A Rolling-Horizon Optimal Selection Algorithm for Disaggregating Flexibility in Heterogeneous Asset Portfolio

Janez Gregor Golja
Laboratory of Energy Policy,
Faculty of electrical engineering,
University of Ljubljana
Ljubljana, Slovenia
janezgregor.golja@fe.uni-lj.si

Dr. Tomi Medved
Laboratory of Energy Policy,
Faculty of electrical engineering,
University of Ljubljana
Ljubljana, Slovenia
tomi.medved@fe.uni-lj.si

Abstract - The energy system is undergoing rapid transformation driven by widespread electrification and increased integration of distributed generation. In this context, flexibility is quickly gaining importance as a means for ensuring secure and cost-effective grid operation. Aggregators play a central role in enabling market participation of small-scale assets by pooling distributed resources and offering their combined flexibility into markets that individual assets could not access on their own. Responding to a flexibility request requires translating a market-level signal into asset-level setpoints that respect technical constraints under uncertainty, while minimizing activation costs. Within the OPENTUNITY project, we address this challenge with a rolling-horizon genetic algorithm that identifies the lowest cost feasible combination of assets to deliver the requested flexibility. The algorithm integrates a unified cost modelling framework capturing energy price impacts, asset degradation, user comfort, and forecast uncertainty, combined with real-time feasibility checks of asset operational limits. Results across multiple portfolio configurations confirm the practical relevance of the approach for realistic aggregator use cases in developing flexibility markets.

***Index Terms* -** Flexibility, Aggregators, Disaggregation, Genetic algorithms

I. INTRODUCTION

The ongoing energy transition, characterized by rapid electrification of transport and heating and increasing penetration of distributed generation, is fundamentally changing the structure of electricity demand and supply. This shift introduces new challenges for grid operators, who must balance a system with growing shares of intermittent and distributed resources. Flexibility, meaning the ability of assets to adjust their consumption or generation in response to external signals, is increasingly recognized as a key tool for maintaining this balance efficiently and cost-effectively, reducing the need for costly infrastructure upgrades or conventional reserve capacity [1]. Emerging flexibility markets, at both transmission and distribution level, are creating new monetization opportunities for owners of small-scale assets such as electric vehicle (EV) charging stations, heat pumps (HPs), photovoltaic (PV) systems, and battery energy storage systems (BESS). However, individual assets are typically too small to meet minimum bid sizes or independently fulfill the technical requirements of these markets. There are also legal and contractual barriers to consider. Aggregators address all the barriers by pooling large numbers of distributed assets and jointly offering their aggregated flexibility, effectively acting as intermediaries between asset owners and system operators. The aggregator's operational workflow spans several stages, from portfolio onboarding and forecasting through bid formation to real-time activation and settlement. The most operationally demanding of these is the activation stage: upon receiving a flexibility request, the aggregator must disaggregate the market-level signal into per-asset setpoints that collectively deliver the required amount of flexibility. This disaggregation must respect technical constraints, preserve user comfort, and minimize activation costs while managing forecast uncertainty. As portfolio size and asset diversity grow, the problem becomes increasingly combinatorial and difficult to solve in real time with classical optimization methods.

II. RELATED WORK

Coordinating distributed flexibility resources for market participation has been studied from several angles. Frameworks for representing and aggregating the flexibility of heterogeneous assets (such as zonotopic or polyhedral formulations) provide rigorous ways to characterize the feasible flexibility set of a portfolio and derive disaggregation mappings from portfolio-level decisions to device-level setpoints [2]. Multi-objective disaggregation approaches further extend the latter by distributing flexibility across assets while balancing competing objectives such as cost, comfort, and reliability [3]. Reviews of flexibility products and local markets highlight the growing role of aggregators as intermediaries and the practical challenges of coordinating many small assets under evolving regulatory frameworks [4], [5], [6]. Optimization methods for distributed energy resources (DER) range from exact formulations to metaheuristics. Mixed-integer and convex programming approaches can provide optimality guarantees but scale poorly when asset types are heterogeneous, constraints are nonlinear, or the model must be updated as new devices are added. Stochastic and robust methods explicitly model forecast uncertainty and offer recourse at the cost of increased computational complexity [7]. In contexts where the problem is non-convex or involves mixed decision spaces, metaheuristic methods such as genetic algorithms (GAs) offer greater modelling flexibility and can handle heterogeneous constraints naturally, at the cost of losing formal optimality guarantees. Rolling horizon formulations address the need to balance foresight with real-time responsiveness by repeatedly solving a short-horizon subproblem as new information becomes available [8], [9]. This structure has been applied to microgrid scheduling, demand response, and stochastic energy management, and is well-suited to flexibility disaggregation where asset states evolve continuously and forecasts are updated at each step. However, existing rolling horizon approaches typically assume homogeneous asset types, fixed cost structures, or tractable problem formulations, and do not address the combined challenge of heterogeneous assets, dynamic internal cost modelling, and comfort-aware feasibility enforcement that arises in realistic aggregator portfolios. This paper addresses the identified gap by proposing a rolling-horizon genetic algorithm with a unified marginal cost framework that handles all of the highlighted aspects jointly.

