Summary

This paper explores how fleets of electric vehicles (EVs), can be aggregated and scheduled as a single unit, called a FlexOffer, in order to participate in electricity markets, specifically the day ahead spot market and the manual frequency restoration reserve market.

As renewable energy sources are a large part of today's electricity production, the balance responsible parties are in need of flexible demand-side resources like EV charging in order to balance fluctuation in the energy consumption. In other words, EVs can be charged, when we have an abundance of energy, or they can be idle when the power grid struggles to keep up with the energy demand.

The key modeling tools used in this paper are FlexOffers, which are used for capturing flexibility in electric devices, such as batteries, heat pumps, washing machines and EVs. We use two types of advanced FlexOffers, Total energy constraint FlexOffers and dependency-based FlexOffers.

In the danish electricity markets, there are minimum bid size of 1 mWh, which is way more than what an EV can charge. Therefore FlexOffers representing a single EV must be aggregated into larger ones, in order to meet market requirements. To do this we experiment with known clustering techniques such as bottom up agglomerative clustering and k-means, and for aligning the clustered FlexOffers we use the known start alignment technique, and we introduce our own alignment techniques called flexibility alignment and a naive version of it, called fast flexibility alignment.

Next we have developed an idealized two-stage linear programming model that we run with perfect price forecasts. The first stage makes bids in both the spot and reserve markets, based on a price forecast. Stage 2 then adjusts operations depending on reserve activations of the EVs. To do this, we experimented with different strategies, namely:

- **Sequential scheduling:** Here we first schedule our bids on the reserve market, and afterwards we schedule the bids on the spot market. The idea is, that we want our bids in the spot market, to be based on the bids we made in the reserve markets. This means that we might not be able to make as much money as possible in the spot market, but hopefully the gains in the reserve market will make up for it.
- **Joint scheduling:** Here we co-optimize both markets at the same time, in order to capture opportunities in both markets, and not be restricted to one of them.
- **Spot-only:** In this case we only make bids on the spot market, and ignore the reserve market.

We then run experiments with fleets of electric vehicles of up to 500.000 EVs over a one-month period, which shows:

- Dependency-based FlexOffers consistently outperform total energy constraint FlexOffers, regardless of which bidding strategy is used.
- Joint optimization yields the highest savings, with up to 92.14% of the theoretical optimal.
- Fast flexibility alignment reaches much better a runtime than the ordinary flexibility alignment, while it is only very slight worse in terms of savings.

Our thesis improves on this domain by: (1) Improving aggregation time with a new fast flexibility alignment. (2) Inputting aggregated FlexOffers into a linear programming problem (LP) that co-optimizes day-ahead, mFRR capacity, and activation bids according to realistic driver behavior. (3) We show that this approach scales to 500.000 EVs with good performance.

Optimizing EV flexibility for spot and mFRR market participation

Raymond Asoklis Kronborg[1], Mikkel Müller Smedt[1], Andreas Ravnholt Jensen[1]

[1]Department of computer science, Aalborg university

Abstract

The continuing transition towards an electrical grid which has a greater proportion of its electricity being generated from renewable energy sources (RES) introducing more fluctuating energy prices. Flexibility in the electrical demand could introduce considerable savings, shifting demand from hours of scarcity to hours of abundance. This paper considers a Balance responsible party managing a fleet of electric vehicles (EVs) participating in the danish electricity markets. The EVs are modeled as so-called FlexOffers, models which describe the inherent flexibility of the EVs. The Total energy constraint and Dependency-based FlexOffer are the considered variants. This paper experiments with different ways to cluster, align and then schedule the FlexOffers in the danish spot market, mFRR reserve / activation market, in order to maximize profit from participating in the mFRR market and minimize the cost of the purchased electricity used to charge the EVs from the spot market.

1 Introduction

1.1 Background and motivation

In recent years, the Danish power grid have achieved good progress in integrating renewable energy sources (RES). Wind and solar generation are replacing power plants at an increasing rate with renewable energy sources having accounted for almost 50% of Denmarks final electricity consumption in 2022 [1]. Such a high level of electricity coming from RES also increases the need for a grid that can handle this uncertain level of generation. Specifically, the grid needs to be flexible such that it can handle sudden spikes in energy production. An example of flexibility could be demand side flexibility from EV owners, who can change their charging patterns. Transmission system operators (TSOs) therefore seek flexibility. I.e, the ability to adjust generation and consumption easily in order to maintain balance when RES production is variable. In Denmark, official forecasts of future energy systems predict a large rise of almost a million electric vehicles by 2030 [2].

Demand-side resources, specifically electric vehicle (EV) charging, can be used as a scalable and effective way to provide this flexibility. EVs are essentially large mobile batteries with somewhat predictable charging patterns. Shifting EV charging can either help absorb excess electricity production from RES or decrease electricity consumption during hours with excess demand, which contributes to balancing the grid. In the future, bidirectional charging using V2G (vehicle-to-grid), can even return energy to the grid if needed [2]. In the long run, this means that a large fleet of EVs will be a substantial resource of flexibility for the power grid.

Utilizing all EVs in a fleet at large scale requires strong coordination. Individual EV owners can only charge small amounts of energy and are uncoordinated. On the contrary, third party aggregators can combine the individual charging profile of many individual EVs into a single large virtual EV. Generally these aggregators serve as market actors and can combine small flexible energy assets into a larger unified asset [3]. By aggregating the fleet into a few larger unified assets, it reaches the size that is required to participate in energy and reserve markets. Thus it is possible for the aggregators to bid and schedule EV charging sessions when prices are cheap, which also often coincides with a high level of RES production. EV owners are thus incentivized to allow this controlled charging from a third party aggregator for cheaper prices and use of RES. Furthermore, in Denmark the TSO (Energinet) encourages EV aggregators to participate in ancillary services like reserve markets. By allowing them to bid into reserve markets, the aggregator can provide balancing for the TSO [4].

1.2 FlexOffer Model description

To succesfully integrate EV fleets flexibility for proper market participation requires precise modeling, aggregation and scheduling of each individual EV charging session. There are multiple ways to represent these flexible EV charging loads formally, but recent research introduces the so called *FlexOffer* (FO) model [5]. A FlexOffer (FO) in our case encodes a predicted EV charging session, capturing the flexibility of it [6]. A description of a FO is available in [7].

