sowie des Schützens der Gridblocks vor Fremdzugriff (Hardening von Soft- und Hardware). Zudem wichtig sind Mechanismen zur Prüfung der Authentizität von Direktiven sowie die Implementation des GB Controllers, welcher energietechnische Risiken abfangen- und vor externen Einflüssen geschützt werden muss.

Dezentrale Verwaltung

Bei der dezentralen Verwaltung handelt es sich um den Zugriff auf den physikalischen Gridblock, der «private Cloud», durch den Endbenutzer. Dadurch erhält der Endbenutzer die Möglichkeit, seine privaten Daten (z.B. Smart Home) einzusehen, zu konfigurieren und seine lokalen Sensoren / Aktoren zu Verwalten. Der Zugriff kann direkt (LAN, WiFi, Bluetooth) oder indirekt (DVPN) stattfinden.

Sicherheitsrisiken dabei sind Fremdzugriff und Missbrauch durch Dritte sowie Missbrauch durch den Endkunden selbst (Privilegien Eskalation). Ein Angreifer könnte hier in die Privatsphäre des Endbenutzers eindringen (p5) oder gar Echtzeitmessung (p1) und Echtzeitsteuerung (p2) beeinflussen. Es ist zu bedenken, dass der Grossteil der möglichen Angriffsszenarien auf einzelne Geräte zielt, was den «Impact» der Risiken in diesem Datenfluss deutlich schmälert.

Die getroffenen Massnahmen hierbei beinhalten zum einen die Absicherung des direkten, lokalen Zugriffes «in house». Dazu gehören Technologien zur Absicherung von lokalen Verbindungen (BLE, LAN, WiFi) sowie das Konzept der dezentralen IAM Funktionalität der private cloud. Zudem wird der indirekte Zugriff «extern / online» mit entsprechenden Technologien abgesichert (VPN, TLS, End-to-end Verschlüsselung). Die Daten werden auch hier mittels separater Subdomain von anderen Datenflüssen (z.B. 4-Werte, Infrastruktur Management) getrennt.

Infrastruktur Management

Um alle Systemkomponente Instandhalten zu können, müssen Statusinformationen zu Soft- und Hardware aller im System verwendeten Geräte verfügbar sein. Sichere Softwareupdates müssen es ermöglichen selbst dezentrale Infrastruktur auf dem neusten Stand zu halten. Zudem müsse logische Identitäten (z.B. Subdomain, Gridblock IDs) bei Auf- und Abbau sicher zugewiesen und entfernt werden können ohne Raum für Missbrauch zu bieten.

Die Sicherheitsrisiken dabei umfassen unerkannte Defekte und Störungen, Sicherheitslücken aufgrund veralteter Soft- oder Hardware sowie Einführung von Schadsoftware durch unsichere Softwareupdates. Insbesondere aufgrund des letzten Punkts betrifft dieser Datenfluss alle formulierten Security Policies.

Im Zentrum der Massnahmen steht das so genannte «Logistics and Inventory Management» (kurz LAIM) System, welches die angesprochenen Funktionen zentral verwaltet. Sämtliche zentralen



Funktionen (z.B. LAIM, DVPN) werden dabei anhand best-practises zum Betrieb von IT Backends geschützt. Zudem ist die logistische Sicht der Geräte strikt getrennt vom 4-Werte / Gridblock- sowie Endbenutzer Datenfluss. Das LAIM System ist sich somit nur der Geräteidentität und dessen Soft- und Hardware Status bewusst, nicht aber der Gridblock(s) die darauf betrieben werden oder der Endbenutzer, welche die Geräte besitzen. Eine weitere Architektur Massnahmen ist, dass die gesamte Infrastruktur (LAIM, DVPN, etc.) als mobile, autarke «Site» behandelt wird. Diese kann dadurch je nach Bedürfnis vom aliunid Partnern (z.B. EVU oder Produzent) selbst betrieben werden. Dies führt zu zusätzlicher Aufspaltung der zentralen Funktionen und reduziert den Impact bei einem Sicherheitsinzident.

10.6.7 Datenschutz

Dieses Kapitel befasst sich mit den Datenschutzrechtlichen Herausforderungen des aliunid Systems. Dabei werden zuerst die geltenden Rahmenbedingungen und daraus resultierenden Herausforderungen beschrieben. Danach wird darauf eingegangen mit welchen Lösungsansätzen das aliunid System den Rahmenbedingungen im Datenschutz gerecht werden will. Zudem werden Erkenntnisse aus dem bisherigen aliunid Feldtest eingebracht.

Durch den Einsatz von Internet of Things Systemen findet beim aliunid System ein, stellenweise, tiefes Eindringen in die Privatsphäre statt. Dies bedingt, dass das System die gängigen Schweizer Gesetzgebungen zum Datenschutz (Bundesverfassung / Datenschutzgesetz) einhält. Dazu kommt, dass in der Energiebranche zusätzliche Bestimmungen zu beachten sind (i.e. Stromversorgungsordnung), welche insbesondere die Art und Weise der Echtzeiterhebung von Energiedaten limitiert.

