Master controller concept for power flexible energy systems

Thomas Egsgaard Kallesen[1] Søren Waagø Christiansen[1] Rene Just Nielsen[1]

[1]Added Values P/S, Denmark, {tep, swc, rjn}@AddedValues.eu

Abstract

A generic model of a master controller concept is presented to demonstrate the advantages of aggregating multiple flexible electrical power units. The master controller consists of various submodels with different features. A Power Balancing Controller ensures that the operation does not cause any imbalance with the traded power. A waterfall method that ensures that any ancillary service activations are handled using the appropriate units to obtain cheapest operation. Simulations of the concept include examples of step responses on automatic Frequency Restoration Reserve (aFRR) activations, as well as actual aFRR data from the Danish TSO. The results suggest that if the system consists of slow units, the aggregation of units can increase the capacity offered to faster types of ancillary services. Additionally, by allowing certain units to shut down automatically more capacity can be offered as the units are no longer bound by minimum loads.

Keywords: master controller, district heating, energy systems, ancillary services

1 Introduction

The expected continued expansion of and reliance on non-dispatchable renewable energy sources in the production of electricity has led to greater volatility in the power markets (International Energy Agency 2020). That is, production does not necessarily follow consumption as fewer dispatchable energy sources are available in electricity grids. In order to mitigate this increased volatility, weather forecasting has become an important tool used to predict the production capacity, which is used to set the prices on the day-ahead market.

However, uncertainties in weather forecasts will always cause deviations between planned production and consumption. Here, the importance of ancillary services comes in. They are standardized across the connected grids to ensure that enough capacity is connected to compensate for the uncertainties with the increasing share of renewable energy. The forecast deviation from actual production will only increase in the future. Hence, there is a need for more flexible grid-connected capacity that can participate in ancillary services (Energinet 2024b).

Multiple ancillary services are available on the Danish DK1 grid, which is part of the European Network of Transmission System Operators for Electricity – ENTSO-E (ENTSO-E 2018). DK1 is chosen as the reference in this paper, as it is the market of the authors' workplace. The ancillary services (AS) are distinguished by their response time requirements and the energy delivery that are inversely correlated (Energinet 2025b):

- **Frequency Containment Reserve (FCR)** requires response time from zero to activated load in 30 seconds, but is typically quickly returned to baseline load again resulting in average activations as ratio of capacity of 0.05 %. The load is proportional to grid frequency deviation.
- **Automatic Frequency Restoration Reserve (aFRR)** requires a response time from zero to activated load in 5 minutes. It has an average activation ratio below 12 % compared to the reserved capacity. The transmission system operator (TSO) uses a closed-loop controller to suppress the power imbalance, which makes it automatic, as the name implies.
- **Manual Frequency Restoration Reserve (mFRR)** requires a response time from zero to activated load in 12.5 minutes—typically through the balance responsible party (BRP).

These ancillary services can be performed both by electricity *consuming* or *producing* units as long as they can adjust the power at the required rate. There is even an economic incentive to participate, but this is not covered in this paper. However, often enough, electric units do not comply with any of the above requirements, and certainly not all of them (Danish Energy Agency 2025). Besides, many units have some minimum power limit they can operate at. Thus, some unexploited capacity is available in many power or power-to-heat (PtH) district heating plants that could be optimized with respect to utilization and speed. This can be done by combining the attributes of different units in the plant. For example, a fast but expensive unit can temporarily help a slower but cheaper unit increase its speed so that it can participate in faster ancillary services, such as aFRR. Furthermore, the power consumption or production in a plant can be distributed, so it allows units to shut down and start up dynamically. This gives additional regulation capacity to the plant while maintaining a continuous trajectory of the power.

Most of the literature on the aggregation of grid-connected units focusses on coordinating large populations of small household devices to provide ancillary services. Within this field, (Iria, Soares, and Matos 2019) addresses the optimisation of day-ahead market bidding for aggregated resources. In terms of real-time control approaches, (Meyn et al. 2015) investigates the aggregation of large load populations using optimal control techniques, while (Bella et al. 2021) applies mixed-integer linear programming to the problem of load distribution. Similarly, (Iria and Soares 2019) introduces a fast model predictive control method, though it is still computationally heavy. A common feature of these studies is their focus on aggregating large numbers of small devices or prosumers to provide ancillary services. In contrast, (Tan and Zhang 2017) proposes a state machine approach for the combined control of a battery and wind power plant; however, this method is highly specific to that particular application.

The present work instead targets industrial-scale plants and aims to develop a general-purpose algorithm suitable for real-time control systems. The proposed approach emphasizes low computational requirements and operational transparency, ensuring that it remains accessible and comprehensible to plant operators, thus facilitating troubleshooting. By aggregating several power producing or consuming units with different technical characteristics, the strategy coordinates their setpoints to gain larger overall regulation capacity and faster response times. This is enabled by utilizing the fast response times of the more expensive units to aid the cheaper but slower units.
The final result is a structured Modelica implementation of the control concept, exemplified on a test system combining a heat pump, an electric boiler, and a battery. This implementation demonstrates both the internal architecture of the controller and the methodology applied to develop and validate it.

2 Master controller — working principle

The so-called *master controller* (MC) connects the electrical consumption or production units in the plant portfolio with the power grid, as illustrated in Figure 1.