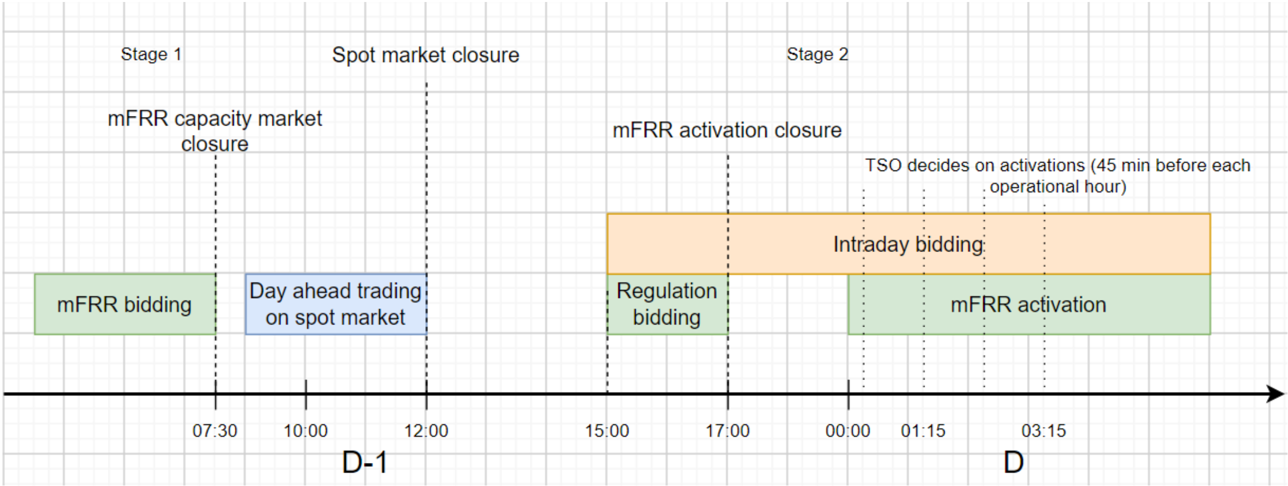

Figure 1: *Timeline of market events and closures*

1.3 Total Energy constraint FlexOffer

The TEC-FO is a variant of the normal FO seen on figure 2. The image illustrates an instance of a TEC-FO. Each bar represents a time unit (eg. one hour), and the the allowable energy to be scheduled to the EV at that time unit. The lower-bound in this example is then 0 kWh, which means we can decide to not charge the EV, and the upper-bound is 11 kWh, which means we could charge up to 11 kWh for this timeslice. The upper-bound constraint is of course determined by the physical properties of the EV. The TEC-FO also enforces a total energy constraint, ensuring that the sum of energy over all time units to be within a bound. For instance, the total energy allocated to this EV has to be between 36 kWh and 40 kWh. The full formal definition is seen in [7].

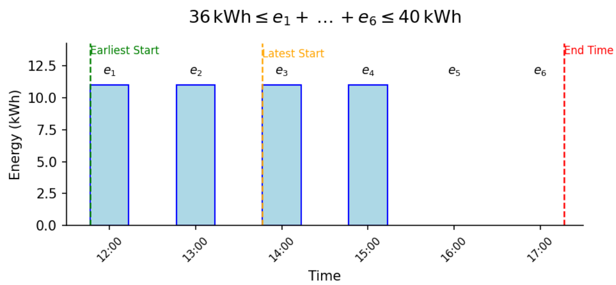

Figure 2: *Example of TEC FO*

1.4 Dependency-based FlexOffer

The Dependency-based FlexOffer (DFO) is a variant of the standard FO model, originally introduced in [8], with the intention of improving the dependency modeling between time steps, in order to better capture the constraints and state of charge (SoC) of a battery-like device e.g. EVs. For each time unit, the DFO has a 2-dimensional polygon describing the allowed minimum and maximum energy usage for that time step, given the total cumulative energy used in the previous time steps (referred to as the *dependency* amount). An example of this can be seen in figure 3.

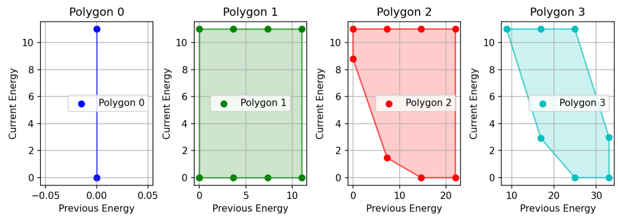

Figure 3: *Example of a DFO*

Each polygon models this, by having the dependency amount on the x-axis, and the allowed energy usage on the y-axis. Given a specific *dependency* amount, the allowed energy usage range are then all the y-values of the points within the polygon on that specific x-value. This way, the DFO also models total energy constraints implicit. The full formal description is also seen in [7].

1.5 FlexOffer lifecycle

The lifecycle of a FO captures the start to end process of utilizing the flexiblity of EVs for market participation. In the construction phase, automated software analyze and forecast the expected charging pattern of each EV - based on user behavior, driving needs and battery level. With this, a FO is constructed, which represents a set of constraints for energy usage, capturing the flexibility of each EV. Flexibility includes both amount and time flexibility, which describe how much energy the EV needs (amount flexibility) and within what time window the charging can take place (time flexibility). The FOs are then sent to the aggregator, who collects multiple FOs during the aggregation phase. In this phase the aggregator combines these into a few unified aggregated FlexOffers (AFOs). These AFOs are then scheduled through optimization to minimize the cost of electricity. This results in a large schedule (a discrete time and energy amount) which is then disaggregated to the original users, respecting the constraints of the individual FOs. The user then executes the assigned schedule (charging their EV by the specified amount in the specified hours).

1.6 Objectives and contribution

Having established the lifecycle of an FO, we now focus on (1) developing methods for scheduling aggregated EV fleets for participation into multiple energy markets, specifically the danish spot market, mfRR capacity market, and the regulation market. (2) improve the speed and quality of aggregation with new aggregation techniques, and (3) maximize multi-market revenue under realistic constraints, but assuming perfect knowledge.

In order to build a multi market scheduling framework, we start by describing the relevant structure of Danish electricity markets; then we survey EV-aggregation and related papers.

2 Related work

2.1 Danish electricity markets

In Denmark, the electricity system operates within a multi-market framework. The market events follow a structured timeline as shown on figure 1. The spot market (*elspot*) is where the largest quantity of energy is traded. Aggregators or Balance response parties (BRP) submit their hourly bids by 12:00 day D-1, a day before realtime operations begin. In the same day, the aggregator can place potential mFRR capacity offers in the *manual Frequency Restoration Reserve (mFRR)* before 07:30 on day D-1. On day D during real time operations, the *mFRR activation market* opens (also refered to as just activation or regulation market). If the TSO determines an imbalance in real time, they may choose to activate the mFRR capacity offers from aggregators that provided capacity reserve about 45 minutes in advance. The aggregator thus receives a regulation response profit for the reserve capacity they delivered, furthermore, they gain a small profit for just providing capacity even if they are not activated. [10] If the aggregator activate the reserves, the fleet should be able to respond within 15 minutes. In reality there are multiple different ancillary markets, however the figure only includes the markets we are considering, hence markets like FCR and aFRR are not included.

Most studies have taken advantage of this two day market structure. Stage 1 (day D-1) decisions are usually made deterministically with forecasted spot and reserve prices while real time operations in stage 2 (day D) are dependent on uncertain activation chances.

2.2 Literature review on EV optimization

[11] provides a comprehensive review of charging of EVs by detailing the key role of aggregators, charging infrastructure, and economic opportunities in different markets. They show that using smart charging, charging costs can be cut by 30% and grid operational costs by 10%. [12] was one of the first to develop a day-ahead market framework integrating EV aggregators. They capture aggregate State-of-Charge (SoC) constraints and bids for flexible energy and reserve capacity. [13] determines an optimal bidding strategy for aggregators in Day-ahead and reserve markets through stochastic optimization. Stochastic programming is common for optimizing multi-market revenue with uncertain prices and is seen in many papers [14] [15]. [16] extends this model by introducing a "reserve capacity ratio" to tune the risk of not being able to meet reserve demand. Another common approach to optimize in the reserve market is to use chance-constraints like in [17] to explicitly set a max limit for probability of failing to deliver promised reserve, and integrate a penalty for delivery failure. Furthermore some studies also include specific battery properties like aging constants [18]. [19] extends further by introducing a stochastic bi-level optimization where an aggregator optimizes its own day-ahead first. EV-owners then decide on charging based on the aggregators prices. Very recent literature has begun to address new reserve market formats that better capture flexibility of EVs. [20] point out that existing reserve markets cannot fully capture the cross-temporal nature of EV flexibility (limited energy capacity), and propose a new "Energy Reserves" product. [21] develop a MILP program for entire EV fleets that co-optimizes day-ahead energy and ancillary servicse (frequency-regulation) bids, while also modeling driver charging patterns, SoC, and network constraints. Lastly [22] uses a data driven approach by using reinforcement learning (RL) to learn optimal trading patterns. Other works combines EV charging sessions together with grid constraints across markets by letting an RL-based controller adapt to charging rates to avoid transformer overloads [23].