Um diesen Anforderungen gerecht zu werden, setzt das aliunid System auf zwei Kernmassnahmen. Zum einen wird ein Grossteil der privaten Daten dezentral auf der „private Cloud“ gehalten, wo nur der Endbenutzer Zugriff hat. Die privaten Energiedaten die für die „Atmung“ im aliunid Netz relevant sind, werden dabei auf dem dezentralen Gerät auf eine 4-Werte Repräsentation „kondensiert“. Die Zweite Massnahme betrifft die Weiterleitung dieser 4-Werte. Diese werden nicht direkt an eine zentralistische Instanz weitergeleitet, sondern schrittweise in der Gridblock Architektur in Gruppen aggregiert und dadurch anonymisiert.

Aliunid geht mit dem Thema Datenschutz proaktiv um hat Gespräche mit verschiedenen Experten dazu geführt. Deren Stellungnahmen deutet darauf hin, dass die von aliunid getroffenen Massnahmen grundsätzlich mit dem Datenschutz vereinbart werden können. Eine Erkenntnis ist, dass es fundamental sein wird, dass die aggregierten Gridblock Gruppen genügend gross gewählt werden, um eine effektive Anonymisierung zu garantieren.

Im Rahmen dieser Erkenntnis wurden im aliunid Feldtests Echtzeitdaten ex-post annotiert um in einer ersten Simulation den Wirkungsgrad des „Anonymisierung durch Aggregation“ zu untersuchen. Dabei wurden erste Schlüsse auf eine mögliche „Gridblock Gruppengrösse“ gezogen.

Im Feldtest konnte aliunid zudem auch Erfahrungen sammeln wie es im Markt um den Stellenwert des Datenschutzes steht. So wurde bei Umfragen etwa klar, dass der Datenschutz zwar ein wichtiger Faktor für Endkunden ist, der Kunde aber nicht bereit ist dafür mehr zu zahlen. Für den Kauf sind nach wie vor andere Faktoren (Kosten, Funktionalität) entscheidend. Dies deckt sich auch mit den Erfahrungen von aliunid was existierende Smart Energy Lösungen betrifft. Viele dieser Lösungen basieren auf einem zentralistischen Cloud Ansatz, der das Datenschutzproblem durch eine Einverständniserklärung des Endbenutzers umgeht.

Eine weitere Erfahrung aus dem Feldtest ist, dass sich im aliunid System zwei datenschutztechnisch entgegengesetzte Interessen gegenüberstehen:

- die dezentrale Datenerfassung von privaten Endnutzern, welche nach der aliunid Gridblock Struktur schrittweise aggregiert und anonymisiert wird



- die Erfassung von Infrastrukturinformationen durch das Energieverteilunternehmen (EVU), welches aufgrund der Energiewende zunehmend mehr Einsicht in die Netzebene 7 benötigt um die Stabilität des Netzes und damit auch das „atmen“ der Flexibilität überhaupt zu ermöglichen

Die Infrastruktur von aliunid eignet sich dazu beide Fälle abzudecken, die unterliegende SIOT Lösung beinhaltet auch bereits eine inhärente Separation der beiden Datenflüsse durch Subdomains. Die Situation erfordert jedoch eine klare Trennung der beiden Teilsystemen in der Verwaltung des Systems um den Datenschutz konsequent einhalten zu können.

10.6.8 Schlussfolgerung

Mit diesem Bericht haben wir aufgezeigt welche Erfahrungen im aliunid Feldtest zum Thema IT Sicherheit und Datenschutz gesammelt wurden. Wir haben Probleme und Risiken identifiziert und gezeigt welche Massnahmen diese Risiken minimieren können und diese teilweise bereits umgesetzt. Es ist für uns offensichtlich, dass es sich hier um keinen abschliessende Sicheheitspezifikation handelt, sondern vielmehr um eine Momentaufnahme.

Eine der grössten Limitationen aus dem Feldtest ist momentan, dass die bisherigen Installationen noch keine ausgeprägten Gridblock Hierarchien abbilden. Erfahrungen zu Anonymisierung durch Aggregation beziehen sich somit nur auf vereinzelte Installationen. Die Umsetzung von Direktiven über mehrere Stufen konnte bisher nur konzeptionell erforscht werden, da Experimente zum Steuern direkt auf den physikalischen Gridblocks erfolgen. Dies betrifft auch die Management Funktionen der Infrastruktur, da es entsprechend wenig Erfahrungswerte zum Linking von Subdomains und Gridblocks gibt. Die Verwaltung der Identitäten (Messpunkte, Gridblocks) beinhalten nach wie vor einige offene Fragen.

Weitere Punkte:

- bei der Hardware (Geräte, Sensoren, Aktoren) herrscht nach wie vor eine grosse Heterogenität, was das Hardening von Soft- und Hardware erschwert
- aliunid bewegt sich als Startup aktuell von der Entwicklung zum Betrieb der Infrastruktur, dies bedeutet das Umsetzen neuer best-practises und Anpassung der Security Models an die Anforderungen des Betriebs
- Viele Sicherheitskonzepte sind zwar klar, müssen aber noch in Produktreife umgesetzt werden (z.B. End-to-end Verschlüsselung im DVPN, dezentraler Einsatz von TLS, Backend best-practises)

Daher ist es unabdingbar, dass aliunid den Feldtest weiterführt, um die offenen Fragen zu klären. Zudem wurden bei den bisherigen Resultate neue Themen entdeckt die als Folge behandelt werden sollten.