III. METHODOLOGY

A. Problem formulatiom

An aggregator participating in flexibility markets coordinates a portfolio of heterogeneous DERs, each with different technical capabilities, operational constraints, and cost structures. When a flexibility request arrives, specifying a required power profile over a delivery horizon, the aggregator must determine which assets to activate, at what power, and in which time slots, in a way that their combined output matches the requested profile at minimum cost. This is the flexibility disaggregation problem. The decision variables are the per-asset dispatch flexibility magnitudes $x_{i,t}[kW]$ for each asset i and timeslot t, where the optimization horizon is divided into discrete intervals of length Δt. The primary objective is to minimize the total cost of delivering flexibility over the horizon W:

$$\min_{X} \sum_{t \in W} \left(\sum_{i \in I} c_{i,t} \cdot x_{i,t} \cdot \Delta t + C^{pen}(s_t) \cdot \Delta t \right)$$

Where $c_{i,t}$ is the marginal activation cost of asset i at slot t, and $C^{pen}(s_t)$ penalizes any shortfall s_t in meeting the requested activation. The solution must respect per-asset feasibility constraints (such as instantaneous power bounds, state-of-charge for BESS, availability windows and temperature limits), as well as portfolio-level delivery requirements. Several factors make this problem practically challenging. First, the portfolio is heterogeneous: assets differ in their flexibility direction and in how their internal state evolves over time — battery state-of-charge (SoC) and HVAC thermal dynamics both affect dispatch decisions across slots in ways that impact feasibility in subsequent intervals. Second, the combinatorial nature of the problem grows rapidly with portfolio size, making it non-trivial to solve in real time with classical methods. Third, activation costs are not directly observable and must be modelled internally by the aggregator to capture the true economic impact of each dispatch decision on energy costs, asset health, and user comfort.

B. Optimal selection algorithm

To address the flexibility disaggregation problem, we propose a rolling-horizon controller that embeds a GA as its core optimizer. At each decision step, the controller receives the current internal states of all assets, the flexibility request profile over a short look-ahead window, energy price forecasts, and per-asset cost parameters. It outputs a dispatch vector assigning a flexibility magnitude to each asset for the current slot, updates all asset states, and then repeats for the next slot. Rather than optimizing over the full delivery horizon at once, the controller operates on a rolling window of fixed length W, as illustrated in Fig. 1. At each step t_1, the algorithm considers the flexibility request and the feasibility caps of all assets over slots $t_1$through t_W, and searches for the least-cost dispatch plan across that window. Once the best plan is found, only the dispatch for the first slot $t_1$is executed. Asset states are then updated, the window shifts forward by one slot to t_2, and the process repeats with refreshed inputs. This receding horizon structure ensures that the controller always acts on the most current asset conditions and price information, while the look-ahead window provides enough foresight to avoid locally greedy decisions that would be costly or infeasible in subsequent slots.