Our thesis improves on this domain by: (1) Improving aggregation time with a new fast flexibility alignment. (2) Inputting AFOs into a linear programming problem (LP) that co-optimizes day-ahead, mFRR capacity, and activation bids according to realistic

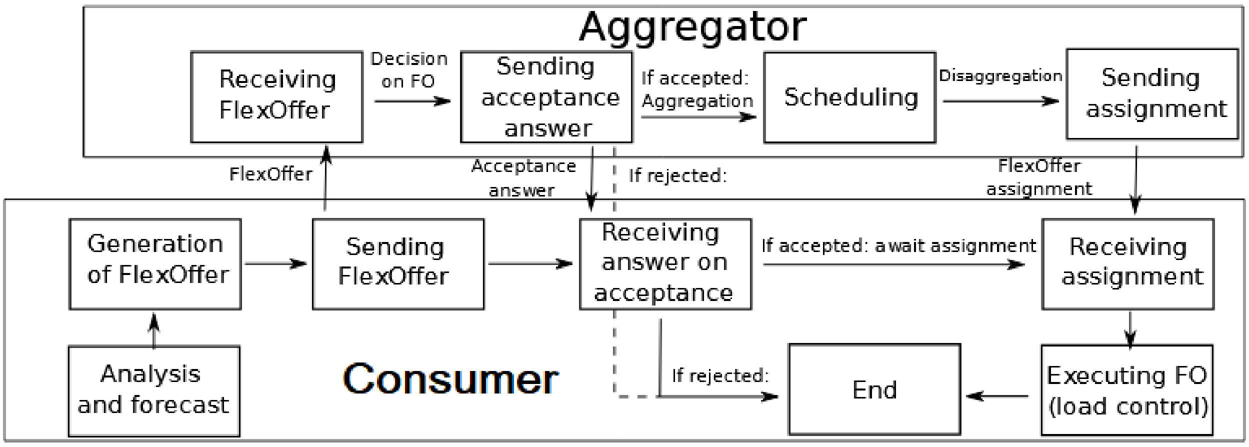

Figure 4: *Life cycle of a FlexOffer. Reproduced from [9]*

driver behavior. (3) We show that this approach scales to 500.000 EVs with good performance.

2.3 Demand response

The concept of demand response (DR) enable end-users to shift energy consumption to support grid needs based on prices or other factors. EVs provide large demand side flexibility by allowing end-users to decide when to charge through a well defined plug-in and plug-out windows. The danish TSO actually incentivizes aggregators to bid in reserve markets [24].

2.4 Aggregation using FlexOffers

Individual EV FOs are usually too small to bid directly into the day-ahead markets, furthermore as EV fleet size grows, the scheduling space expands exponentially. Given n EVs with d discrete start times, the total solution space would be d^n. This makes exact scheduling hard as the fleet grows. Market participation adds further complexity:

- mFRR bids depend on flexible charging capacity.
- Bid granularity and minimum sizes (e.g., 1 MWh spot, 5 MW mFRR) constrain actions, and introduces integer and binary variables, which makes the optimization problem harder.

An aggregator fixes this by combining multiple EV FOs into a single unified FO that meets market requirements. This process is necessary but trades off some of the potential time flexibility of the individual FOs. This means that the AFO has a smaller range of feasible charging time than the joint flexibility of its individual offers [5].

Aggregation requires aligning individual FO profiles together. Aligning all EV schedules at the same start time (a simple method called *start alignment*) yields a very fast runtime, but an imbalanced profile that loses flexibility. [3] More advanced alignment strategies uses heuristics to achieve more balanced FO with a flatter profile (net load), at the cost of extra computations. [3] One study made a set of market aware alignment strategies that fit Nordic day-ahead flexible order. These strategies automatically conforms to the market rules and led to about 20% cost savings for the aggregator [25].

Before alignment, FOs need to be clustered with other similar FOs. (E.g. making sure they have overlapping time windows or energy requirements). Common examples of clustering methods are hierarchical bottom up clustering. Bin packing heuristics can be further used to respect constraints within each group [3].

3 Problem Formulation

Based on previous chapters we define the key problem formulation as

How can flexible EV charging loads be optimally aggregated and scheduled, to maximize the overall revenue in both the spot market (Nordpool) and the manual frequency reserve regulation market (mFRR)

We now model the LP problem and then define how we cluster and aggregate in the next section.

3.1 FlexOffers and Electric Vehicles (EVs) fleet modeling

Let

- $i \in I$: Set of EV flexible loads, each represented as a FO
- $t \in T$: Set of time slots (either 24 (hourly) or 96 (15-minute) intervals for the entire day.

Each FO i represents a charging session with:

- A time window $T_i = \{t_{\text{earliest}}, ..., t_{\text{latest}}\}$,
- Duration d_i,
- Min and max charging bounds $e_{i,t}^{\min}, e_{i,t}^{\max}$.

where

$$e_{i,t}^{\min} \leq e_{i,t} \leq e_{i,t}^{\max}$$

where $e_{i,t}$ is the energy scheduled at time t, and $e_{i,t}^{\min}, e_{i,t}^{\max}$ are the minimum and maximum energy amounts allowed at time t.

We convert these energies to power variables, this aligns with market participation (specifically reserve capacity bids, which are made in kW). We convert energy bounds to average power bounds using the time unit length Δt (eg. Δt = 0.25 denotes 15 minutes):

$$p_{i,t}^{\min} = \frac{e_{i,t}^{\min}}{\Delta t}, \quad p_{i,t}^{\max} = \frac{e_{i,t}^{\max}}{\Delta t}$$

Here, $p_{i,t}$ represents the average charging power for EV i during interval t.

3.2 The optimization problem

We model an aggregator managing a fleet of EVs that can participate simultaneously in the spot market and mFRR (both up and down regulation). The goal is to schedule charging optimally such that revenue is maximized from selling reserve capacity and participating in energy markets, while ensuring that EVs are charged before their deadlines. Instead of constructing a full Mixed Interger Linear Programming (MILP) based on real world uncertainties, which may be unrealistic to solve within practical computational limits, we try to make a realistic deterministic linear programming (LP) problem that runs on perfect forecasted prices.

Parameters:

B_i	Battery capacity of EV i (kWh)
$P_i^{\max}$	Maximum charging power of EV i (kW)
Δt	Duration of time interval (e.g., 1 for 1 hour, 0.25 for 15 min)
E_i^{req}	Required energy for EV i (kWh)
$\bar{\lambda}_t^{\text{spot}}$	Forecasted Day-ahead spot price at t (DKK/kWh)
$\bar{\lambda}_t^{r\uparrow}$	Expected up reserve capacity price at t (DKK/kW)
$\bar{\lambda}_t^{r\downarrow}$	Expected down reserve capacity price at t (DKK/kW)
$\bar{\lambda}_t^{b\uparrow}$	Expected up activation price at t (DKK/kW)
$\bar{\lambda}_t^{b\downarrow}$	Expected down activation price at t (DKK/kW)
π_t^p	Penalty if reserve is not delivered at t
η_i	Charging efficiency of EV i
$\delta_t^{\uparrow}, \delta_t^{\downarrow}$	Precomputed activation reserve indicators $\in \{0, 1\}$
$\mathrm{R}_t^{\uparrow}, R_t^{\downarrow}$	The ramp limit for activating up and down reserve

As we are using perfectly forecasted prices, the imbalance will always be 0, and as such the penalty term π_t^p is negligible in our use case.

Decision variables We start by selecting the necessary decision variables from the danish energy market day D-1 from figure 1.