10.7 Details for breathing experiment in BFH

Author: Dimitris Kouzoupis, aliunid

As briefly explained in Section 3.4, an optimal control problem formulation is based on the dynamic model of the controlled system, an objective function and a set of constraints that need to be satisfied at the solution. In this section we go into details on how the above elements have been implemented for the demonstrator of Section 2.2.

10.7.1 Models and forecasts

In order to calculate the optimal control trajectories of the next 24 hours, the optimization routine needs to be aware of:

- The current state of the controlled system (e.g., state of charge of the battery), either readily available by measurement devices or computed using a state estimator (e.g., a Kalman filter).
- The models that describe the dynamic evolution of the controlled system as a function of time and control inputs.
- Forecasts of inputs to the system that cannot be influenced by the controller (e.g., solar radiation).

In this section, we briefly describe the models and forecasts considered for the experiment.

10.7.1.1. Battery model

Most battery models in the literature are linear models, which seem to capture well enough the dynamics of the battery, i.e., the evolution of the state of charge as a function of the power that is provided or extracted. For the purposes of the experiments, the following linear model has been considered:

$$SoC_{k+1} = SoC_k + \frac{100}{E_{max}} P_k^s \Delta t_k + \frac{100\eta^c}{E_{max}} P_k^c \Delta t_k + \frac{100\eta^d}{E_{max}} P_k^d \Delta t_k.$$

In the equation above, $k = 0, \dots, N$ denotes the discretized time index, N the total number of intervals, SoC the state of charge of the battery in percentage ranging from 0 to 100, E_{max} the battery capacity in kWh, $P_k^s \leq 0$ the self-discharging power in kW, Δt_k the size of the k^{th} time interval in hours, η^c the battery charge efficiency, $P_k^c \geq 0$ the charging power in kW, η^d the battery discharge efficiency and $P_k^d \leq 0$ the discharging power in kW.

The unknown parameters P_k^s, η^c, η^d of the model are identified using linear least squares and historical data of 4 months.

10.7.1.2. Global irradiance and PV production forecast

There are several factors that affect the energy generation of a solar panel, such as the global irradiance and the temperature of the module. For the purposes of the planned experiments, a number of simplifications have been made to facilitate the identification process. The reason for this decision is two-fold: on the one hand, the error in weather forecast is expected to be too high to justify a perfect model for the estimation of the PV production based on the weather. On the other hand, the performance of the controller is not expected to deteriorate significantly with moderate errors on the PV forecast several hours in the future, assuming the trends are captured well enough.

The theoretical output P^{th} of a PV system is given by:

$$P^{th} = P^{nom} \frac{G}{1000}$$



with P^{nom} the peak power of the PV installation in kW and G the global irradiance in W/m^2 . The actual power of the installation is calculated by multiplying the theoretical power with the performance ratio α^{PV} , i.e.,

$$P^{PV} = -\alpha^{PV} P^{th}$$

where the minus sign represents generation instead of consumption of energy. The scaling factor α^{PV} depends on the PV technology and varies with respect to global irradiance and module temperature.

For the purposes of the planned experiments, a constant performance ratio is assumed, which was identified using historical data from <http://www.agrometeo.ch> on global irradiance near the demonstrator in Biel and recorded measurements of the PV production provided by the battery software. A comparison of the simplified model with identified parameter $\alpha^{PV} = 0.84$ and the real measurements for a period of one week is shown in Figure 72. The live solar radiation forecast is obtained every hour by the API of <https://www.solcast.com>.

10.7.1.3. Forecast of non-shiftable loads

Accurate prediction of non-shiftable loads is not realistic in household level due to the high degree of uncertainty. However, the predictive controller of Section 4.2 is expected to work well enough even with less accurate predictions. Since the most relevant factor for the prediction is intuitively the day of the week and time of the day, the forecast is implemented by averaging historical data for the same quarter of the day of the same day of the week. As the effect of different weather conditions throughout the year is also expected to have a significant impact on the consumption, the averaging of the data should only be implemented for a number of weeks in the past rather than longer periods.

