FRAUNHOFER INSTITUTE FOR ALGORITHMS AND SCIENTIFIC COMPUTING SCAI

1 *Organic redox flow batteries, a technology to store large quantities of energy from renewable sources.*
© Fraunhofer ICT

EAO – A FRAMEWORK FOR OPTIMIZING DECENTRALIZED PORTFOLIOS AND GREEN SUPPLY

Contact

Dr. Tobias Pfingsten
Head of Quantitative System Engineering, STEAG GmbH
tobias.pfingsten@steag.com

Dr. Daniel Oeltz
Fraunhofer SCAI
Schloss Birlinghoven
53757 Sankt Augustin
Phone +49 2241 14-4224
daniel.oeltz@scai.fraunhofer.de

1. Summary

Background and target:
The optimization of single assets or a portfolio of assets is a ubiquitous task in energy and commodity trading. Assets may be of various types, such as storage facilities (e.g. batteries or water reservoirs for power, heat or gas storage), decentralized or large power plants to produce or convert a commodity, transport via shipping or pipeline as well as sourcing or supply contracts partly with terms such as minimum or maximum offtake.

Implementation for a specific problem of the above class is typically done using linear programming, mixed integer programming or dynamic programming. The problem setting is well known and there are many publications on specific applications as well as commercial software solutions. However, for some problems commercial software packages do not provide enough flexibility to model all relevant features properly and implementation work must be invested to solve such problems adequately. Hence we have seen people repeatedly working on seemingly different problem settings which could probably be solved with much less effort using a modular framework.

The aim of this technical report is to formulate a unifying way of describing such tasks and to discuss a modular framework that enables practitioners to optimize single assets as well as simple or complex portfolios without the need to resort to specific (and often expensive) software applications or the need to implement the mathematics themselves. We believe that the solution presented here, and implemented in the Python package EAO, can be useful in many real life setups – from building a virtual power plant from decentralized assets to scaling

power purchase agreements (PPA) in green power supply or optimizing a portfolio of sector-coupling assets.

Green power supply as an illustration:
In this report we explore a problem setting that we have been facing recently: In decentralized power, wind and photovoltaic (PV) assets are combined with new storage solutions such as batteries and diverse small scale CHP (combined heat and power) units with local heat demand. As more and more companies work on minimizing their carbon footprint, they ask suppliers to deliver power directly from green sources without detours via certificates. Since volatile renewable generation typically does not match load profiles, the challenge for suppliers is to create and steer a cost optimal portfolio of renewable sources, flexible sources and storage that guarantees a certain fraction of green power for the client.

2. Introduction

2.1 Asset optimization in energy and commodity trading

Asset optimization is a common task in energy and commodity trading. Let us start by describing some specific examples without going into much detail – we will rather show later on how their common structure can be exploited to build a generic framework to solve all of the following examples with the same approach.

Decentralized power:
Power generation has been facing drastic changes in recent years. Large scale generation assets such as coal or nuclear power plants have been and are being replaced by a large number of smaller decentralized assets. To take PV as an example, the World Energy Outlook 2020 [10] forecasts the installed capacity of solar PV worldwide to surpass that of coal by 2026[1].

Decentralized assets may be wind and solar PV generators with volatile generation, small combined heat and power plants (CHP) that produce power in combination with heat (typically consumed locally and leading to complex restrictions), power to heat or power to gas assets or batteries to just name a few. At the same time, flexibility in consumption is used more extensively. The term "prosumer" subsumes the idea that the clear distinction between production and consumption becomes fuzzier as consumers may themselves invest in small assets (down to batteries in households) and monetize their flexibility in consumption. Operators face the need to optimize a large portfolio of hundreds of such assets. The challenge often lies in the fact that all assets are somewhat special, e.g. the operator needs to take into account the specific restrictions of each CHP plant with its own combination of heat storage facilities. The operator will therefore try to translate the joint behavior of all assets into one "virtual power plant" that can be handled and marketed by traders.

Sector coupling:
Commodity sectors in energy have traditionally been treated separately in daily operation. With the wider application of volatile renewable generation, CHP plants as well as newer applications such as power to heat and power to gas, a separation no longer holds. To name two examples: flexibility from heat storage generates a flexibility on the power side via power to heat and CHP, that may directly compete with batteries. Flexible power to heat or power to gas can use surplus renewable power and links power, heat and gas markets.