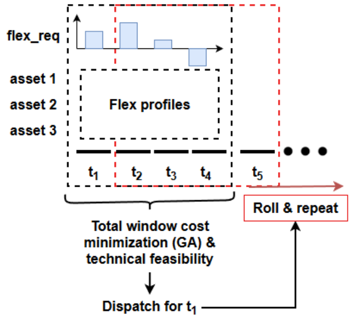

Figure 1. Rolling horizon illustration

Before the optimizer runs, per-slot feasibility caps $\bar{p}_{i,t}$ are computed for each asset, bounding the search space to activations each asset can physically deliver given its current state. Within each window, the search space consists of all feasible dispatch matrices $X$of size $n_{assets} \times W$, where $x_{i,t}$is the flexibility magnitude of asset $i$in slot t. We employ a GA for this search because the combined space of heterogeneous assets, inter-temporal constraints, and non-convex cost structures makes gradient-based methods impractical. As shown in Fig. 2, the GA maintains a population of candidate dispatch matrices and evolves them over successive generations.

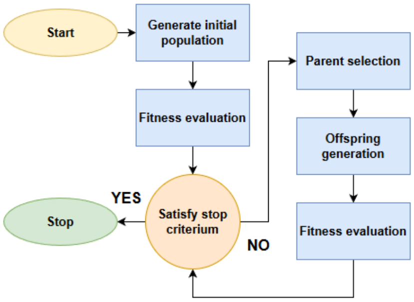

Figure 2. GA workflow

Each candidate is evaluated according to the fitness function:

$$F(X) = \sum_{t\in W}\sum_{i\in I} c_{i,t}\cdot x_{i,t}\cdot \Delta t + \sum_{t\in W} C^{pen}(s_t)\cdot \Delta t$$

where $c_{i,t}$is the per-asset marginal cost defined in Section III-C and $C^{pen}(s_t)$ penalizes any shortfall s_t relative to the requested profile. Fitter individuals are selected via tournament selection and recombined through column crossover, exchanging coherent time blocks between parent plans, followed by random perturbations of set-points. Every offspring is passed through a repair step that clips activations to the feasibility caps and rescales the plan to meet the requested total where possible. This ensures the population remains feasible throughout. Once the generation budget is exhausted, the best plan's first column is dispatched and the rolling horizon advances.

C. Cost and asset modelling

The controller ranks using marginal cost signals $c_{i,t}$ that capture economic impact, comfort, delivery reliability, and asset health — analogous to the merit order principle in wholesale markets. The total marginal cost has five components:

- Energy costs capture the net effect of shifting energy in time relative to the baseline. When an asset is curtailed, consumption or generation is reduced now but must be compensated later; when it is boosted, energy is consumed now to avoid or defer future use. In both cases, the cost depends on the difference between the current price and the expected future price at which the energy will be replaced or recovered. We estimate this future reference price v_{ref} by taking a low or high percentile of the forecasted day-ahead price curve over a look-ahead window of N slots:
  $$\Delta\lambda_{i,t} = \lambda_t - v_{ref}, v_{ref} = percentile(\lambda_{t+1}, \dots, \lambda_{t+N})$$
  For upward flexibility the term is positive when current prices exceed future replacement cost; for downward it reverses, with efficiency losses η applied where relevant.
- Opportunity costs capture value that is lost when an asset is activated. The most direct example is PV curtailment: each kWh of curtailed generation represents lost export revenue at the feed-in tariff λ_{FIT}. Unlike the exchange energy, this cost does not depend on the market prices and is a fixed cost, which makes PV curtailment consistently expensive relative to other downward flexibility options.
- Degradation cost applies to storage assets and reflects the wear caused by each charge-discharge cycle, modelled as a fixed cost per kWh c_{deg}, which is derived from the battery replacement cost and expected cycle life.
- Comfort cost is used for Heating, Ventilation and Air-Conditioning (HVAC) loads and penalizes activations that push indoor temperature toward the boundary of the comfort band $[T_{min}, T_{max}]$. We define a normalized deviation from the band midpoint T_{mid} as:
  $$z_t = \frac{T_t - T_{mid}}{w}, \qquad w = \frac{T_{max} - T_{min}}{2}$$
  The comfort cost scales linearly with $|z_t|$, becoming a penalty when the activation pushes temperature away from T_{mid} and a reward when it moves temperature back toward comfort:
  $$K_{i,t} = \kappa \cdot |z_t| \cdot effect\_sign$$
  Where κ is a weight in €/kWh.
- Uncertainty cost reflects the risk of non-delivery associated with assets whose baseline forecasts are less reliable:
  $$R_{i,t} = \gamma_i \cdot \sigma_i$$
  Where σ_i is the historical baseline forecast error and γ_i is the cost scaling factor in €/kWh. Assets with larger forecast errors receive a higher effective marginal cost, steering the optimizer toward more predictable resources.