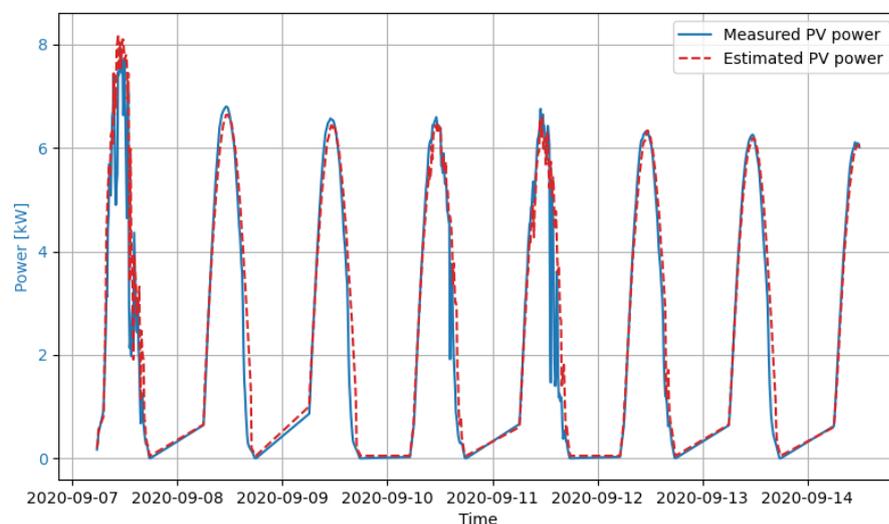


Figure 72: Comparison of simulated and measured PV production.

10.7.2 Cost functions

Two relevant cost functions have been considered. Namely, minimization of energy costs and minimization of CO2 emissions. In the case of cost minimization, the following pricing system is considered (source: www.diewerke.ch):

- High tariff: 0.18CHF/kWh (Monday to Friday 07:00 to 20:00, Saturday 07:00 to 13:00)
- Low tariff: 0.12CHF/kWh (All remaining times of the week)
- Feed-in tariff: -0.05CHF/kWh (For feeding energy from the PVs to the grid)



In the case of carbon footprint minimization, aim it to (a) minimize the electricity coming from the grid and (b) if necessary, buy that electricity at a timepoint that has the smallest impact on CO2. Note that the CO2 emissions of the own PV production can be added as a constant term in the objective, but it will not change the result of the optimization since it cannot be affected by any action of the controller. The forecast of the CO2 emissions is calculated using the algorithm of Section 3.2.3.

10.7.3 Constraints

The one set of constraints that all optimal control problems have in common is the satisfaction of the model equations along the prediction horizon. In the most general form these constraints can be written as:

$$x_{k+1} = f_k(x_k, u_k)$$

for $k = 0, \dots, N - 1$, with $x_k \in \mathbb{R}^{n_x}$ the differential states and $u_k \in \mathbb{R}^{n_u}$ the control inputs. For the experiments of Section 4.2, the only state of the system is the state of charge of the battery and its evolution is described by the equation of Section 10.7.1.1. The controls u_k in the same experiment are the charging and discharging power of the battery. Note that two controls are introduced in order to model different charge and discharge efficiency. These controls are mutually exclusive, i.e., only one of them can be non-zero. Although this could be imposed by introducing auxiliary integer variables, it is not necessary as charging and discharging a battery simultaneously can never be an optimal solution when there is an associated cost for both directions of interaction with the grid¹³. This fact is formally proved in [29]

Another common set of constraints in optimal control problems are the so-called algebraic constraints, i.e., equality constraints coupling the variables of the same time interval (instead of variables of successive intervals as above). Such constraints take the form:

$$y_k = g_k(x_k, u_k)$$

for $k = 0, \dots, N - 1$, with $y_k \in \mathbb{R}^{n_y}$ the algebraic variables. In the context of Section 4.2, algebraic constraints can be considered the balance of power on the grid block, i.e.,

$$P_k^{g2b} + P_k^{b2g} = P_k^d + P_k^c + P_k^{PV} + P_k^L$$

with $P_k^{g2b} \geq 0$ the power flowing from the grid to the household, $P_k^{b2g} \leq 0$ the power flowing from the household to the grid and $P_k^L \geq 0$ the forecast of the non-shiftable load. Recall that $P_k^{PV} \leq 0$ is the forecast of the PV production, $P_k^c \geq 0$ the charging power set-point at time interval k and $P_k^d \leq 0$ the discharging power set-point at the same interval. Note again the need to introduce two separate variables P_k^{g2b} and P_k^{b2g} due to the fact that the price for feeding in energy to the grid is typically different than the price for consumption, as discussed in Section 10.7.2. Similar to the modelling of the battery, these two variables are mutually exclusive but as long as one chooses one of the cost functions that have been discussed above, the optimizer will not converge to a solution that does not satisfy this condition. This is confirmed by the experiments of Section 4.2.2.

To respect the physical limitations of the battery, the charging and discharging power is constrained based on the available specifications, i.e.,

$$-P_{min} \leq P_k^d \leq 0, \quad 0 \leq P_k^c \leq P_{max}$$

for $k = 0, \dots, N - 1$.

A further constraint that needs to be satisfied during the experiments is to keep the state of the charge of the battery between meaningful bounds, e.g., from 0% to 100%, optionally including a safety margin. In practical applications, it is often common to relax any constraints on the states of the system

¹³ In CO2 minimization, an order of magnitude smaller cost is associated to feeding energy back to the grid to avoid simultaneous charging and discharging without altering the optimization results.



using slack variables, to avoid problem instances becoming infeasible (meaning that the optimizer cannot find any solution that satisfies the constraints). This is usually implemented as:

$$\begin{aligned} SoC_{min} - \underline{s}_k &\leq SoC_k \leq SoC_{max} + \bar{s}_k \\ \underline{s}_k &\geq 0, \bar{s}_k \geq 0 \end{aligned}$$

while at the same time penalizing with a high quadratic penalty any deviation of the variables $\underline{s}_k, \bar{s}_k$ from zero.