The balance responsible party (BRP) is a third-party company responsible for power bidding and trading on behalf of the plant owner and for handling ancillary service requests from the transmission system operator (TSO). The BRP distributes AS activation signals from the transmission system operator directly to the master controller of its customers. A load scheduling system generates optimized base load setpoints for the individual plants, based on day-ahead prices, operating characteristics, plant availability, weather forecasts, and heat demand forecasts, among other factors.

Figure 2 illustrates the internal control concept of the master controller which can be divided into several parts introduced below. The parts are explained in the same order in which the setpoints are processed throughout the MC:

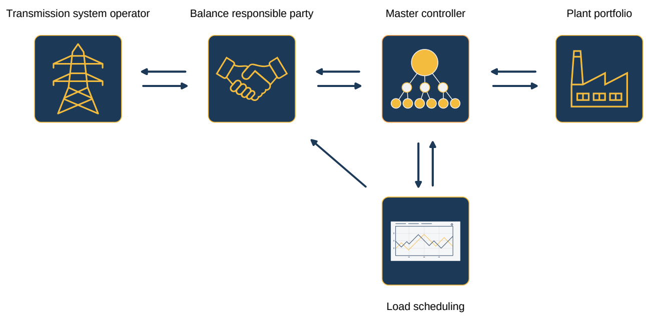

Figure 1. Simple overview of the master controller principles.

- **Load Schedule (LS)** handles the setpoints to each unit received from the load scheduling system. The setpoints are ramped according to the mFRR requirements and provide the operating point (baseline) of the system.
- **Balance Responsible Party (BRP)** handles ancillary service activations to the plant. The BRP setpoints are not individual for each production or consumption unit as the strategy requires the system to be considered as one aggregated unit. Each of the AS — FCR, aFRR and mFRR — are ramped according to their respective requirements, added together, and forwarded to the power balance controller.
- **Power Balance Controller (PBC)** ensures that the power consumption is equal to the traded power including AS. This is necessary since non-dispatchable consumption from pumps, lights, or similar ("house load") can be hard to predict. The PBC uses a closed-loop controller to steadily reach a zero deviation between the actual and traded power.
- **Waterfall method** distributes the AS and balance controller output to the units in a prioritized order. Higher-priority units respond to activations before lower-priority units.
- **Output signal handling** Each unit receives their load plan and the respective power correction from the waterfall method. Additional unit-related control such as output magnitude and ramp limitation, and unit start/stop is also handled here.

3 Model

The model is built of components from the Modelica Standard Library (MSL) (Modelica Association 2020) and blocks from an in-house developed library. It is divided into sub-models that correspond to each of the

 DOI 10.3384/ecp218253

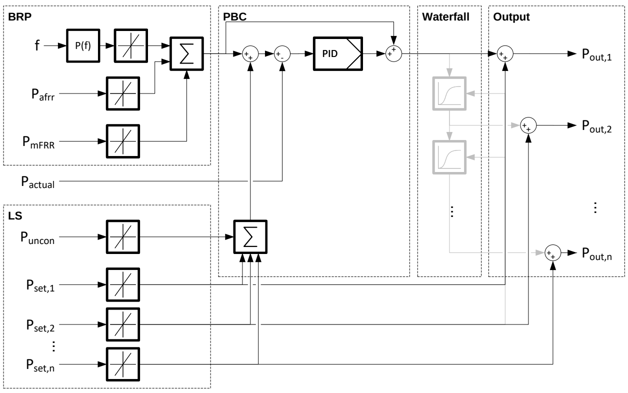

Figure 2. Overview of the master controller principles.

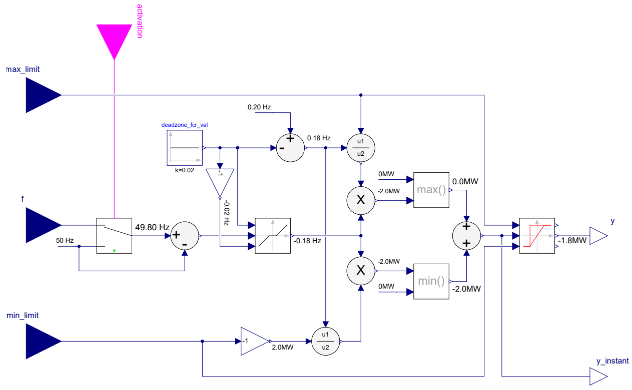

Figure 3. Diagram of FCR input handling

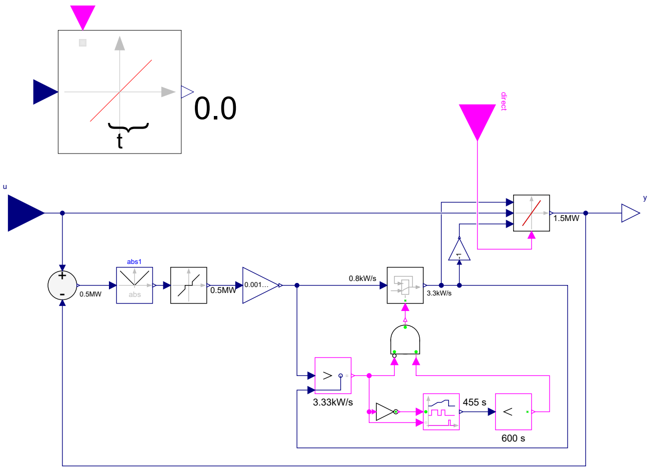

Figure 4. Modelica icon and diagram view of the constant timed ramp model including its icon.

framed sections in Figure 2.