The portfolio in our simulations includes four representative asset types, each with distinct technical capabilities and cost components:

- EV charging locations provide upward flexibility only, by curtailing active charging relative to the baseline demand profile. The per-slot cap equals the forecasted baseline consumption, falling to zero when the site is unavailable.
- BESS: supports both directions, with per-slot caps determined by available energy within SoC bounds and maximum charge/discharge power.
- PV systems provide downward flexibility through curtailment only, with the per-slot cap equal to the forecasted generation baseline.
- HVAC units support both directions within a comfort-constrained envelope: upward flexibility through consumption curtailment and downward through pre-heating or pre-cooling. Caps are computed slot by slot using a first-order thermal model, excluding any activation that would violate the comfort band at the next step.

The cost components active for each asset type and the specific parameter values used in the simulations are summarized in Table I in Section III-D.

D. Scenario definition

The simulations use real measurement data shared by project partners, sampled at 15-minute resolution over a full day (96 slots). PV generation is represented by a single profile with a peak of approximately 11 kW, typical of a residential rooftop installation. Two HVAC profiles are used: a larger commercial-scale unit and a smaller household heat pump, with outdoor temperature obtained from Open-Meteo as an input to the thermal model. Three EV charging location profiles are drawn from the considered car-sharing fleet, each representing a site with multiple chargers and a relatively stable demand pattern. The BESS state is fully determined by the algorithm's dispatch decisions. Day-ahead electricity prices for the same date serve as the price input to the cost model.

Flexibility request profiles are constructed by sampling a fraction of the portfolio envelope and shaping it into activation blocks with varying duration, magnitude, and direction, always remaining feasible by design. Two portfolio configurations are evaluated. The baseline scenario consists of one asset of each type, allowing the algorithm's behavior to be clearly interpreted in terms of individual asset contributions. The large portfolio comprises 47 assets: 10 prosumer households each with PV, HVAC, and a household BESS, two larger buildings with scaled PV and a larger HVAC unit, one large BESS, and three EV charging locations.

Table I shows the key cost parameters used in the simulations.

Table I. asset cost parameters with values

Parameter (€/kWh)	EV	BESS	PV	HVAC
Energy (DA prices)	✓	✓		✓
Opportunity (FIT)			0.06	
Degradation c_{deg}		0.06		
Comfort κ				0.08
Uncertainty σ	0.30		0.05	0.12

IV. RESULTS

A. Baseline case (small portfolio)

Table II summarizes the overall delivery performance of the baseline scenario.

Table II. Baseline scenario: overall delivery performance

E_{total} (kWh)	E_{loss} (kWh)	C_{total} (€)	C_{avg} (€/kWh)
U: 29.1	0	2.791	0.0291
D: 28.18			

The distribution of flexibility across asset types is shown in Table III and visualized in Fig. 3, which shows the signed dispatch aggregated by asset type over the full day, with positive values indicating upward and negative values downward activations.

Table III. Baseline scenario: flexibility contribution by asset type

	PV	HP	BESS	EV
$\%_{up}$	0	6.58	55.60	36.82
$\%_{down}$	51.13	7.64	41.23	0

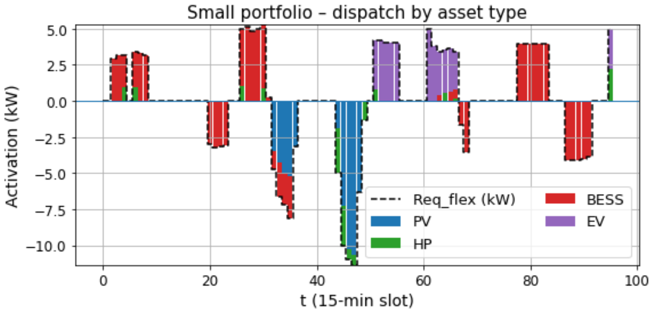

Figure 3. Baseline scenario: dispatch by asset type

In the upward direction, BESS provides the majority of flexibility at 56%, followed by EV at 37%, with HP contributing a modest 7%. PV provides no upward flexibility by design. In the downward direction, PV dominates at 51% during the central part of the day, with BESS covering 41% and HP the remaining 8%. To understand why the algorithm distributes flexibility in this way, Fig. 4 shows the asset baselines and realized flexibility caps alongside BESS SoC and HP indoor temperature trajectories; Fig. 5 shows the corresponding realized activation costs by asset and direction.