To conclude with the section on the constraints, we recall the initial condition constraint that is also common to every optimal control problem. That is essentially the constraint $x_0 = \hat{x}$, imposing the beginning of the state trajectory to be equal to the current measurement (or estimate) of the states. In the content of the experiments of Section 4.2.2, this boils down to providing the current state of charge of the battery to the optimization routine.

10.7.4 Summary of optimization problem

In this section, we put together the elements of Sections 10.7.1, 10.7.2 and 10.7.3 to formulate the optimization problem that we aim to solve.

The optimization problem that calculates the optimal charging/discharging profile for the battery demonstrator of Section 2.2 is therefore summarized below:

$$\begin{aligned} &\underset{SoC, P^{bat}, P^{grid}}{\text{minimize}} \sum_{k=0}^{N-1} \begin{bmatrix} \bar{s}_k \\ \underline{s}_k \end{bmatrix}^\top \bar{W} \begin{bmatrix} \bar{s}_k \\ \underline{s}_k \end{bmatrix} + \begin{bmatrix} P_k^{g2b} \\ P_k^{b2g} \end{bmatrix}^\top \begin{bmatrix} q_k^{g2b} \\ q_k^{b2g} \end{bmatrix} + \begin{bmatrix} \bar{s}_N \\ \underline{s}_N \end{bmatrix}^\top \bar{W} \begin{bmatrix} \bar{s}_N \\ \underline{s}_N \end{bmatrix} \\ &\text{subject to } SoC_0 = \hat{SoC} \\ &SoC_{k+1} = f(SoC_k, P_k^d, P_k^c, \Delta t_k), && \text{for } k = 0, \dots, N-1, \\ &P_k^{g2b} + P_k^{b2g} = P_k^d + P_k^c + P_k^{PV} + P_k^L, && \text{for } k = 0, \dots, N-1, \\ &P_k^{g2b} \geq 0, && \text{for } k = 0, \dots, N-1, \\ &P_k^{b2g} \leq 0, && \text{for } k = 0, \dots, N-1, \\ &-P_{min} \leq P_k^d \leq 0, && \text{for } k = 0, \dots, N-1, \\ &0 \leq P_k^c \leq P_{max}, && \text{for } k = 0, \dots, N-1, \\ &SoC_{min} - \underline{s}_k \leq SoC_k \leq SoC_{max} + \bar{s}_k, && \text{for } k = 0, \dots, N, \\ &\bar{s}_k \geq 0, && \text{for } k = 0, \dots, N, \\ &\underline{s}_k \geq 0, && \text{for } k = 0, \dots, N, \end{aligned}$$

with $\bar{W} \in \mathbb{R}^{2 \times 2}$ a diagonal matrix with diagonal elements equal to 10^5 , q_k^{g2b} the prices or CO2 content of electricity, q_k^{b2g} the fixed feed-in tariff in case of cost minimization and a small constant in case of CO2 minimization and $f(\cdot): \mathbb{R}^4 \rightarrow \mathbb{R}$ the linear battery model as described in Section 10.7.1.1. P^{bat} contains the optimal trajectories of charging and discharging, i.e., the controls, and P^{grid} the trajectories of the power interaction with the grid, i.e., the algebraic variables.

The problem stated above is a convex quadratic program (QP), which can be solved efficiently and reliably using open-source software. Convexity implies that there is a globally optimal solution, i.e., an objective that achieves the smallest possible value. However, there can be infinitely many trajectories that yield that optimal value. To avoid having degenerate solutions, a common approach is to add a small (e.g., in the order of 10^{-6}) quadratic penalty on all the optimization variables and convert the problem to a strongly convex QP.

Note that adding integer controls (e.g., switching on/off the domestic hot water tank) to the optimization routine makes the problem significantly more challenging from a numerical point of view. Since there is a large interest in solving such problems, reliable commercial software exists but its pricing is often excessively high.



10.8 Details of CO₂ calculation

Author: Matthias Meier, aliunid

10.8.1 Chosen calculation approach

The chosen approach follows the one applied by [30]. Starting from a balance equation for electricity in Switzerland

$$consumption_{CH} = production_{CH} + import_{CH} - export_{CH} - loss_{CH} \quad (10.1)$$

and assuming transmission losses to be negligible, one can write a balance equation for carbon emissions:

$$x_{CH} \cdot consumption_{CH} = \sum_m I_{CH,m} \cdot production_{CH,m} + \sum_j x_j \cdot import_{CH,j} - \sum_k \tilde{x}_{CH} \cdot export_{CH,k}, \quad (10.2)$$

where

x_{CH} = CO₂ intensity in gCO_{2eq}/kWh of electricity consumed in Switzerland

\tilde{x}_{CH} = CO₂ intensity in gCO_{2eq}/kWh of electricity exported by Switzerland

x_j = CO₂ intensity in gCO_{2eq}/kWh of electricity imported from neighboring country j