3.1 Input handling

Input handling comprises the BRP and LS functionality and the master controller inputs consist of load schedules and ancillary services signals. The load schedules are ramped to comply with the mFRR response times, and the resulting signal is routed to each corresponding unit to define their operating points. The ancillary services, which are later distributed to the units by the waterfall method, are ramped to meet their individual timing requirements. Thus, the aFRR input signal is ramped so that it can reach reserved capacity in less than 5 minutes. FCR is ramped so that the reserved capacity can be reached in 30 seconds. This differentiation in ramping, depending on the source of the activation, ensures that the reacting units and their components do not react faster than necessary.

Most setpoint sources are simply treated with a ramp. The exception is FCR, which is derived from the grid frequency deviation. Figure 3 shows the FCR calculation model. It converts the frequency measurement deviation from 50 Hz, into a proportional power response based on reserved capacity in the system. A rate limiter placed at the output performs the final setpoint shaping.

The input of the load schedule and the mFRR setpoint is ramped differently than the other ancillary services. The ramp times are set to 10 minutes to ensure compliance with the formal 12.5 minute requirements. mFRR responses that deviate from the 10 minute ramping are expressed as bills to the participant, depending on the current imbalance price. Therefore, these inputs are handled with the model shown in Figure 4, which is a modified rate limiter that adjusts the slope to ensure that the input value is reached after a prescribed time. It works by continuously calculating the slope required to reach the input from the current output and then inserting it into the rate limiter model. An *analogue memory* model in the centre of the figure saves the slope until a larger slope is derived from the calculation, at which point the updated maximum slope is replaced. With this method, the model works with any arbitrary input.

3.2 Power balance controller

The Power Balance Controller (PBC) compares the actual (measured) total power consumption with the combined setpoint of AS and the sum of LS setpoints. It uses an integrator followed by a PID controller to adjust the combined setpoint, so that the electric energy deviation becomes zero within a 15 minute horizon when the integrator is reset. See Figure 5. The 15 minutes cycle is introduced in the model because the BRP issues the electric power bill to the plant owners on the basis of 15 minutes energy accumulation.

Typically, it is difficult to predict the intermittent house load and, therefore, the PBC can help to precisely reach the setpoint, removing any possible economic consequences when deviating from the traded power. Some units have to actually change their consumption, but that is distributed by the waterfall method (subsection 3.3). The controller output is added to the AS feed-forward setpoint as shown in Figure 2.

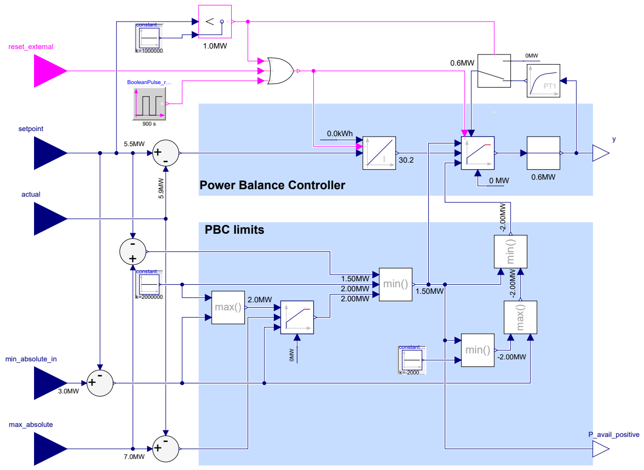

Figure 5. Modelica graphical model of the power balance controller.

In Figure 6, the system used later on in Section 4 is subjected to a step change in the ancillary services input. This is done both with and without the power balancing controller to compare the responses. Looking at the response without PBC, it is clear that there is a non-zero steady-state error. This is simply because, in this case, the system is essentially an open-loop system that does not correct the (non-dispatchable) house load of the system.

The steady-state error is cancelled when the PBC is activated. However, it introduces a relatively large overshoot that settles at the setpoint after 5–10 minutes. The overshoot is intentional and ensures that the accumulated energy deviation from the day-ahead energy purchase is zero over a 15 minute horizon. This means that steady state must be reached in 15 minutes, and that the overshoot actually compensates for energy that was not covered in the beginning of the response. This is illustrated by the two equally sized hatched areas. Normally, the overshoot is not going to be this large. It is purely due to the step input, which is not a realistic input and is chosen here to show the concept.

3.2.1 Power limitation

The power balance controller is also used to limit total power consumption or production. This is useful if faults occur on a transformer or at the electricity provider. When the power limit is reduced (input "max_absolute" in Figure 5) or the power consumption exceeds its limit, the power limitation proactively reduces the PBC output (via its upper limit input signal), potentially driving it to a negative value. This ensures that the sum of all

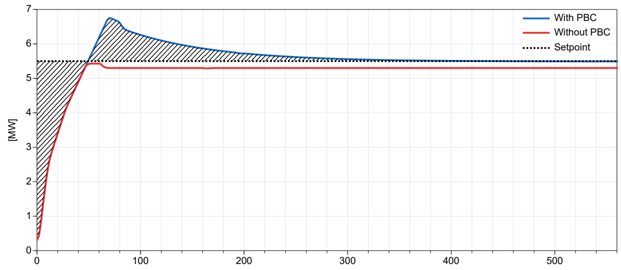

Figure 6. Demonstration of the power balance controller.

setpoints for the units is immediately reduced to comply with the transformer power limits. The power reduction is distributed to the units in the same way as the output from PBC.

If the setpoint is inaccurate compared to the actual power consumed or produced, then the above proactive (feed-forward) power limitation may be insufficient. Thus, an additional feedback controller is used to ensure that the actual power does not exceed the power limit. The minimum of the proactive power limitation and the controller limiter is the final maximum power used in the PBC.