Gas supply:
Gas supply is typically governed by long-term supply contracts with complex terms such as minimum offtake volumes, seasonal storages, pipelines and consumption from gas power plants and clients. As markets liberalize, a joint optimization of all assets (contracts or physical) is essential.

Other fuel markets:
The above arguments hold as well for other fuel markets such as coal. While coal may be declining in importance, in the current transition phase, the impact of sector coupling is highly prevalent in this market: steam coal is practically only used for power generation.

[1] *Note that several scenarios are analyzed.*

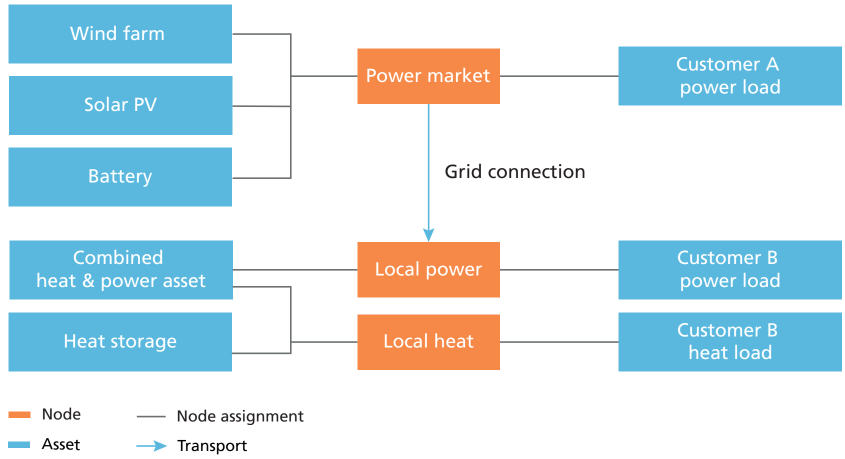

3

Traditionally, coal plants were used for base or mid power with very stable generation volumes, and thus very stable coal consumption. Today, coal consumption is highly volatile, as coal plants are pushed out of the market when renewable generation is high, heat demand is low or gas prices are low enough so that producing power from gas is cheaper. Consequently, a joint and robust optimization – of power plants, long-term sourcing contracts, storage at various locations, shipping and rail transport – that takes uncertainty into account, is essential.

Green energy supply:
In the early years of power from wind, PV and biomass, these have mostly been running under regulated feed-in-tariffs, where offtake and remuneration were guaranteed by regulation. As those technologies,in particular PV, have become competitive, they require new marketing channels. At the same time, consumers are putting more importance on minimizing their carbon footprint. New products have been developed to meet those new demands: green power purchase agreements (PPAs) guarantee the asset owner a fix price. Structured downstream contracts guarantee a minimum fraction of green energy to the consumer. As large-scale batteries are becoming cheaper, they are used to bridge the time difference between generation and consumption.

We explore green energy supply in more detail. The example serves us as an illustration of the main features of energy asset optimization and, at the same time, we believe that we can shed some light on the main features and challenges of these new products.

2.2 Framework to solve this family of problems

As we will outline in the following, literature is vast and commercial software packages that solve specific tasks exist. However, from our own practical experience, what seems to be missing is an approach that takes advantage of the common structure of most applications in asset optimization in energy and commodity trading. Despite available textbooks and literature, as well as commercial software, problems are often solved in-house and much energy is wasted in reinventing the wheel.

The common structure is key:
The main structure of the framework is illustrated in figure 3. The assets to be optimized and the involved commodities may be very different, but they have in common that:

- Assets may have a complex structure of restrictions and states. However, we are interested in their "dispatch", the quantity of a commodity that goes in or out of the asset at each point in time. There can be various different commodities involved at the same time.
- Assets do not interact directly. The connection of assets in a portfolio is given implicitly via the flow of commodities between the assets.
- The main structure of a portfolio is given by nodes (i.e. virtual trading locations per commodity) and connections (transport) between those.
- In each node, commodities are conserved, i.e. the dispatch of all assets and transports must add to zero for each commodity for each point in time.