$I_{CH,m}$ = CO₂ intensity in gCO_{2eq}/kWh of production type m in Switzerland

$production_{CH,m}$ = production of type m

$import_{CH,j}$ = import from country j to Switzerland

$export_{CH,k}$ = export from Switzerland to country j

For the purpose of this report, we assumed that the carbon intensity x_{CH} of consumed electricity is equal to the carbon intensity of exported electricity \tilde{x}_{CH} . This is a non-trivial assumption and is discussed in the next section. Given this assumption however, the consumption term on the left-hand side in the above equation can be rewritten by Eq. (10.1) (neglecting the losses) and the export term cancels out on both sides, thus leaving us with the following equation:

$$x_{CH} \left(\sum_m production_{CH,m} + \sum_j import_{CH,j} \right) = \sum_m I_{CH,m} \cdot production_{CH} + \sum_j x_j \cdot import_{CH,j} \quad (10.3)$$

Solving for x_{CH} we then arrive at Equation (3.1) which we used to determine the carbon intensity of consumed electricity in Switzerland.

In contrast to [30] we then make a simplification and assume that the carbon intensities x_j of imported electricity correspond to the carbon intensities of electricity *produced* in the respective countries (Equation (3.2)). Therefore, we neglect the fact that the carbon intensity of electricity of neighboring countries in principle also depends on their imports. An in-depth analysis of the justification of this assumption has not yet been made by aliunid. In general, it can be stated that it is justified if for each neighboring country at least one of the following two points is satisfied:

- The ratio between the amount of electricity it imports and the amount of electricity it produces is small.
- The deviation in carbon intensity between electricity it imports and electricity it produces is small.



For Germany - the country with the greatest impact on electricity import related carbon emissions of Switzerland - a rough analysis showed that the first point is clearly satisfied.

10.8.2 Discussion of the chosen approach

As intuitive as the approach described in the previous section may seem, the way of accounting for electricity imports when deriving the CO₂ intensity is ambiguous and choosing the right approach is also a matter of the goal that is supposed to be met. Concretely, the ambiguities lie in the following two domains:

- Distinction between CO₂ import and CO₂ transit.
- Determination of the carbon intensity of imported electricity.

The two points are discussed separately in the following although they are not completely independent of each other.

Distinction between CO₂ import and CO₂ transit

Cross-border flows at the Swiss borders are not only motivated by the electricity requirements of Switzerland, but Switzerland also strongly acts as a transfer platform for electricity trade between neighboring countries. It is therefore not obvious what part of carbon emissions shall be accounted to Switzerland. Merely from the physical cross-border flows it is impossible to identify what share of the import/export imbalance corresponds to “pure transit” (e.g., Italy buying electricity from Germany, delivered through Switzerland) and what share is caused by “sequential trading” (e.g., Swiss traders buying electricity from Germany and Italian traders buying the same amount of electricity on the Swiss market at the same time).

By setting the carbon intensity of electricity consumed in Switzerland equal to the carbon intensity of electricity exported by Switzerland (see previous section), we chose an approach that is somewhat in between the two cases. Namely we imply that all electricity entering Switzerland immediately “mixes” with the rest of Swiss electricity. As a consequence, the carbon intensity of electricity that is entering Switzerland purely for transit is “diluted” by the carbon intensity of the Swiss production mix. On the other hand, electricity that is leaving Switzerland due to trading with Switzerland gets diluted by imports from other countries entering at the same time. Overall, however, given that exact information about the underlying trading is not available, this seems to be a reasonable approach since it somewhat averages over the two cases.

Determination of the carbon intensity of imported electricity

As described in 10.8.1, for the scope of this report we assumed that the carbon intensity of imported electricity from a neighboring country is equal to the carbon intensity of all the electricity produced in that country (Equation (3.2)). Apart from the fact that this approach neglects imports of second order, it is in principle also blind to the underlying economic interconnections.

An alternative approach was applied by Romano et al. [15] by taking the marginal costs of electricity production into account. Their argumentation is that our neighboring countries, left by themselves, would optimize their electricity production based on marginal costs, meaning that they would dispatch power plants according to the merit order until the domestic demand is covered. Any dispatch of power plants with higher marginal costs than the last plant required to cover the domestic demand is then regarded to be caused by foreign demand. As a consequence, the carbon intensity of Swiss electricity would be estimated to be significantly worse since most of the technologies of high marginal costs in Germany - the country contributing the largest share to Switzerland's electricity imports - are of fossil nature.



While the approach of Romano et al. is focusing on countries as actors – which is definitely desired when assessing political measures and estimating national responsibilities for carbon emissions – for the purpose of impacting carbon emissions via load shifting it might be more appropriate to focus on the actual market members as actors. In that case, one can hardly argue that a Swiss trader buying electricity on the German market has a smaller impact on carbon emissions than its German counterpart buying the same amount of electricity. And the approach chosen for this report might actually quite well reflect the fact, that the economical coupling, due to the costs associated to cross-border transmission capacity, naturally decreases with every additional border that needs to be crossed.