3.3 Waterfall method

The *waterfall method* distributes one setpoint correction among several units according to a denoted priority associated with each unit. Generally, it means that one (total) setpoint is delivered to the unit with priority 1. The deviation between the setpoint and the *expected* dynamic response from the first unit is sent to the unit with priority 2. Similarly, the expected control deviation from unit 2 is forwarded to the unit with priority 3. This pattern is repeated until the setpoint is zero or the system runs out of available units.

The waterfall method is demonstrated with a system of three units: a slow heat pump as the first priority. A faster electric boiler as second priority, and a fast, but low-energy battery as third priority. The step responses from the output of the PBC sub-model to the individual units are shown in Figure 7. This corresponds to a 1 MW step change in ancillary services.

The heat pump receives the entire correction but has a large deviation at the beginning of the transition. The estimated deviation is passed on to the electric boiler which covers more of the transients. Lastly, the expected deviation of the electric boiler (shown as a short pulse) is passed as a setpoint to the battery The combined responses of all three units cover most of the delivery as shown with the output P_{total}.

This waterfall method consists of mainly two parts to

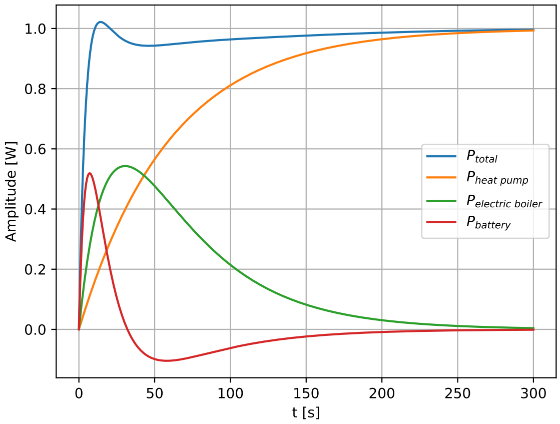

Figure 7. Step response of linear system with three units and two first order unit estimations.

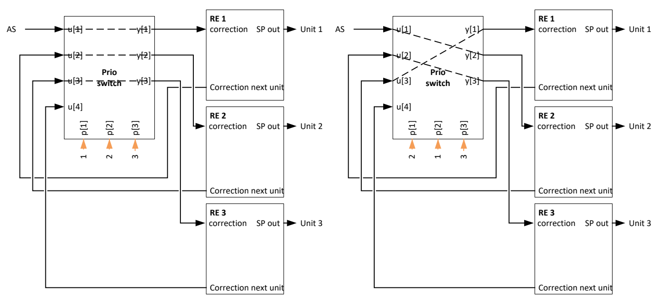

Figure 8. Different configurations of the priority switch

operate:

- **Priority switch** routes setpoints from the unit with first priority to the unit with the next priority and so on. It allows dynamic switching of priorities between units.
- **Response estimation** is the function that is repeated for each unit to estimate the power residual that *cannot* be covered by the unit. This residual is forwarded to the next unit in the order.

3.3.1 Priority switch

The priority switch is a function that connects each output to one of the inputs, such that the setpoint is handled by each unit in the order specified by the list of priorities. The priority switch is required in the case that different priorities are needed. otherwise, simple routing between each Response Estimation is sufficient. Figure 8 illustrates the layout of the waterfall method comprising the three mentioned units and shows the functionality with two different prioritization lists. In the figure to the left, the priority list is $\{heat pump = 1, boiler = 2, battery = 3\}$. This is the simplest configuration as it connects the outputs with the corresponding indexed input, indicated by the dashed lines inside the "prio switch" block. The priority list can be any incremented list of integers starting from 1 $(1 : n_{units})$. Thus, another valid case is $\{heat pump = 2, boiler = 1, battery = 3\}$ which is the case shown in the figure to the right.

Reversing The priority switch features a Boolean input that enables automatic reversal of the priority order during operation. This is useful because a participant may not want the highest priority unit to react as the first unit in both up- and downward directions. Using heat pump and electric boiler as an example, the heat pump produces much cheaper heat. Therefore, it is desirable to use the heat pump as priority 1 when increasing its load. But when the combined load is decreased, it is beneficial to reduce the load of the electric boiler before the heat pump. Therefore, when ancillary services move from a positive setpoint to a negative one, the priority order can be reversed to ensure the cheapest unit is always running. Additionally, the model gives the option to include only the first n units in the reversing. This is again useful if a system uses a battery to only adjust for the last corrections as the last priority. Then it is undesired if it suddenly acts as the first priority due to the reversing action.

3.3.2 Response estimation

It is commonly known that real measured responses introduce delays and more complex higher-order dynamics. While implementing a feedback control scheme for the priority switch inputs using real measured responses would be quite simple, introducing these delays in the control of the units would result in unnecessary activation of all prioritized units, which is undesirable. This is the case, because the delays in the measurements of the first priority unit will cause a small amount of the signal to be forwarded to the subsequent unit and so on, even if that unit is not necessary to meet the response requirements.

This phenomenon is also shown in the top graph of Figure 9. The figure shows a comparison between using real measured responses as feedback instead of estimations (top graph) and using an approach based on a rate limit block (bottom graph). The figure uses the example featuring a heat pump, an electrode boiler, and a battery energy storage system. In this test, the electric boiler is parametrized to meet the FCR requirements, eliminating the need for a faster battery in this specific case. However, it is clear that on the top graph, which uses real measurements, the battery still contributes, even though it is unnecessary.