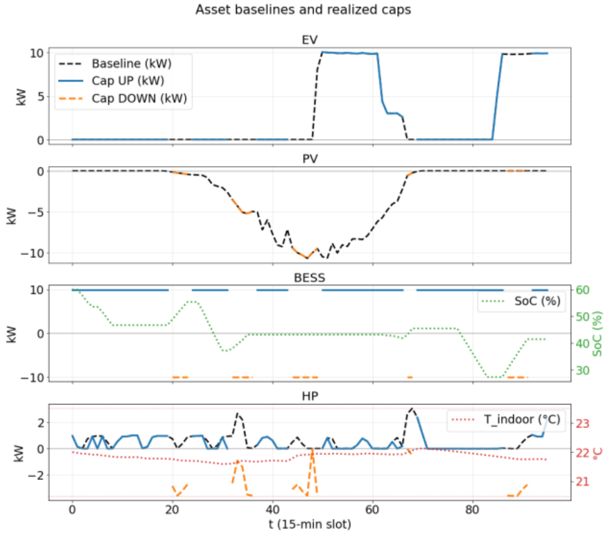

Figure 4. Baseline scenario: asset baselines, realized flexibility caps, SoC and indoor temperature

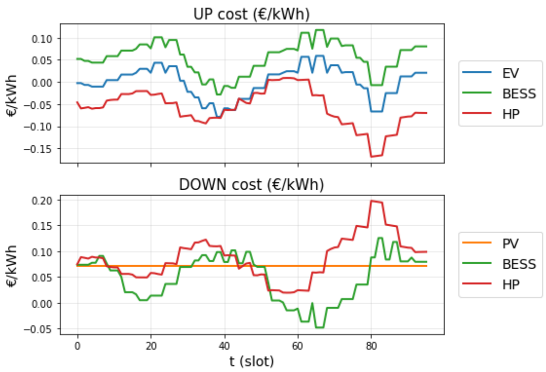

Figure 5. Baseline scenario: realized activation costs by asset and direction

Several observations explain the dispatch pattern. For upward flexibility, EV charging demand is absent before slot 40, so

BESS covers the earlier requests where EV caps are zero. Once EV demand appears, it becomes the preferred upward resource given its lower cost. HP contributes modestly throughout, constrained by its low operating power (its baseline rarely exceeds 2 kW) rather than by cost.

For downward flexibility, PV dominates between slots 30 and 60 where its cost is fixed at the feed-in tariff, making it consistently the cheapest option. Once PV drops to zero, BESS takes over, with HP filling a small share of the third downward block where BESS cost rises.

BESS SoC and HP indoor temperature remain within their defined bounds throughout, confirming correct constraint enforcement.

B. *Large portfolio case*

Table IV summarizes the overall delivery performance of the large portfolio scenario. The total requested flexibility is substantially larger than in the baseline reflecting the larger request profile constructed from the expanded portfolio envelope.

Table IV. Large portfolio scenario: overall delivery performance

E_{total} (kWh)	E_{loss} (kWh)	C_{total} (€)	C_{avg} (€/kWh)
U: 233.96	0	13.0898	0.1364
D: 333.39			

The flexibility contributions by asset type are shown in Table V and Fig. 6, which shows the signed dispatch aggregated by type over the full day.