10.8.3 ENTSO-E data correction

Since Switzerland is not legally bound to publish their data on the ENTSO-E transparency platform and many production plants fall below the publication limit of a capacity of 100MW, Switzerland's publication of production per production type is not complete. More than 30% of the data is not published on ENTSO-E, when comparing with the *Schweizerische Elektrizitätsstatistik 2019 (SES)* [11]. This is shown in Figure 73.

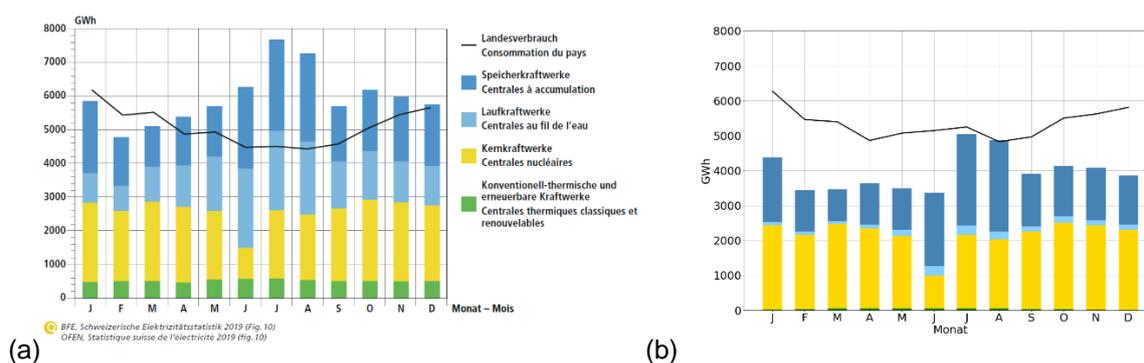


Figure 73: (a) Electricity production per category according to Schweizerische Elektrizitätsstatistik 2019 [11]. (b) Electricity production according to the data published on ENTSO-E.

A closer look into the data reveals that most of the deviation clearly lies in the two categories “Laufkraftwerke” (Run-of-River plants) and “Konventionell-thermische und erneuerbare Energien” (conventional-thermal and renewable sources). The procedure to correct the two categories is described in the following.

Run-of-river plants

In case of the run-of-river plants, Switzerland publishes production data of only a small fraction of its plants. To derive the amount of missing capacity, the production according to the data published on ENTSO-E is compared with the mean expected production of all run-of-river plants according to the *Statistik der Wasserkraftanlagen der Schweiz (SWKS)* as of 31.12.2019 [31]. This leads to a deviation by a factor of 9.1 (see Table 18). Assuming the published run-of-river plants (Löttschen, Amsteg, Laufenburg, Rheinfelden and some not specifically listed plants below 100MW) represent the weather conditions for the other run-of-river plants in a reasonable manner, the production according to ENTSO-E is linearly scaled by this factor.



Table 18: Comparison of production of run-of-river plants as published on ENTSO-E with the mean expected production according to the SWKS [SWKS2019].

Production 2019 according to ENTSO-E	1947 GWh
Mean expected production according to SWKS	17802 GWh

Conventional-thermal and renewable sources

According to SES, conventional-thermal and renewable sources produced roughly 6000 GWh in the year 2019. Of these 6000 GWh, 500 GWh of solar and wind production are published on ENTSO-E. For the remaining 5500 GWh it is assumed that they are mostly produced by non-fluctuating energy sources and aliunid therefore corrects it by adding band energy with a constant power of 628 MW.

The result after correction of the two categories is shown in Figure 74.

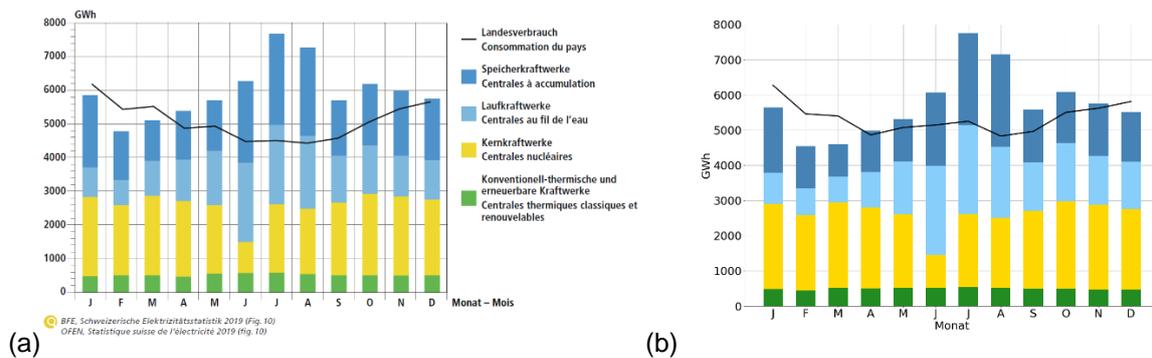


Figure 74: (a) Electricity production per category according to [11]. (b) Electricity production according to *corrected* data from ENTSO-E.

10.8.4 CO₂ emissions per production type

Below is the description of the sources and procedure to get the carbon intensities per production types of Switzerland and its neighboring countries. The results are shown in Table 19Table 19.