For that reason, the dynamic response of the units is estimated with rate limit blocks as they react without any delay. Although the implementation is not as simple as using measured responses, using rate limit blocks prevents the lowest priority unit from being activated every time. This is shown in the bottom graph of Figure 9.

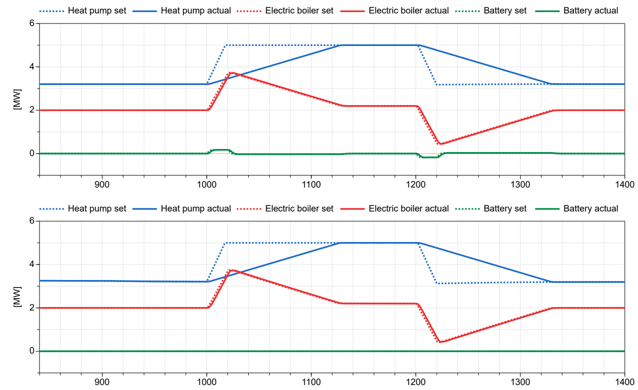

Figure 9. Comparison of waterfall method using real feedback (top figure) or estimated feedback (bottom figure).

The response estimation (RE) model utilizes information on ramping and limiting capabilities of the corresponding unit to estimate the response of the unit. The setpoint that is fed back to the priority switch to be routed to the next unit is the input setpoint subtracted by the estimated response of the unit. This can be formulated with Equation 1:

$$r_{corr,i+1} = (r_i + r_{corr,i})(1 - \hat{G}_{p,i}) \tag{1}$$

Where r_i is reference (load schedule) for unit i. $r_{corr,i}$ is the correction setpoint from the priority switch and $\hat{G}_{p,i}$ represents the estimation model of the unit. The combined setpoint that is sent to the unit is $u_i = r_i + r_{corr,i}$. A block diagram of the example with three units (without the priority switch) is shown in Figure 10.

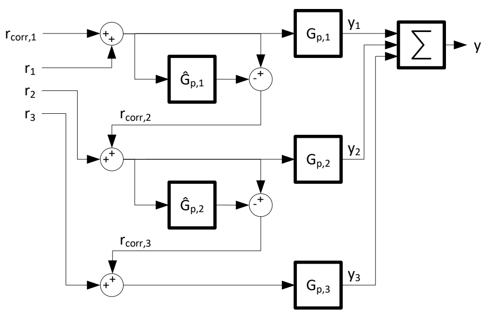

Figure 10. Block diagram of the waterfall principle

In the simulation, $\hat{G}_{p,i}$ is simply modelled with a power rate limiter and absolute limits of the unit. The rate limiter is sufficient to model the unit, as the master controller ramps the setpoint before it is sent to the unit. However, the unit typically also ramps the input itself if the input ramp is faster than the capability of the unit. Therefore, it is also important for this strategy to receive up-to-date information on the ramping capability of each unit as it is used as input in the response estimation.

If the actual ramping slopes differ from that used in the unit response estimations, the responses can become less accurate.

3.4 Automatic shutdown

A function can be implemented to automatically shutdown a unit in the output signal handling module mentioned in Section 2. This function can cause the unit to shut down when ancillary services combined with the planned setpoint reach the minimum limit of the unit. When this occurs, the waterfall method ensures that other units with lower priority temporarily replace the lost power during the ramp, resulting in almost continuous control.

3.5 Ancillary service validation

A validation function is modelled to ensure that activated ancillary services are within the response envelope required by the transmission system operator (TSO). The model calculates the upper and lower bounds required in the response and uses the constant-timed ramp model in Figure 4.

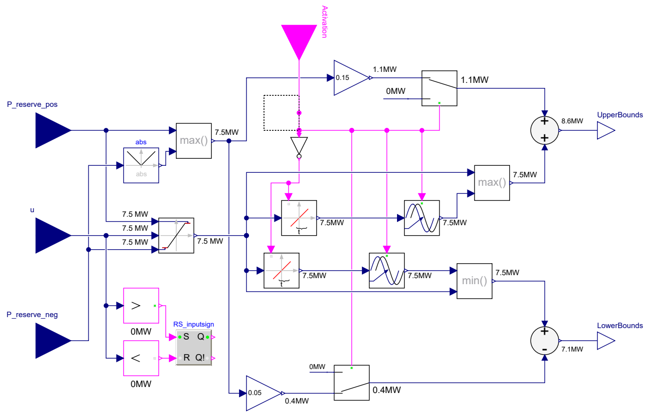

Figure 11. Modelica graphical model of the ancillary services validation.

This validation model (Figure 11) is based on the pre-qualification requirements in DK1, described in (Energinet 2025c). The model is applied to the setpoint components (load schedule, aFRR, mFRR) that cover the relevant ancillary services, after which the upper and lower bounds of each setpoint component are added together and plotted with the actual total response of the system for comparison. An example of bounds applied on a square-wave input setpoint is shown in Figure 12.

The envelope shows that the *increasing* source signal (the square wave) is allowed to rise immediately and must increase at least as fast as the blue ascending slope of the lower boundary. In the opposite direction, the source signal can decrease immediately and must decrease at least as fast as the blue descending slope. Similarly, the envelope allows for a certain steady-state deviation from the setpoint, indicated by the vertical margins. The validation model is especially useful when using realistic aFRR data that are updated continuously.