In stage 1, decisions are made before market closures on day D-1. This entails bidding for reserve capacity in the MFRR markets as well as procuring energy in the spot market, and are modeled by variables:

$$p_{i,t}^{r,\uparrow} \geq 0 : \text{Upward reserve bid (kW) in time } t,$$

$$p_{i,t}^{r,\downarrow} \geq 0 : \text{Downward reserve bid (kW) in time } t,$$

$$p_{i,t}^{\text{spot}} \geq 0 : \text{power at time } t \text{ (kW) bought in the day-ahead}$$

These variables must be fixed before 07:30, and 12:00 respectively. Stage 2 captures how the EV fleet responds in real-time on Day D. In the model, reserve activations depend on whether a precomputed indicator signals an expected call by the TSO. This indicator functions is simply based on forecasted expected prices. The key decision variables that we want to optimize are:

$$
\begin{aligned}
p_{i,t} \geq 0 &\quad : \text{Actual charging power for EV } i \text{ at } t, \\
\text{SoC}_{i,t} \geq 0 &\quad : \text{State-of-Charge of EV } i \text{ (kWh) at } t, \\
p^{b,\uparrow}_{i,t} \geq 0 &\quad : \text{upwards activation amount (kW) at } t, \\
p^{b,\downarrow}_{i,t} \geq 0 &\quad : \text{downwards activation amount (kW) at } t, \\
s^{\uparrow}_{i,t}, s^{\downarrow}_{t} \geq 0 &\quad : \text{slack variables for unmet reserve at } t
\end{aligned}
$$

3.3 Objective function

$$
\begin{aligned}
\max \quad & \sum_{i\in I}\sum_{t\in T} \underbrace{\bar{\lambda}^{r\uparrow}_t p^{r,\uparrow}_{i,t} + \bar{\lambda}^{r\downarrow}_t p^{r,\downarrow}_{i,t})}_{\text{Reserve revenue}} + \underbrace{\bar{\lambda}^{b,\uparrow}_t p^{b,\uparrow}_{i,t} + \bar{\lambda}^{b,\downarrow}_t p^{b,\downarrow}_{i,t})}_{\text{activation revenue}} \\
& - \underbrace{(\bar{\lambda}^{spot}_t \cdot p^{\text{spot}}_{i,t} \cdot \Delta t)}_{\text{Spot market}} - \underbrace{\pi^p (s^{\uparrow}_t + s^{\downarrow}_t)}_{\text{Penalty}}
\end{aligned}
$$

3.4 Constraints

Battery Dynamics and SoC Constraints:

$$SoC_{i,t+1} = SoC_{i,t} + \eta_i p_{i,t} \Delta t$$

$$0 \leq p_{i,t} \leq P^{\max}_i, \quad 0 \leq SoC_{i,t} \leq B_i \quad \forall t$$

$$SoC_T \geq E^{\text{req}}_i + \sum_{t\in T} p^{b,\uparrow}_t \Delta t$$

Energy Balance:

$$p_{i,t} = p^{\text{spot}}_{i,t} - p^{b,\uparrow}_{i,t} + p^{b,\downarrow}_{i,t} \quad \forall t \in T$$

This constraint ensures the actual charging power p_t for EVs at time t matches the procured spot market energy, given potential regulation actions.

Reserve market Feasibility: The aggregator's offered mFRR reserves must be feasible given the spot schedule. Ie., the spotmarket consumption at time t should be made such that the aggregator can still deliver the reserve if called by the TSO. The key constraints are:

Up-regulation (load reduction) capacity:

$$p^{r\uparrow}_{i,t} \leq p_{i,t}, \qquad \forall t.$$

Down-regulation (load increase) capacity:

$$p^{r\downarrow}_{i,t} \leq p^{max}_t - p_{i,t}, \qquad \forall t.$$

The first constraint ensures that in period t the aggregator cannot promise to reduce more load than it is actually consuming and the second constraint ensures that the aggregator has headroom to increase consumption by $p^{r,\downarrow}_t$ if down-regulation is called.

Reserve Activation Feasibility:

$$p^{b,\uparrow}_{i,t} = 0 \quad \text{if } \delta^{\uparrow}_t = 0 \qquad p^{b,\downarrow}_{i,t} = 0 \quad \text{if } \delta^{\downarrow}_t = 0$$

$$p^{b,\uparrow}_{i,t} + s^{\uparrow}_{i,t} \geq p^{r,\uparrow}_{i,t} \cdot \delta^{\uparrow}_t, \quad p^{b,\downarrow}_{i,t} + s^{\downarrow}_{i,t} \geq p^{r,\downarrow}_{i,t} \cdot \delta^{\downarrow}_t \quad \forall t \in T$$

we ensure that the activated up/down reserve $p^{b,\uparrow}_t, p^{b,\downarrow}_t$ can meet the reserve commitments $p^{r,\uparrow}_t, p^{r,\downarrow}_t$, if the TSO activates them (i.e., when $\delta^{\uparrow}_t = 1$ or $\delta^{\downarrow}_t = 1$). If activation is not fully met, the slack variables $s^{\uparrow}_t, s^{\downarrow}_t$ capture the amount of shortfall, which then means the aggregator gets a penalty

Aggregate Fleet-Level Constraints: Across all EVs, the total scheduled charging power at time $t \in T$ is:

$$p_t = \sum_{i\in I} p_{i,t}$$

$$p^{\min}_t = \sum_{i\in I} p^{\min}_{i,t}, \quad p^{\max}_t = \sum_{i\in I} p^{\max}_{i,t}$$

$$p^{\min}_t \leq p_t \leq p^{\max}_t$$

These bounds define the aggregator's safe range for scheduling and market participation.

Ramp Rate Constraints Let $R^{\uparrow}_i$ be the per-slot ramp limit.

$$|pr^{\uparrow}_{i,t} - pr^{\uparrow}_{i,t-1}| \leq R^{\uparrow}_t \tag{1}$$

$$|pr^{\downarrow}_{i,t} - pr^{\downarrow}_{i,t-1}| \leq R^{\downarrow}_t \tag{2}$$

DFO Energy If we use a DFO instead of FO we have to ensure that each DFO time unit has a polygon constraint on feasible power $p_{i,t}$, based on accumulated energy:

$$y^{(i,t)}_{\min}\left(\sum_{\pi=1}^{t-1} p_{i,\pi}\right) \leq p_{i,t} \leq y^{(i,t)}_{\max}\left(\sum_{\pi=1}^{t-1} p_{i,\pi}\right)$$

where $y^{(i,t)}_{\min}$ and $y^{(i,t)}_{\max}$ are bounds interpolated from the dependency polygon of DFO i at time t. $y^{(i,t)}_{\min}$ and $y^{(i,t)}_{\max}$ are themselves functions of the cumulative energy consumed up to time $t1$, i.e. $\sum_{\pi=1}^{t-1} p_{i,\pi}$.

4 Clustering, Aggregation and Scheduling

In this section, we will describe the two different alignments we have utilized for aggregation in this project, specifically start alignment [3] and flexibility alignment[7]. Flexibility alignment is our own contribution[7], which is described in section 4.2. These alignments are a core part of the aggregation phase, as it is in these algorithms that we decide how to align different FO objects, in regards to time and energy consumption.

4.1 Start alignment

Start alignment is very simple and intuitive to understand. In start alignment, we simply position all the incoming FOs that we are aggregating over, and position them all at the earliest possible start time for the individual FO. An intuitive illustration of start alignment can be seen in figure 5.