Prototype assets:
Assets may be quite complex in reality. However, we have found that there are a few prototype assets that can be used in many different circumstances, if they are defined in a generic way:

- Storage: for any commodity, storages are mathematically very similar. They are defined by the capacity (maximum dispatch) in or out, storage size, costs and efficiency. Physically it could be anything from a battery to a pile of coal

3 *Illustration of a portfolio setup in energy asset optimization. The starting point are assets, which may have a complex structure. As they are added to the portfolio, they may interact via their consumption or supply of a commodity. The commodities used in this illustration are heat and power. Nodes are locations or virtual trading points, where the commodity is exchanged between assets. Nodes themselves may be linked by (potentially limited) transport. In each node, the sum of flows must be zero for each point in time for each commodity.*

- Basic contracts: in its basic form, a commodity may be bought or sold at a given price up to a certain quantity per time step.
- Complex contracts or production assets: in their more complex form, contracts and production assets can be represented by basic contracts in combination with more restrictions. Those are typically minimum and maximum offtake over certain time frames or certain limits for each point in time.
- Transport: whether pipeline, cable or shipping, transport is defined as getting a commodity from one location (node) to another in a specified time and at given efficiency and cost.

New assets are often similar to those of the base type and come with just a few more specific restrictions or cost components. In this case we found it very useful to have a framework at hand that allows us to inherit from the base classes to be able to specifically add the missing functionality instead of building new assets from scratch. To name two examples, we implemented scaled assets that allow the user to optimize the overall capacity of an asset against fixed costs and we introduced structured assets where the user may "wrap" several assets in one more complex asset.

Formulation as a linear programming problem:
Our framework makes use of linear programming (LP) or alternatively mixed integer programming (MIP). The very first step is to define a common time grid for the time range analyzed, which is used for all assets. This discretizes the asset dispatches to a number of points in time. In this section we keep the discussion at a descriptive level. Please refer to the separate sections on mathematical details.

Combination of assets into a portfolio:
Once the internal mechanics of assets are formulated as restrictions of an LP, assets can be combined to form a portfolio. Here, the common structure as described above can be exploited. The main idea is that assets only interact via the flow of commodities they generate, i.e. their dispatch. In the formulation of LPs, this means that assets come with their own set of variables and restrictions. Without any interaction, those LPs can be combined in one large LP for the whole portfolio. The interaction comes into play only as we introduce the common set of restrictions, that the sum over all dispatches in each node must be zero for each time step.

The idea is trivial, but it is essential for being able to reduce complexity in setting up a generic framework for portfolio optimization. We can define detailed and specific assets separately and then combine them modularly in any portfolio structure.

The EAO Python package:
We implemented a modular Python framework named EAO[2], based on the above principles. This package enables the practitioner to use standard assets, define new specific assets, combine and optimize them in a portfolio or structured assets for specific or multi commodity use cases.

The EAO software package is not trying to compete with commercial software applications. Instead, we are trying to bridge the gap between "implementation" from scratch for each application and a "black box solution". It offers a modular tool box in Python, that provides standard assets such as contracts, transport and storages and a translation to obtain a closed LP formulation on the level of the portfolio. It also to manages serialization and input or output. New assets are easily added by an experienced Python user. Its mathematical basis is described in a separate section.

2.3 Literature
Throughout the report we deliberately put only limited focus on a thorough literature review. The application of LPs in the optimization of large-scale portfolios is by no means new and has been widely discussed. Our goal is rather to provide a lean introduction to our approach and to directly dive into its application for green energy supply as an important topic in the industry today.

In our experience, LPs are very well suited for the optimization of assets in energy and commodity markets. Problems may quickly become large, but target functions and restrictions mostly translate to LPs in a very direct manner. Dynamic programming is the method of choice when optimizing single assets in greater detail.

---

[2] *See https://github.com/EnergyAsset-Optimization/EAO*

In our implementation we use CVXPY [4], a package to formulate convex optimization problems that serves as an interface to common solvers. Other packages can be integrated as well. Through CVXPY, our approach directly extends to mixed integer programs (MIPs). In the definition of our assets we have avoided MIPs where possible, as we have found that careful approximations in problem definition can spare us much trouble if the problem can be kept convex.

Linear programming in energy and commodity markets:
The application of LPs in energy is not new at all and omnipresent. Besides earlier works such as [6], the topic is described in text books on finance and commodity trading (e.g. [7]) as well as in a very large number of articles that describe specific implementations. Commercial software packages are available for specific purposes such as market simulation and optimization of single assets. From our experience, "black-box" software packages are a good choice for companies and teams with limited experience in modeling. However, very often companies implement their own solutions from scratch as they feel there is the necessity for a specific solution that suits their requirements.