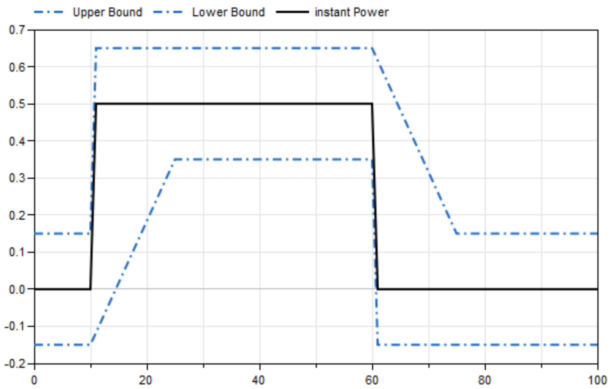

Figure 12. Modelica graphical model of the ancillary services validation envelope.

4 Simulation results

To demonstrate the functionality of the master controller, the concept is exemplified with three aggregated units in a PtH district heating system: a slow but inexpensive heat pump, a faster, but more expensive electric boiler and a battery with limited energy capacity. Participating in the AS markets requires some flexibility, so the example here assumes that a sufficient large heat storage tank can adjust for heat production variations.

The characteristics of the three units are shown in Table 1.

Table 1. Key model parameters.

Quantity	*Value*	*Unit*
Heat pump rate limit	± 0.0125	MW/s
Electric boiler rate limit	± 0.05	MW/s
Battery rate limit	± 0.1	MW/s
Heat pump max limit	5	MW
Electric boiler max limit	5	MW
Battery max limit	1	MW
Heat pump min limit	2	MW
Electric boiler min limit	0.4	MW
Battery min limit	-1	MW
Heat pump priority	1	-
Electric boiler priority	2	-
Battery priority	3	-

Heat pump response times vary considerably depending on, e.g., compressor type. (MAN 2024) shows a heat pump that can participate in part in FCR and (Danish Energy Agency 2025) shows that heat pumps can typically ramp with 5 % in 30 seconds. Thus, to show the concept working on slower types of units, the rate limit of the heat pump in the table is selected to correspond to 400 seconds. (PARAT 2020) is an example of a fast ramping electric boiler with full FCR capabilities. Instead of 30 seconds, the selected rate limit in the table corresponds to 100 seconds, to show that the performance of the equipment does not need to be maxed out. Based on the notion that batteries regulate fast, 10 seconds is selected for the battery. Thus, the rates have been selected so that the heat pump is too slow for aFRR activations, whereas the electric boiler is not.

Other notable features in the simulations are

- The heat pump and electric boiler are configured to swap priority when $r_{corr,1} < 0$.
- A simple State-Of-Charge (SOC) proportional controller is added to the battery setpoint to ensure the battery is within capacity limits of 1 MWh.
- The heat pump can automatically shut down as described in Section 3.4.

In the following sections, the system is simulated mainly with aFRR responses, since its requirements are more strict than mFRR. FCR is even more strict, but in DK1 the activated bids are smaller and less profitable than aFRR, and thus from an economic perspective not as interesting.

4.1 Up- and down regulation

This simulation adds an aFRR input to the control scheme. The combined load schedule is around 5 MW, and the aFRR signal contributes first with an additional 5 MW in upward direction followed by 4 MW in downward direction. This is close to the maximum capacity possible with the chosen set-up given the minimum limits of the units. See Figure 13.

The combined electric power consumption is shown on the top and stays within the allowed envelope during the simulation. The second plot shows how the response is distributed between units. The heat pump reacts first, followed by the electric boiler when the heat pump reaches its limit. The battery reacts when the electric boiler reaches its limit. During the load reduction, starting after 24 minutes, the heat pump and electric boiler swap priority, causing the electric boiler to back down first, followed by the heat pump and then the battery. The responses from different units are delivered without any notable bumps in the combined response.

Note that the battery only contributes when the other units reach the limits and not every time there is a correction to the overall setpoint. This shows the advantage of

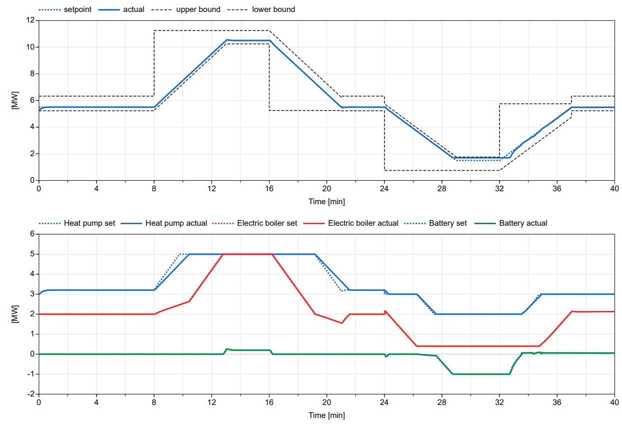

Figure 13. Simulation of an aFRR response on the example system.

using rate limiters to model expected responses.

4.2 Automatic startup and shut-down

The simulation test in Figure 14 shows the option of letting the heat pump automatically stop when the setpoint is below minimum. Similarly, it starts when the setpoint is above the minimum again. Thus, allowing the system to regulate in the downward direction with 5.5 MW, contrary to the 4 MW in Figure 13. In addition, the shutdown dynamics have been modelled to be faster than normal rates, whereas the startup rate is assumed to be equal to the continuous rate limit.