From the figure it is very clear to see, that we might

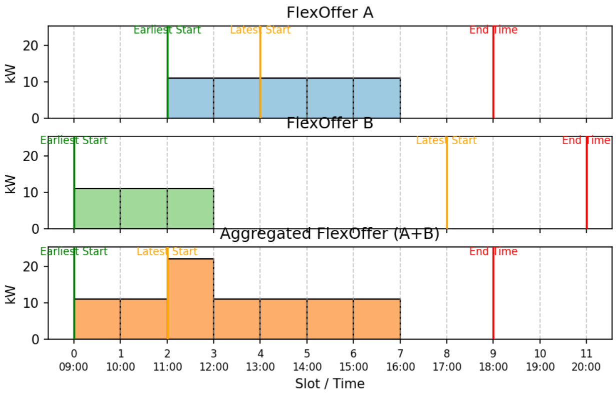

Figure 5: *Illustration of start alignment*

get aggregated FOs which are very large in terms of the duration of the aggregated FO. This means that we are forced to spend energy in a lot of hours, and it is likely that we might be forced to charge in hours, where the energy prices are high. Start alignment runs in linear time $O(n \cdot d)$ where n is the number of FOs and d is duration of the FO.

4.2 Flexibility alignment

With flexibility alignment, we aim to preserve as much flexibility of the individual FOs as possible. Production of electricity from RES cannot be moved in time, but flexible energy consumption can [7]. As such, preserving as much flexibility in the aggregation process allows for better utilization of energy from RES.

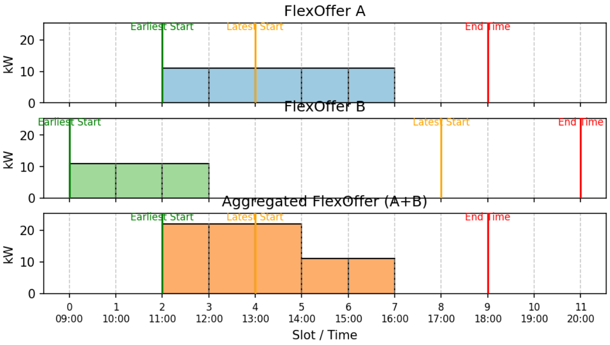

Figure 6: *Illustration of flexibility alignment*

We calculate the flexibility of a FO using the following equation:

$$\text{Flexibility}_f = \frac{1}{d} \sum_{t \in T} \left(e_t^{\max} - e_t^{\min} \right)$$

Where e_t^{max} corresponds to the maximum consumption of the timeslice in the FO at time t. Likewise, e_t^{min} corresponds to the minimum consumption of the same timeslice in the FO. T is the set of all timeslices in the FO, and $d = |T|$ represents the duration of the FO. This definition is a refinement of the energy flexibility metric introduced in [26]:

$$\text{EF} = \sum_{t=1}^{d} \left(e_t^{\max} - e_t^{\min} \right)$$

Our version uses a normalization by duration d, yielding a *per-time-unit* average, enabling comparison across FOs of variying lengths and allowing for benchmarking at a system level. For a given set F of FOs which we want to align using flexibility alignment, the alignment which maximizes the total flexibility is selected:

$$\text{Flexibility align}(F) = \text{Maximize} \sum_{f \in F} \text{Flexibility}_f$$

While flexibility alignment heuristics have been explored in previous papers, we extend them in a few ways [7][3]. First we introduce a tunable offset variable, enabling users to choose the number of candidate offset positions to be evaluated when aligning two FOs. This allows a trade-off between computational cost and flexibility. This makes our method adaptable for different fleet sizes and runtime requirements. Furthermore we also use a min-priority queue to always merge and keep track of the least flexible FlexOffers first. We simply call this fast flexibility alignment as show in algorithm 1.

Algorithm 1 Fast Flexibility Alignment

Input: List of FlexOffers $\mathcal{F} = [F_1, F_2, \ldots, F_n]$, number of candidate offsets k
Output: Aggregated FlexOffer F_A
Initialize a priority queue Q with $(\text{flexibility}(F_i), F_i)$ for all $F_i \in \mathcal{F}$
while $|Q| > 1$ **do**
  Pop the two least flexible offers F_a, F_b from Q based on est and lst
  $F_{\text{best}} \leftarrow$ None, minCost $\leftarrow \infty$
  for each offset δ in k candidate offsets **do**
    $F_{\text{temp}} \leftarrow \text{Merge}(F_a, F_b, \delta)$
    cost $\leftarrow \text{Flexibility}(F_{\text{temp}})$
    if cost $<$ minCost **then**
      minCost $\leftarrow$ cost, $F_{\text{best}} \leftarrow F_{\text{temp}}$
    end if
  end for
  Insert F_{best} into Q with updated flexibility score
end while
return Final aggregated FlexOffer from Q

Where the Flexibility metric measures the amount balance in the FlexOffer that we want to optimize. we simply measure the average power in each of the profile P, and divide by the duration of the FO. Here we try to stack as much amount flexible as possible in a few hours, while still retaining time flexibility allowing for shift.

$$\text{AbsBalance}(F) = \sum_{i=1}^{P} \left| \frac{p_i^{\min} + p_i^{\max}}{2} \right|$$

We merge flexoffer F_b on F_a with a given offset δ like this:

- Determine the combined length from the earliest start of the two to the latest end when one is shifted with δ
- Initialize a new profile with that length.
- Flexoffer $F_a's$ min and max power slices are added directly to the new profile.
- Flexoffer $F_b's$ power slices are then added with the shifted offset δ taken into account.
- The new EST of the AFO is the latest of both individual ESTs.
- The new LST is the earliest of the individual LSTs among the FOs
- The total energy constraints are summed from the two original offers.

4.3 Clustering

In order to group similar FOs, we experimented with a set of different clustering algorithms:

4.4 Bottom up agglomerative clustering

In this approach we are essentially building a tree-like structure. The FOs are sorted based on its structure, such as earliest start time and latest start time, and a tree is then build on top of this sorted list of FOs. We continue building up the tree, until we reach the number of clusters we desire. This is done by always merging, the most similar clusters.

4.5 k-means clustering

In the k-means algorithm, the goal is to separate X datapoints (Our FOs), into k clusters. A point belongs to the cluster, which cluster center is closest to the point. Whenever a point is added to a cluster, it moves the center of the cluster it was added to, which means that we need to run quite a few iterations of the algorithm, until we are sure that the cluster centers are not moving anymore [27].

We split the clustering from aggregation and allow different sets of clustering methods like, k-means and dbscan in addition to our implementation of bottom up clustering. All clustering are based on the same set of features that is [earliest start, latest start, end time, and average energy requirement in each slice]. Our implementation also supports dynamic clustering, where the number of clusters is selected based on quality criteria such as silhouette score. This yields an overall data-driven approach to clustering. To further improve runtime, we use parallel aggregation for FOs. Here we distribute clustering and alignment across multiple CPU cores.

4.6 Dynamic clustering

In order to improve clustering quality, we implement a dynamic clustering strategy. Here we iterate over a number of cluster counts $k \in [k_{min}, k_{max}]$. We evaluate the results of each clustering by its silhouette score, which measures how well separated the clusters are. The number k that yields the highest score is then the actual number of clusters used the final clustering (which will be final number of AFOs).

The silhouette score is defined as

$$s(i) = \frac{b(i) - a(i)}{max(a(i), b(i))}$$

where $a(i)$ is the mean intra-cluster distance between point i to its cluster center, where the distance is measured as the difference between features such as [est, lst, end, and energy balance]. Similarly $b(i)$ is the mean nearest cluster distance.