Neighboring countries

Except for hydro pumped storage, the carbon emissions per production types of the countries DE, AT, FR and IT are all taken from [32], which is based on the ecoinvent database v3.4 [10]. In case of hydro pumped storage, determining carbon emissions is difficult since they depend on the carbon intensity of the electricity when the water is pumped into storage. Here we rely on the values of the ElectricityMap project [30], which calculates grid carbon intensities for several countries based on flow tracing [32]. They are therefore able to calculate reasonable long-term national averages of grid carbon intensity which can be used as estimated carbon intensity for electricity from hydro pumped storage.

Switzerland

Carbon emissions of nuclear, hydro run-of-river and hydro reservoir are again taken from the ecoinvent database [ecoinvent] whereas carbon emissions from solar and wind power are taken as the averages of the values of the neighboring countries from [32]. For hydro pumped storage, an average of the neighboring countries is taken as well but weighted by the net imports per country of the year 2019.



For the remaining category of conventional thermal and renewable energy, the different origins need to be dissected. From Figure 75 we see that 50% of the category conventional thermal and renewable are *diverse Erneuerbare* - of which another 8% is solar and wind which is already published on ENTSO-E and taken care of separately (see above). 30% are *konv. thermisch (nicht erneuerbar)* and 20% are *konv. thermisch (erneuerbar)*. The resulting ratios between the non-published conventional thermal and renewable sources (628 MW, see 10.8.3) is then used to derive the carbon intensity of the combined band energy (listed as “conv. thermal and renewable” in Table 19).

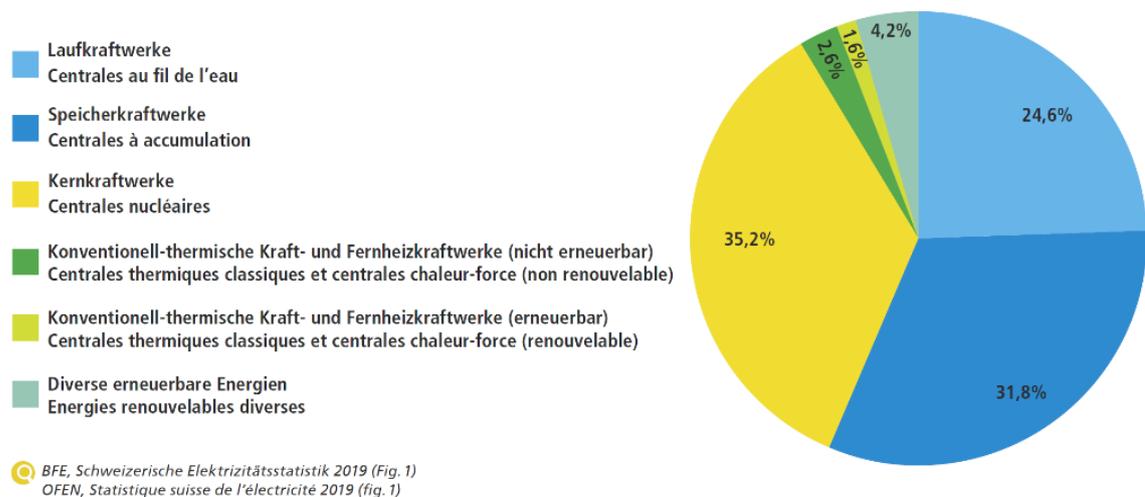


Figure 75: Energy production per category, from [11].

Thereby, the non-renewable part of the thermal plants is assumed to be mainly fossil gas (Alstom Open Cycle Gas Turbine, Birr, Aargau [33]), whereas the renewable part is assumed to be biomass. For both, the corresponding average carbon intensities of the neighboring countries from [32] are taken. The remaining non-thermal renewable sources are, according to SES, biomass as well and the same value from [32] is taken.



Table 19: CO₂ intensity in gCO₂eq/kWh per production type and country used for the calculation of the CO₂ intensity of electricity consumed in Switzerland (Section 3.2). Empty fields are not required in the calculation.

	CH	AT	FR	DE	IT
<i>biomass</i>		53.8	53.8	53.8	53.8
<i>fossil brown coal</i>				1170	
<i>fossil coal derived gas</i>				1170	1150
<i>fossil gas</i>		614	588	533	532
<i>fossil hard coal</i>		986	1090	1170	1150
<i>fossil oil</i>		1160	953	877	1060
<i>geothermal</i>		81.8		81.8	81.8
<i>hydro run-of-river</i>	13	4.42	4.42	4.42	4.42
<i>hydro water reservoir</i>	13	6.97	6.97	51.4	6.97
<i>hydro pumped storage</i>	198	173	48	281	328
<i>nuclear</i>	12		12.9	11.3	
<i>other</i>		323		657	443
<i>other (renewable)</i>		107		110	
<i>solar</i>	97.3	107	90.9	110	81.3
<i>waste</i>		427	427	427	427
<i>wind (offshore)</i>				20	
<i>wind (onshore)</i>	18.3	17.8	15.6	20	19.7
<i>conv. thermal and renewable</i>	208				