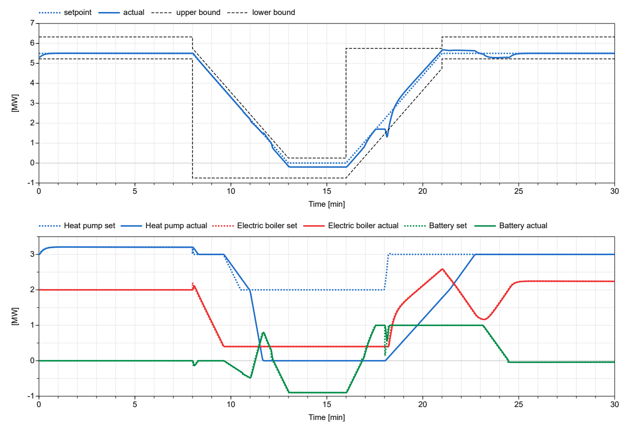

Figure 14. Simulation of auto-stop response on the example system.

Only minor bumps in power consumption are seen in combined consumption. However, after approximately 18 minutes, the total consumption decreases as the battery reaches its maximum power limit. With the current concept, the electric boiler cannot contribute from the waterfall method, since at this moment the priority is 1 for the electric boiler and 2 for the heat pump, due to the reversing. In this case, the PBC is utilized to increase the setpoint of the electric boiler before the total consumption deviates for a prolonged period of time.

4.3 Realistic aFRR activations

The entire system is tested with actual DK1 aFRR activations gathered from the Danish TSO data base (Energinet 2025a). The data covers ranges between ±150 MW, but is then scaled with 0.035 to match the showcased capability (around 5 MW in each direction).

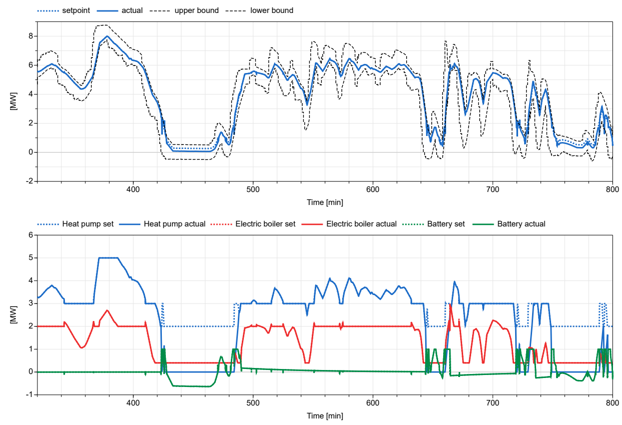

Figure 15. Simulation of an realistic aFRR response on the example system.

The simulation is shown in Figure 15 which shows several elements introduced in the paper. The waterfall method coordinates the units as expected even when using this continuous setpoint. Evaluation and visualization of the upper and lower boundaries of aFRR activations show that the response is, in fact, within the requirements. Occasionally stopping and starting the heat pump is not handled perfectly as there are some quick spikes in the consumption, but they stay within the boundaries. During the large load reductions, the automatic stopping of the heat pump happens fairly often. This could indicate potential issues with real heat pumps if they are not built for frequent on and off cycles. Lastly, the SOC controller for the battery is seen to apply small and slowly changing setpoints. This is compensated by some of the other units by the PBC.

5 Discussion

This work presents the concept of aggregating the electric power units of a PtH district heating plant. It shows that the master controller concept offers several advantages in terms of providing flexible operation to the grid through ancillary services. Since participation in ancillary service can be economically beneficial, it should be taken advantage of as much as possible. This work shows that the master controller concept provides the following potential advantages.

 DOI 10.3384/ecp218253

- **Faster responses:** The power capacity in dynamically slow units can now participate in fast ancillary services.
- **Automatic shutdown:** Units with a relatively high minimum power level are no longer a limiting factor when aggregating, as long as there are other units that can replace its power level when it is off. If some units have a minimum power level of 50% (which is not uncommon (Danish Energy Agency 2025)), then the capacity in ancillary services can be practically doubled.
- **Power balancing:** Uncertainties in actual power consumption can be adjusted with the power balance controller to produce or consume the exact amount of power that has been traded, reducing the imbalancing bill.

It is important to note that this work is still quite theoretical, with multiple tests performed on estimated data instead of real-world data. However, the work presents a general structure of the master controller concept from which it is possible to create more tailored solutions based on the setup and needs of individual PtH plants.

The work highlights a potential economic benefit to be gained from implementing the different parts of the master controller. If a slow unit, such as a heat pump, cannot ramp fully within the 5-minute aFRR window, adding a 1 MW battery to help with transients increases the system's regulation capacity by 1 MW. The battery alone cannot be pre-qualified to deliver 1 MW due to its limited energy content (Energinet 2025c), so both units benefit from aggregation. The prices vary and evolve a lot, but at the time of writing, the price of the aFRR up-regulation capacity averaged around 70 €/MW/h throughout the year (Energinet 2024a). This can be directly translated into additional income when the battery is added to the portfolio. If it is assumed that the battery can only help half the time during 6 months of cold weather, then it is still possible to earn 151,000 € throughout the year. Additional income can be earned in the energy activation market, but it requires a more extensive analysis. An interesting next step would be to further explore this potential economic benefit with specific real-world PtH district heating plant setups and also adding actual imbalance costs, AS prices, etc.

The theoretical nature of the work means that some real-world complications are not taken into account. As mentioned, this work emphasises that some units can help other units decrease response time, but this is exclusively relevant if the response time is slower than the AS requirements to begin with. If all units are technically able to meet all requirements on their own, this concept could become less beneficial.