Algorithm 2 Dynamic Clustering

Input: List of FlexOffers $\mathcal{F} = [F_1, F_2, \ldots, F_n]$, Range of clusters count $k = [k_{min}, k_{max}]$
Output: List of k clusters

$k_{best} \leftarrow$ None
score_best $\leftarrow -\infty$
for each k in$[k_{min}, ..., k_{max}]$ **do**
  clusters, labels $\leftarrow$ cluster_offers(offers, k)
  score $\leftarrow$ eval_silhouette_score (clusters, labels)
  if score $>$ score_best **then**
    score_best $\leftarrow$ score
    $k_{best} \leftarrow$ k
  else
    break
  end if
end for
final_clusters, _ $\leftarrow$ cluster_offers(offers, best_k)
return Final Clusters to be aggregated

4.7 Scheduling

After aggregation produces a set of AFOs, they are input in the scheduling phase where they are bid and dispatched in the energy markets. Depending on a chosen strategy, we solve one of three linear programming problems. We assume perfect price forecast knowledge, to isolate the aggregation logic impact on costs from forecasting errors.

In all cases, the scheduling LP is built with AFOs, and not FOs. This is done to limit the number of variables and constraints, such that the scheduling runtime is negligible compared to the aggregation. Decision variables and parameters are seen in section 3.2.

- **Joint:** A single LP allocates energy for spot market, mFRR reserve and activation market simultaneously, maximizing profits.
- **Sequential:** First a single LP optimizes reserve capacity bids, and only then does a second LP allocate energy on the spot market energy, while ensuring that we meet the fixed reserve bids from the first LP pass.
- **Spot only:** A single LP program optimizes for spot energy only, minimizing the cost. This is used as a baseline for comparing the other modes.

5 Evaluation

In this section we evaluate the runtime performance and economic benefits in multiple markets of our aggregation and scheduling framework.

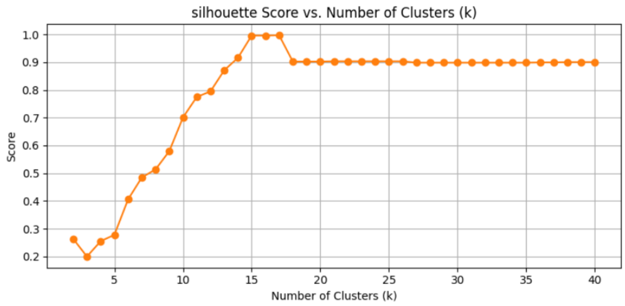

Figure 7: *Optimal cluster amount*

5.1 Method

To evaluate the performance of our EV scheduling framework, we simulate an electric vehicle (EV) fleet over a one-month period, using synthetic input data. For each day, we generate FOs or DFOs based on sampled EV behavior and then optimize their participation in the day-ahead spot and mFRR reserve markets, through clustering and aggregation and the different LP programs described in section 4.7.

Fleet Simulation and Offers Each EV in the fleet generates a flexible charging schedule using probabilistic sampling of:

- Arrival and departure times (sampled between 18:00 and 10:00) drawn from a log-normal distribution [28] [29].
- Arrival SoC, drawn from a uniform distribution [30] between 0.3 and 0.5.
- Target SoC bounds is set to a uniform distribution between 0.8 and 0.9

Each day we compare 3 market strategies,

- Joint: A single LP solves a co-optimization problem for both spot and mFRR simultanously.
- sequential: optimize for reserve market first, then spot-market.
- spot only: optimize only the spot costs in the day-ahead.

All 3 strategies are evaluated under two time resolutions (15 min and 60 min). We use real prices from Nordpool and Energinet CSV archives for spot, mFRR and balance activation prices in the DK1 area for the entire month of January of 2024.

Aggregation and scheduling We aggregated our EVs into 15 clusters based on the clustering results of our dynamic clustering shown on fig 7, which shows that 15 clusters yields the best score for well separated clusters for 500.000 EVs. We then aggregated each cluster to convert them into AFOs.

Baseline and Optimal Models Our **baseline** mode charges each EVs naively starting at the earliest start time, without reserve participation (only spot market) until their charging requirements are met.

The **theoretical optimal** strategy solves a full LP on non-aggregated DFOs participating in all markets with the same market rules.

Metrics We use the following metrics:

- **Total savings (%)**:

$$\text{Savings} = \frac{C_{\text{baseline}} - C_{\text{scheduled}}}{C_{\text{baseline}}} \tag{3}$$

where C_{baseline} is the cost under a naive spot-only baseline, and $C_{\text{scheduled}}$ is the optimized schedule cost.

- **Percent of theoretical optimal (%)**:

$$\text{\% of theo. opt} = \frac{C_{scheduled}}{C_{opt}} \tag{4}$$

This measures the economic gap between our aggregated strategy and an ideal MILP using non-aggregated DFOs. This metric is only computed for small fleet sizes ($N \leq 1000$) as it is computationally inefficient for larger fleet sizes. We therefore assume that the % of theoretical optimal we get is independent of the fleet size. It is important to note that here we still use perfect foresight, so that the difference is truly between non-aggregated and aggregated FOs and DFOs.

- **Runtime (seconds)**: Aggregation and scheduling times are measured.

Scenario Configuration We run all possible combinations of settings for each day, where each setting is a possible combination of the following variables. We do this for a month then average values for each day.

For each scenario and day:

- We simulate the fleet and generate FOs/DFOs.
- We cluster, align, and aggregate offers.
- We schedule them via LP (joint or sequential).
- We evaluate performance and savings in each market.

5.2 Results and discussion

Runtimes Figure 8 and 9 show the total different runtime from FlexOffer creation to scheduling of aggregated offer for 15 minutes granularity and 60 minute granularity respectively for different aggregation setting. At 15 minute granularity, the normal flexibility alignment grows steeply with fleet size while flexibility_fast, start, and DFO remain under ~80 seconds at 500 k EVs. Furthermore the agglomerative clustering is only slightly slower than k-means, however it is still respectable, since we can prove that agglomerative respects earliest and latest start time thresholds, whereas k-means cannot guarantee those thresholds are always maintained.

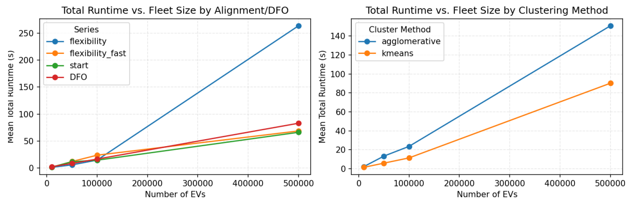

Figure 8: *Mean total runtime versus fleet size at 15 minute resolution. (a) Comparison of four alignment/DFO aggregation methods—classic flexibility (blue), fast flexibility (orange), start alignment (green), and DFO aggregation (red). (b) Comparison of two clustering algorithms—k-means (blue) and agglomerative (orange).*

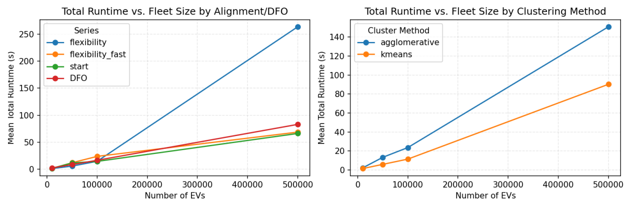

Figure 9: *Mean total runtime versus fleet size at 1 hour resolution.*

Figure 10 breaks down the execution time of each step (excluding the initial simulation of EVs), doing this we see the bottleneck as the fleet size increases. The key takeaway is that after about 50 000 EVs, the time spent aggregating FOs is much higher than both clustering and scheduling. This means that optimizing aggregation will yield the largest gains for large-fleets. It also means that for large fleets, more precise clustering could be a good idea since it is not the bottleneck.