Stochastic programming and robust optimization:
Uncertainty plays an important role in energy and commodity trading. Most assets such as flexible generation assets or storages can be seen as complex options that have an extrinsic value that should not be ignored in optimization. We leave quantitative finance aside here and concentrate on the question how uncertainty can be treated in optimization. Stochastic Linear Programming (SLP) was proposed by [2] early on to deal with uncertainty in a principled manner. The approach is very clear, but if extended to a multi-stage case, quickly creates problems that cannot be handled practically. [12] gives a very nice introduction to the topic. Other authors explore this framework for various applications: [3] for portfolio optimization and [13] for energy. Particularly for hydro power generation, stochastic planning techniques are important and are in use in countries such as Brazil and Colombia with a large dependence on hydro power. See e.g. [5] for a review of optimization algorithms used in hydro power.

As compared to other approaches, the SLP approach is easily applied to various setups. Other approaches could be dynamic programming or least squares Monte Carlo. Dynamic programming is, from our experience, a great alternative when optimizing single assets. As long as the state space can be represented with a limited number of states, it allows us to implement basically any transition rule and is extremely fast. However, as soon as just a few assets are combined in a portfolio, the state space basically explodes.

Least squares Monte Carlo, similarly, may be a good choice for specific applications such as the optimization of a single storage, options, etc. However, we have not seen so far an approach that would suit generic types of larger portfolios.

In our software package, we have implemented the two-stage setup where sampling is used to optimize a decision "today" under an uncertain future. We believe that for many cases this approach is good in practice, if the user keeps in mind its limitations. Please refer to the separate sections on mathematical details.

Robust optimization such as the maxi-min approach aims at improving the worst case in a distribution (see e.g. [9]). For a limited number of price samples it is easily implemented in an LP approach. The maxi-min approach is included in the EAO package together with some examples to show where it is a good optimization target and where it is not.

3. Green energy supply

3.1 Customer efforts to minimize their carbon footprint

Traditionally, customers in the power sector have been passive. Via full supply contracts, they resorted to utilities to securely provide them with power when needed. This has been changing in past years as customers started to monetize their flexibility in demand, e.g. allowing utilities or other service providers to cut demand in return for discounts. This exists throughout the world, a great example being the demand response program from PJM, the regional transmission operator in the eastern USA ([8]). In PJM, demand response is an

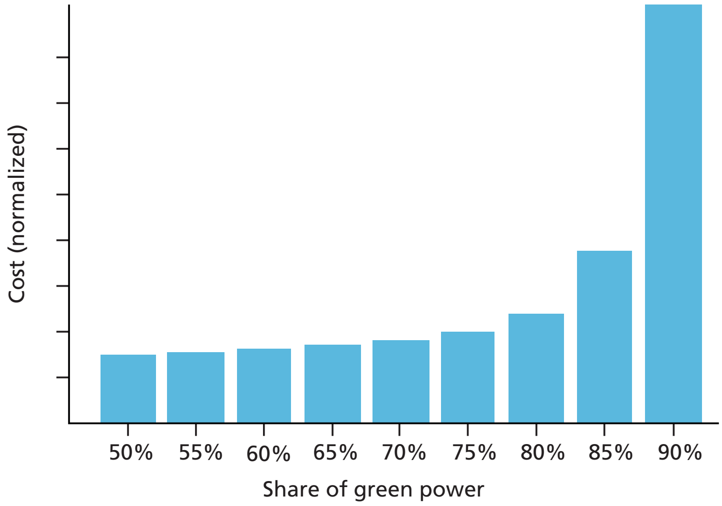

important factor in securing energy supply.

As customers are putting their carbon footprint on the agenda, supply with green energy is gaining importance. Major deals include Microsoft's deal to secure 500 MW solar PV[3] as well as similar activities by Google[4]. PPAs are a well-known concept in energy markets. Customers, traders or utilities guarantee to take power over a long period at a fixed price. This enables developers to invest in assets, as the secure income allows them to obtain financing for their investment. Green PPAs are a vehicle to secure green power supply for a company and provide a new way of financing green assets (mostly wind and PV) in addition to regulators' subsidy schemes.