Another point made in this work is that it is possible to utilize a larger capacity by automatically shutting down units instead of being restricted by a minimum load on all units. Here, it has already been mentioned that this requires other units to help in meeting all requirements. However, the work does not account for any complications related to possible shutdown/startup restrictions. Such restrictions could potentially be due to, i.e., wear and tear, minimum downtimes during shutdown, or other unit obligations tied to shutdown/startup of the unit. As an example, it is quite common for some of the units used in the simulations (e.g., heat pumps) to have minimal downtime between operational cycles. In addition, there could be a limit on the number of shutdowns allowed within a day or other period of time. Any such restrictions are not taken into account in this work, but could cause the automatic shutdown function to be less useful.

6 Conclusion

This paper presents a master controller concept that aggregates electricity producing and consuming units of a power-to-heat district heating plant. The concept has shown potential to offer improved participation in ancillary services. Aggregating the units enhances the ability of some units to meet certain AS delivery time requirements. Using the described Waterfall Method, the faster units can assist the slower ones, ensuring that the overall response complies with the requirements. Thus, it is possible to increase the total capacity that can be offered to AS. Additionally, by utilizing the ability to completely shut down units, the aggregated system can also increase the total capacity offered to AS.

Along with these improvements associated with AS, the master controller concept can also increase performance on other parameters. The power balance controller component assists the system in maintaining zero deviation between purchased and used electricity. This can help prevent the undesirable costs associated with imbalances. The Waterfall Method in conjunction with the priority switch also enables the system to use the correct units for different AS activations. This entails that the cheapest units are utilized the most and the more expensive units are used last.

The potential benefits provided by the concept are shown using a simple example with a heat pump, an electric boiler, and a battery energy storage system. This example indicates that the concept offers a framework in connection with operating a power-to-heat district heating plant.

References

Bella, Alessio La et al. (2021). "A mixed-integer distributed approach to prosumers aggregation for providing balancing ser-

vices”. In: *Electrical Power and Energy Systems* 133. DOI: 10.1016/j.ijepes.2021.107228.
Danish Energy Agency (2025). *Technology Data - Generation of Electricity and District Heating*. Tech. rep. URL: https://ens.dk/en/analyses-and-statistics/technology-data-generation-electricity-and-district-heating.
Energinet (2024a). *Årlig statistik for reserver og modhandel 2024*. Tech. rep. URL: https://energinet.dk/media/fi1iozqw/arlig-statistik-for-reserver-2024.pdf.
Energinet (2024b). *SCENARIERAPPORT 2024 – 2040 Forventninger til fremtidens Systemydelser*. Tech. rep. URL: https://energinet.dk/media/aj1frc1q/151124da_forventninger-til-fremtidens-systemydelser-2024-2040.pdf.
Energinet (2025a). *Energi Data Service*. Tech. rep. URL: https://www.energidataservice.dk/tso-electricity/PowerSystemRightNow.
Energinet (2025b). *Overview of Ancillary Services*. Tech. rep. URL: https://energinet.dk/el/systemydelser/introduktion-til-systemydelser/oversigt-over-systemydelser/.
Energinet (2025c). *Prequalification of Units and aggregated portfolio*. Tech. rep. URL: https://en.energinet.dk/media/gkijvg05/prequalification-of-units-and-aggregated-portfolios-february-2025.pdf.
ENTSO-E (2018). *Electricity Balancing in Europe*. Tech. rep. URL: https://eepublicdownloads.entsoe.eu/clean-documents/Network%20codes%20documents/NC%20EB/entso-e_balancing_in%20_europe_report_Nov2018_web.pdf.
International Energy Agency (2020). *Power Systems in Transition*. Tech. rep. URL: https://www.iea.org/reports/power-systems-in-transition.
Iria, José and Filipe Soares (2019). “Real-time provision of multiple electricity market products by an aggregator of prosumers”. In: *Applied Energy* 255. DOI: 10.1016/j.apenergy.2019.113792.
Iria, José, Filipe Soares, and Manuel Matos (2019). “Optimal bidding strategy for an aggregator of prosumers in energy and secondary reserve markets”. In: *Applied Energy* 238, pp. 1361–1372. DOI: 10.1016/j.apenergy.2019.01.191.
MAN (2024). *MAN heat pumps Transcritical compression cycle (TCC)*. Tech. rep. URL: https://www.man-es.com/docs/default-source/document-sync/man-heat-pump-hpu328c592b5d584f9eba47f4b6bed3574e.pdf?sfvrsn=9bd416fb_14.
Meyn, Sean P. et al. (2015). “Ancillary Service to the Grid Using Intelligent Deferrable Loads”. In: *IEEE Transactions on Automatic Control* 60.11. DOI: 10.1109/TAC.2015.2414772.
Modelica Association (2020). *Modelica Standard Library*. Tech. rep. URL: https://github.com/modelica/ModelicaStandardLibrary.
PARAT (2020). *High Voltage Electrode boiler for Steam and Hot water*. Tech. rep. URL: https://www.parat.no/en/products/industry/parat-ieh-high-voltage-electrode-boiler.
Tan, Jin and Yingchen Zhang (2017). “Coordinated Control Strategy of a Battery Energy Storage System to Support a Wind Power Plant Providing Multi-Timescale Frequency Ancillary Services”. In: *IEEE Transactions on Sustainable Energy* 8.3. DOI: 10.1109/TSTE.2017.2663334.