5.3 Savings

Figure 11.a illustrates how alignment strategies affect cost savings compared to the naive baseline in a boxplot. Here we average over all market modes. The median saving for for flexibility alignment is about

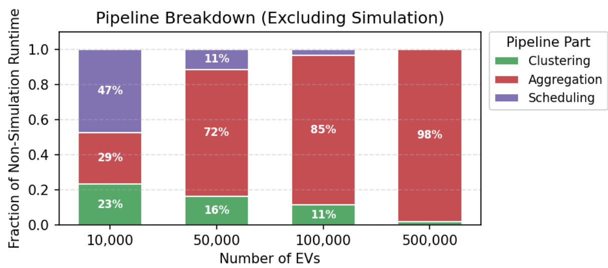

Figure 10: *Percentwise runtime of each part in the pipeline excluding simulation (drawing charging windows from distributions, loading prices, etc.*

MODE	Joint, sequential, spot
TIME_RES	15 min, 1 hour
NUM_EVS	10000, 100000, 500000
NUM_CLUSTERS	15
ALIGNMENT	Flexibility, Flexibility_fast , start
CLUSTER_METHOD	agglomerative, kmeans

Table 1: *We run all possible settings for every day in a month, then take the average for every day*

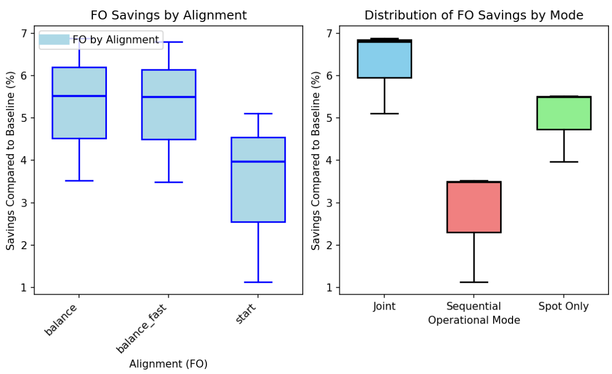

Figure 11: *a) Cost savings compared to a naive baseline. b) FO savings by mode*

5.6% DKK compared to the naive baseline. Furthermore we see that about 50% of all savings lie between 4.5 to 6 % indicated by the box. This captures market volatility as prices can change much from day to day. Choosing different alignments seem to have a slight affect on the overall flexibility on the market, but the main takeaways are that flexibility alignment is better than start alignment, and that flexibility_fast heuristic loses almost nothing compared to the slow flexibility heuristic. Here we have set the number of candidate offsets to be 5 for fast_flexibility. DFOs are not shown in the image as they do not use alignments, however when we use DFOs we can capture up to 40% savings in best case scenarios, showing that DFOs are much better at capturing and keeping flexibility when aggregated. Figure 11.b shows how much of a difference the reserve and activation market presents when we use the joint model compared to our sequential reserve first model, which is even worse than spot only. This shows that optimizing for reserve first, even with perfect knowledge of the activation prices is a bad strategy, since we mostly ignore better spot prices.

pct of theo.opt and Market Contributions Table 1 summarize the effectiveness of our aggregation and scheduler versus the actual theoretical limit we could attain by scheduling unaggregated DFOs in the joint solver in all markets. DFOs substantially outperforms the FO in all settings, and the joint optimization is generally better than sequential for profits. Lastly we also note that clustering only has a minor effect compared to the type.

Table 2 summarizes market contributions under both joint and sequential optimization. As expected we see that when we run reserve first in sequential, we gain much more potential activation revenue. The model sacrifices spot revenue to gain more flexibility for activation bids

Table 2: *Average % of Theoretical Maximum, grouped by Type, Mode, and Clustering Method (FO aggregation is flexibility_fast)*

Type	Mode	Cluster Method	Avg. % Optimal
DFO	joint	dbscan	91.02
DFO	joint	k-means	90.13
DFO	joint	agglomerative	92.14
DFO	sequential	dbscan	86.65
DFO	sequential	k-means	82.98
DFO	sequential	agglomerative	87.92
FO	joint	dbscan	73.90
FO	joint	k-means	70.73
FO	joint	agglomerative	73.88
FO	sequential	dbscan	70.15
FO	sequential	k-means	67.88
FO	sequential	agglomerative	70.91

Table 3: *Average Market Contributions (%) by Type and Mode*

Mode	Spot (%)	Reserve (%)	Activation (%)
joint	89.25	7.20	3.55
sequential	24.24	50.54	25.22

6 Conclusion

We have shown that smart aggregation of EVs flexibility via FOs makes it feasible for aggregators to participate in the danish spot, mFRR reserve, and balance activation market by making profitable scheduling decision in under 1 hour for 500.000 EVs and beyond for a 15 minute resolution. Specifically, by aligning individual EV charging windows into AFOs we showed through multiple simulations that a flexibility_fast alignment can be used to aggregate multiple FOs into 15 AFOs in less than a minute, and still be profitable in the 3 mentioned markets, boosting bids by up to 70% of theoretical max profits for FOs and up to 92.14% for DFOs.

The key takeaways is that fast_flexibility is usually better than normal flexibility alignment given that it runs much faster and keeps almost all of the flexibility. Furthermore if we aggregate down to 15 AFOs, then co-optimizing a joint market is better than a sequential market, since the savings are higher and the runtime is negligible. Lastly we observe that DFOs always achieves higher savings and only takes slightly longer to compute than FOs with flexibility_fast.

6.1 Future work

Renewable Uncertainty: Extend the co-optimization model to incorporate wind and solar forecast errors (e.g., via chance-constrained programming), enabling more aggressive yet safe bids in reserve markets.

Uncertain Flexoffers: incorporate Uncertain FlexOffers into clustering and aggregation and measure its potential savings compared to DFO and FO. [9]

Grid Constraints: Extend hierarchical aggregation such that that we respect local distribution limits in areas (by making individual areas their own clusters)), ensuring that AFOs remain feasible as EV penetration grows.

Scalability Trials: Make a pilot approach to validate that the runtime remains for low for 15 minute resolutions in real time scenarios with potential challenges (such as communication latency, data privacy, etc.).

6.2 Github

The complete code for this project can be found at `https://github.com/orgs/cs-25-dt-10-03/repositories`. All repositories belonging to the organization, is a part of this project.

References

[1] Luc Pelkmans. "Implementation of bioenergy in Denmark – 2024 update". In: (20 december, 2024).

[2] Christian Adelhardt. "Vehicle Grid Integration Reseach". In: (). URL: `https://energinet.dk/media/qmhnfcam/vehicle-grid-integration-research-rethinking-energy-international-alignment.pdf#:~:text=1,grid%20in%20time%20is%20impossible`.

[3] Laurynas Šikšnys et al. "Aggregating and Disaggregating Flexibility Objects". In: *IEEE Transactions on Knowledge and Data Engineering* 27.11 (2015), pp. 2893–2906. ISSN: 1041-4347. DOI: 10.1109/TKDE.2015.2445755.

[4] Kristian Sevdari, Peter Bach Andersen, and Mattia Marinelli. "Aggregation and Control of Electric Vehicles AC Charging for Grid Services Delivery". English. In: *IEEE Transactions on Smart Grid* 16.2 (2025), pp. 1523–1534. ISSN: 1949-3053. DOI: 10.1109/TSG.2024.3492391.

[5] Torben Bach Pedersen, Laurynas Šikšnys, and Bijay Neupane. "Modeling and Managing Energy Flexibility Using FlexOffers". In: *2018 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. 2018, pp. 1–7. DOI: 10.1109/SmartGridComm.2018.8587605.

[6] Faria. "Participation of an EV Aggregator in the Reserve Market through Chance-Constrained Optimization". In: *Energies* (2020).

[7] Raymond Kronborg, Mikkel Smedt, and Andreas Jensen. "Survey of flexoffer aggregation methods". In: *Last semester project* (2024).