Green PPAs or own assets are an important starting point for green power supply. A limitation lies in the fact that power is not easily stored. It has to be produced basically at the same time as it is consumed. Since the demand profile is typically fixed just as the production from wind and PV assets is given by whether conditions, the fraction of green power in consumption is limited to the intersection of these profiles. When the sun is not shining or the wind is not blowing, the consumer needs to resort to "grey" power from the market. Depending on the load profile, this typically limits green power to roughly 40 % - 60 % of consumption[5].

Batteries or other storage technologies provide the opportunity to maximize the share of green power in consumption. The power supply for a whole country, guaranteeing 100 % green power, while ensuring security of supply at the same time, requires an optimal interplay of various technologies (see [1] for an economic perspective).

Looking at solutions for single customers, the problem is simpler: the customer or its power supplier can build a portfolio from renewable energy assets and storage to reach the desired share of green power – depending on the cost sensitivity of the client. In times of peak load and limited renewable generation, the client can resort to the spot market and security of supply is no issue.

The illustration in figure 4 shows results for an illustrative, yet realistic setup. We chose hourly load and renewable generation data for Denmark[6] in combination with cost estimates for renewable sources and battery storage from [11]. We assume a given load of 100 GWh with a given target of generable share in supply over a full year. Further, we assume that the client would build assets or buy production from PV and wind assets and potentially apply a battery to maximize green share. The EAO framework was then used to find an optimal mix of wind, PV and battery capacity to meet the minimum green share of supply at lowest costs.

The main purpose of the illustration is to introduce the use case of the optimization scheme in green sourcing. It will be elaborated in more detail in the following. The simple case of figure 4 already shows, that beyond a certain share of green power (here roughly 80 %) costs rapidly increase. In this example, a rapidly growing battery capacity is needed to bridge the gap when wind is not blowing. Depending on the region, the optimal share of onshore or offshore wind, PV and battery capacity is essential to minimize costs.

3.2 PPAs and structured "green supply" contracts

Let us introduce the idea of "green supply" contracts that are currently discussed in the power industry. Their main features are the following:

- A full supply contract for the client's load, which is typically given as per 15 min or hourly time interval.
- A minimum share of green power is guaranteed in the supply contract, e.g. over the course of the year.
- Green supply realized via specific green assets (contracted via green PPAs and/or own assets), directly matching load and supply for each point in time.
- When the load and green generation are not balanced, the supplier may trade at the spot market.

4 *Costs of power supply at different shares of green power*

[3] From pv-tech, 'Microsoft announces firm's single largest green energy PPA', July 22, 2020
[4] From sustainability.google
[5] The fraction is highly dependent on the load profile. In this illustration we assumed that green PPAs match the total load per year.
[6] Load profiles from Denmark taken from energidataservice.dk (2020), pricing zone DK1