Table V. Large portfolio scenario: flexibility contribution by asset type

	PV	HP	BESS	EV
$\%_{up}$	0	16.97	51.01	32.82
$\%_{down}$	27.20	16.28	56.52	0

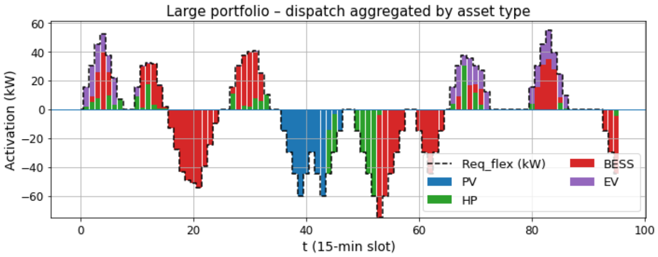

Figure 6. Large portfolio scenario: dispatch aggregated by asset type

BESS remains dominant in both directions at 51% upward and 57% downward, closely mirroring the baseline and reflecting its unconditional availability and bidirectional capability. HP contribution increases substantially from 7% to 17% upward and 8% to 16% downward, driven by the twelve HP units in the large portfolio including two powerful commercial units, providing meaningful flexibility across a wider range of slots. Three EV locations operating on different schedules spread upward flexibility more evenly across the day than the single baseline location, maintaining a similar 33% overall share. PV downward contribution drops from 51% to 27% despite greater total installed capacity, as the proportionally larger downward request extends well beyond daylight hours, shifting the remaining load onto BESS and HP.

V. CONCLUSION

This paper presented a rolling-horizon GA for optimal flexibility disaggregation across heterogeneous asset portfolios. The algorithm combines a unified marginal cost framework with asset-specific feasibility models that enforce technical and comfort constraints at every step. Evaluated on a baseline four-asset portfolio and a realistic 47-asset configuration, the algorithm achieves full delivery of the requested flexibility in both cases, with asset prioritization consistent with the defined cost structures and availability constraints.

Several directions remain open for future work. First, the current formulation assumes asset-level metering; extending it to incorporate portfolio-level smart meter feedback for real-time setpoint correction would make it directly applicable to local flexibility markets. Second, the uncertainty cost term currently uses a static forecast error estimate per asset type. Updating this dynamically after each activation would create a self-tuning weighting that naturally steers dispatch toward historically reliable assets. Third, a systematic sensitivity analysis of the cost parameters would help aggregators understand how to calibrate the model for their specific portfolio and market context.

ACKNOWLEDGMENT

The solutions presented in this work have been developed within the research activities of the OPENTUNITY project (Grant Agreement No. 101096333) and the STREAM project (Grant Agreement No. 101075654), both funded by the European Union under the Horizon Europe research and innovation programme.

References

[1] ENTSO-E, "System Flexibility Needs for the Energy Transition," ENTSO-E AISBL, Brussels, Belgium, 2024.

[2] F. Müller, J. Szabó, O. Sundström and J. Lygeros, "Aggregation and Disaggregation of Energetic Flexibility From Distributed Energy Resources," *IEEE Transactions on Smart Grid,* vol. 10, no. 2, pp. 1205-1214, 2017.

[3] D. Danner, R. Huwa and H. Meer, "Multi-Objective Flexibility Disaggregation to Distributed Energy Management Systems," *ACM SIGEnergy Energy Informatics Review,* vol. 2, 2022.

[4] J. Villar, R. Bessa and M. Matos, "Flexibility products and markets: Literature review," *Electric Power Systems Research,* vol. 154, pp. 329-340, 2018.

[5] J. Xiaolong , W. Qiuwei and J. Hongjie , "Local flexibility markets: Literature review on concepts, models and clearing methods," *Applied Energy,* vol. 261, p. 114387, 2020.

[6] D. Badanjak and H. Pandžić, "Distribution-Level Flexibility Markets — A Review of Trends, Research Projects, Key Stakeholders and Open Questions," *Energies,* vol. 14, no. 20, 2021.

[7] S. Vinothine and N. W. A. Lidula , "Microgrid Energy Management and Methods for Handling Forecast Uncertainty," *Energies,* vol. 15, no. 22, p. 8525, 2022.

[8] D. Aguilar, J. Quiñones, L. R. Pineda, J. Ostanek and L. Castillo, "Optimal scheduling of renewable energy microgrids: A robust multi-objective approach with machine learning-based probabilistic forecasting," *Applied Energy,* vol. 269, p. 123548, 2024.

[9] J. Silvente, G. Kopanos, V. Dua and L. Papageorgiou, "A rolling horizon approach for optimal management of microgrids under stochastic uncertainty," *Chemical Engineering Research and Design,* vol. 131, pp. 293-317, 2017.