[8] Laurynas Šikšnys and Torben Bach Pedersen. "Dependency-based FlexOffers: scalable management of flexible loads with dependencies". In: *Proceedings of the Seventh International Conference on Future Energy Systems*. e-Energy '16. Waterloo, Ontario, Canada: Association for Computing Machinery, 2016. ISBN: 9781450343930. DOI: 10.1145/2934328.2934339. URL: `https://doi.org/10.1145/2934328.2934339`.

[9] Fabio Lilliu et al. "Uncertain FlexOffers: a scalable, uncertainty-aware model for energy flexibility". English. In: *e-Energy 2023 - Proceedings of the 2023 14th ACM International Conference on Future Energy Systems*. e-Energy 2023 - Proceedings of the 2023 14th ACM International Conference on Future Energy Systems. Funding Information: This work was supported by the DomOS and FEVER projects, funded under

the Horizon 2020 programme, GAs 864537 and 894240 respectively, and Flexible Energy Denmark, funded by Innovation Fund Denmark (Case No. 8090-00069B). We thank Hamdi Ben Hamadou for proofreading. Publisher Copyright: © 2023 Owner/Author.; 14th ACM International Conference on Future Energy Systems, e-Energy 2023 ; Conference date: 20-06-2023 Through 23-06-2023. United States: Association for Computing Machinery (ACM), June 2023, pp. 30–41. DOI: 10.1145/3575813.3576873.

[10] EnergiNet. "ANCILLARY SERVICES TO BE DELIVERED IN DENMARK - TENDER CONDITION". In: (2024).

[11] Omid Sadeghian et al. "A comprehensive review on electric vehicles smart charging: Solutions, strategies, technologies, and challenges". In: *Journal of Energy Storage* 54 (2022), p. 105241. ISSN: 2352-152X. DOI: https://doi.org/10.1016/j.est.2022.105241. URL: https://www.sciencedirect.com/science/article/pii/S2352152X22012403.

[12] Miguel A. Ortega-Vazquez, François Bouffard, and Vera Silva. "Electric Vehicle Aggregator/System Operator Coordination for Charging Scheduling and Services Procurement". In: *IEEE Transactions on Power Systems* 28.2 (2013), pp. 1806–1815. DOI: 10.1109/TPWRS.2012.2221750.

[13] Stylianos I. Vagropoulos and Anastasios G. Bakirtzis. "Optimal Bidding Strategy for Electric Vehicle Aggregators in Electricity Markets". In: *IEEE Transactions on Power Systems* 28.4 (2013), pp. 4031–4041. DOI: 10.1109/TPWRS.2013.2274673.

[14] Peter A.V. Gade et al. "Load shifting versus manual frequency reserve: Which one is more appealing to thermostatically controlled loads in Denmark?" In: *Electric Power Systems Research* 232 (2024), p. 110364. ISSN: 0378-7796. DOI: https://doi.org/10.1016/j.epsr.2024.110364. URL: https://www.sciencedirect.com/science/article/pii/S0378779624002529.

[15] Shaofeng Lu et al. "Stochastic bidding strategy of electric vehicles and energy storage systems in uncertain reserve market". In: *IET Renewable Power Generation* 14 (Feb. 2021). DOI: 10.1049/iet-rpg.2020.0121.

[16] Han Wang et al. "Optimal operating regime of an electric vehicle aggregator considering reserve provision". In: *Energy Reports* 8 (2022), pp. 353–362. URL: https://www.sciencedirect.com/science/article/pii/S2352484722004103.

[17] António Sérgio Faria et al. "Participation of an EV Aggregator in the Reserve Market through Chance-Constrained Optimization". In: *Energies* 13.16 (2020). ISSN: 1996-1073. DOI: 10.3390/en13164071. URL: https://www.mdpi.com/1996-1073/13/16/4071.

[18] Dante F. Recalde Melo et al. "Robust Electric Vehicle Aggregation for Ancillary Service Provision Considering Battery Aging". In: *IEEE Transactions on Smart Grid* 9.3 (2018), pp. 1728–1738.

[19] Homa Rashidi et al. "A Stochastic Bi-Level Scheduling Approach for the Participation of EV Aggregators in Competitive Electricity Markets". In: *Applied Sciences* 7 (Oct. 2017), p. 1100.

[20] Christoforos Menos-Aikateriniadis et al. "Towards a New Energy Reserves Market: The Role of EV Aggregators". In: *IEEE Access* 13 (2025), pp. 52176–52194. DOI: 10.1109/ACCESS.2025.3552838.

[21] S. de la Torre, J.A. Aguado, and E. Sauma. "Optimal scheduling of ancillary services provided by an electric vehicle aggregator". In: *Energy* 265 (2023), p. 126147. ISSN: 0360-5442. DOI: https://doi.org/10.1016/j.energy.2022.126147. URL: https://www.sciencedirect.com/science/article/pii/S036054422203033X.

[22] Feiye Zhang, Qingyu Yang, and Donghe Li. "A deep reinforcement learning-based bidding strategy for participants in a peer-to-peer energy trading scenario". In: *Frontiers in Energy Research* 10 (Jan. 2023), p. 1017438. DOI: 10.3389/fenrg.2022.1017438.

[23] Abdullah Zishan, Moosa Haji, and Omid Ardakanian. "Adaptive Control of Plug-in Electric Vehicle Charging with Reinforcement Learning". In: June 2020. DOI: 10.1145/3396851.3397706.

[24] Kristian Sevdari, Peter Bach Andersen, and Mattia Marinelli. "Aggregation and Control of Electric Vehicles AC Charging for Grid Services Delivery". In: *IEEE Transactions on Smart Grid* 16.2 (2025), pp. 1523–1534. DOI: 10.1109/TSG.2024.3492391.

[25] Emmanouil Valsomatzis, Torben Bach Pedersen, and Alberto Abello. *Day-ahead Trading of Aggregated Energy Flexibility - Full Version*. 2018. arXiv: 1805.02301 [cs.CE]. URL: https://arxiv.org/abs/1805.02301.

[26] Fabio Lilliu, Jonas Brusokas, and Torben Bach Pedersen. "Device-independent metrics for explicit, a priori measurement of energy flexibility". English. In: *e-Energy '24: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 15th ACM International Conference on Future Energy Systems, e-Energy 2024 ; Conference date: 04-06-2024 Through 07-06-2024. United States: Association for Computing Machinery (ACM), June 2024, pp. 268–278. DOI: 10.1145/3632775.3661965. URL: https://energy.acm.org/conferences/eenergy/2024/.

[27] Xin Jin and Jiawei Han. "K-Means Clustering". In: *Encyclopedia of Machine Learning*. Ed. by Claude Sammut and Geoffrey I. Webb. Boston, MA: Springer US, 2010, pp. 563–564. ISBN: 978-0-387-30164-8. DOI: 10.1007/978-0-387-30164-8_425. URL: https://doi.org/10.1007/978-0-387-30164-8_425.

[28] "Lognormal Distribution". In: *The Concise Encyclopedia of Statistics*. New York, NY: Springer New York, 2008, pp. 321–322. ISBN: 978-0-387-32833-1. DOI: 10.1007/978-0-387-32833-1_238. URL: https://doi.org/10.1007/978-0-387-32833-1_238.

[29] Andres Arias, Mauricio Granada Echeverri, and Carlos Castro. "Optimal probabilistic charging of electric vehicles in distribution systems". In: *IET Electrical Systems in Transportation* 7 (Mar. 2017). DOI: 10.1049/iet-est.2016.0072.

[30] Ki-Beom Lee et al. "Deep Reinforcement Learning Based Optimal Route and Charging Station Selection". In: *Energies* 13.23 (2020). ISSN: 1996-1073. DOI: 10.3390/en13236255. URL: https://www.mdpi.com/1996-1073/13/23/6255.