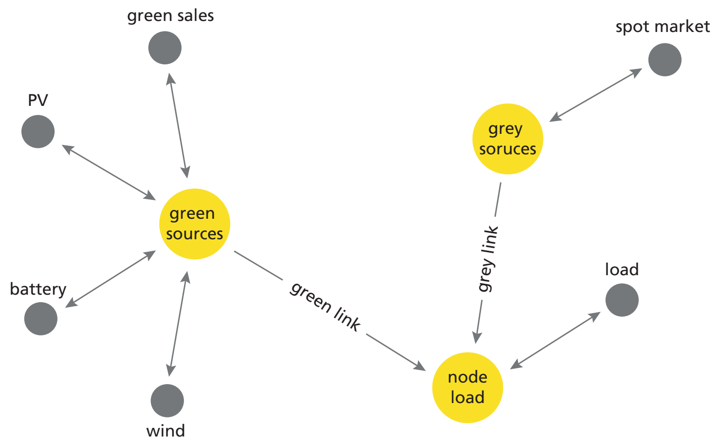

5

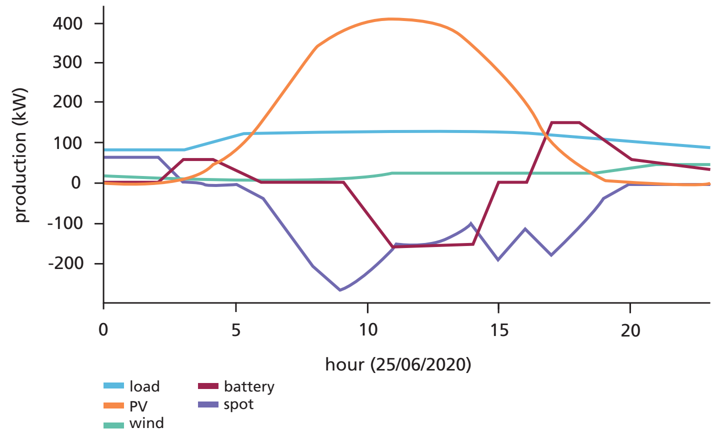

6

- A storage may be used in the portfolio, as long as the supplier can guarantee that it does not mix green and grey power.

In our optimization scheme we implement the contractual setup as shown in figure 5. The portfolio consists of assets and nodes, the main building blocks of the EAO package. Nodes are connected by specific "transport" assets as illustrated in the network graph: green power sources as well as the battery are placed in a specific node ("green sources"). Via a unidirectional link ("green link") the latter is connected to the node, where the load is placed ("node load"). This guarantees that in the course of the optimization no grey power enters the battery, which could be later mixed up with green power. On the green link we impose the restriction that, over the course of the year, consumed green power must exceed the contractual minimum share of green electricity. In addition we allow for the sale of production from the assets ("green sales") in times where there is excess production. The load is also connected to the spot market to be able to fill gaps in times when there is not enough green supply.

The model can be used for various use cases. We can have the algorithm automatically determine an optimal size of the battery and green generation assets to obtain minimal overall costs while meeting the agreed share of green energy. At the same time, overall costs associated with green supply are determined. Once all contracts and short-term forecasts are in place, the model can be used to perform the daily optimization required to schedule battery usage.

3.3 Illustration: An optimized green supply portfolio

In figure 5 we provide an example of a structured supply contract that guarantees a certain share of green electricity directly from wind and PV in combination with a battery. In this section we explore the example in more detail and show that an optimal mix of wind, PV and battery capacity can significantly reduce costs (or maximize the share of green energy[7]):

Simple case of volume matching:
In a first approach we do a very limited optimization. We match the overall yearly demand with 50% wind and 50% PV generation. The battery is then sized as to achieve the given minimum green power share (here 75 %). Figure 6 illustrates the hourly dispatch for a certain day. On this sunny day battery usage is straight forward. It stores electricity when sun is shining for usage in the early morning and evening. The spot market is used to provide backup electricity. As the portfolio is optimized in one go, battery usage is directly optimized against load, PV and wind profiles while exploiting low spot prices.

Optimized portfolio:
An essential cost lever in the portfolio is the battery. Particularly due to longer periods with limited renewable sources, it must be sized to bridge the gap so the minimum share of green power can be met, see figure 7 for a view on associated costs.

In the volume matched example, capacity of wind and PV was given. However, since the hourly profiles of our load, wind and PV are different, an optimized combination may minimize the need for battery capacity, or, respectively, ensure it is used more effectively. In our illustrative example we are using the overall wind and PV profile in Denmark.

5 Network graph of a green supply contract with green PPAs as supply. In order to separate green and grey power we introduce two separate nodes for each as well as the load. Using a yearly minimum transport volume in the green link from assets to load, we ensure a certain minimum share of green power.

6 Illustration of asset dispatch on a sunny day. In the early morning and evening, solar production is low and the battery is used to provide green power stored during the day. At night the spot market is needed to supply demand, while surplus electricity is sold to the market during the day. Note that at the same time, spot purchases vs. battery usage are optimized against the spot price.

[7] Example based on historic hourly load data and investment costs for renewables and batteries for Denmark (pricing zone DK1) for 2020.

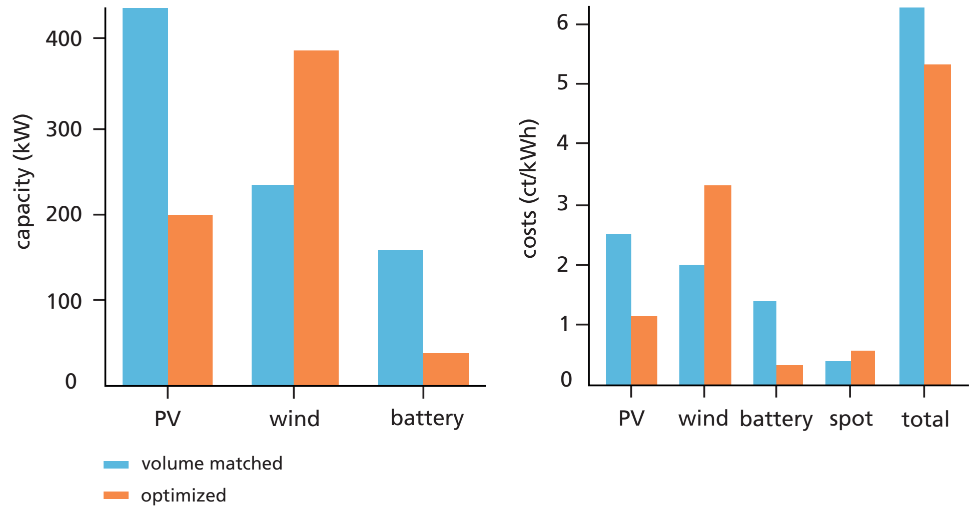

7

However, the effect will be much more profound as the profiles of several different assets (or even technologies) are combined. In the EAO package this is simply done by adding the assets to the portfolio (figure 5).

In comparison to the simple volume matched case, we generate a portfolio where the capacities of wind, PV and battery can be optimized under the restriction that overall production from wind and PV must not exceed load by more than 5 % (to roughly preserve volume matching in total), see figure 7. Costs are reduced by roughly 15 % as a smarter combination of wind and PV reduces battery needs significantly. Looking at the setup from the other direction, the same battery can be used to achieve a higher share of green power.

3.4 Importance of portfolio optimization in green supply

Portfolio optimization is not new, as we discussed in the introduction. In electricity, for example, it has long been used to coordinate power plant dispatch in a market. Take the example of Independent System Operators (ISOs) that are in place in various countries such as Colombia and some regions in the USA. The ISO has access to the technical data of all generation assets and uses an optimization framework comparable to the one described here (while of course much more elaborate in detail) to optimally dispatch every single asset to meet the load in the region.

In liberalized markets such as Europe, an ISO does not exist. Coordination is achieved via a spot market (day-ahead and intraday), where assets are optimized against price signals. In a liquid market, the result is the same as in the planned ISO case. However, from the optimization point of view, the starting point is completely different: in the market setup, each asset is and can be optimized standalone against the market price and no portfolio optimization is required. The market as such creates the link via a repeated iteration between re-optimization of all assets in the market and its effect on the market price.

Green supply that directly connects green sources and demand changes this picture. A supplier that guarantees a certain share of electricity from certain assets to its customer (or a customer that owns green assets to cover his load), is in the need to do an optimization within the portfolio and will resort to the spot market only for backup supply. In this setup an optimization scheme such as the one described here is needed.

3.5 Nodal prices, optimization against the market and auto trading

As long as operators steer a few large assets against the market, they will be able to have in mind their relevant technical restrictions such as marginal costs, gradients, etc. However, at any relevant size of the portfolio, tools are required that help the operator to summarize the behavior of the portfolio in few parameters. Particularly for portfolios of hundreds of decentralized assets with different characteristics, this aggregation is essential. The marginal generation costs of the portfolio for each point in time are the most relevant characteristics when optimizing against the spot market (day-ahead or intraday). At prices below the marginal generation costs, power will be bought and production from own assets reduced – and vice versa.

In the framework we describe here, this aggregation comes along naturally.

Figure 8 provides an illustration using a very simple sample portfolio of a few generation assets with different production costs. In the example they are dispatched optimally to cover a given load. The higher the load the more expensive assets are needed and therefore marginal generation costs are higher. As we are using an LP to solve the problem, the dual of the nodal restriction provides the marginal cost of the whole portfolio.

7 *Optimized green supply portfolio: We compare a portfolio of wind, PV and battery assets to supply a given load with 75 % green power. In the "volume matched" case PV and wind are sized to provide 50 % of sales production each. The battery sized to be able to provide a share of 75 % green power. In the "optimized" case all assets are sized automatically. Observe how optimized sizing alone can already reduce costs by roughly 15 % in this setup.*

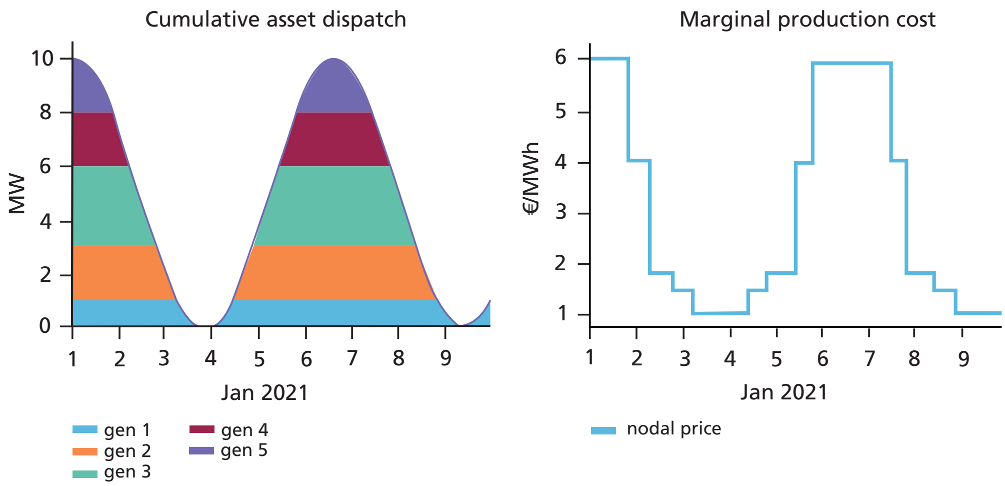

8

The framework applies to any portfolio and any type of asset that can be described as an LP, including more complex contracts or storages. Using the duals for nodal restrictions, the algorithm provides the value of the commodity at each node (or location). The interpretation is as follows: The dual of the nodal constraint for any point in time quantifies how the overall portfolio value changes if an infinitesimal amount of the commodity is made available.

If the operators need to assess the value of a specific deal for their portfolio, they can use an additional run of the optimization to explore the change in the overall portfolio value. The same applies to auto-trading schemes, where an algorithm takes the role of the operator. For such strategies in particular, a generic framework such as the above is required to generate rules that abstract from specific assets.

References

[1] J. Burchardt, P. Gerbert, S. Schönberger, P. Herhold, C. Brognaux, and J. Päivärinta. The economic case for combating climate change. 2018.

[2] G. Dantzig. Linear programming under uncertainty. Management Science, 1(3/4):197–206, 1955.

[3] G. Dantzig and G. Infanger. Multi-stage stochastic linear programs for portfolio optimization. Ann Oper Res 45, 45:59–76, 1993.

[4] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. Journal of Machine Learning Research, 17(83):1– 5, 2016.

[5] A. T. Hammid. A review of optimization algorithms in solving hydro generation scheduling problems. Energies, 13:2787, 2020.

[6] P. F. Hutchins. Energy resource allocation optimization – a linear programming model. 1976.

[7] G. S. M. Burger, B. Graeber. Managing Energy Risk: An Integrated View on Power and Other Energy Markets. Wiley Finance, 2014.

[8] PJM. 2020 demand response operations markets activity report. 2021.

[9] G. Polak, D. Rogers, and D. Sweeney. Risk management strategies via minimax portfolio optimization. European Journal of Operational Research, 207:409–419, 11 2010.

[10] International Energy Agency. World energy outlook 2020.

[11] International Energy Agency. Projected costs of generating electricity 2020. https://www.iea.org/reports/projected-costs-of-generating-electricity-2020, 2021.

[12] S. Sen and J. Higle. An introductory tutorial on stochastic linear programming models. Interfaces, 29:33–61, 04 1999.

[13] S. W. Wallace and S.-E. Fleten. Stochastic programming models in energy. In Stochastic Programming, volume 10 of Handbooks in Operations Research and Management Science, pages 637–677. Elsevier, 2003.

8 *The figure to the left shows how illustrative generation assets are dispatched optimally to cover a given load. They are dispatched according to their marginal generation costs. The right figure shows the marginal generation costs of the whole portfolio computed via the duals of the nodal constraint.*

A Mathematical framework

In the following sections, we introduce the mathematics behind our framework for energy asset optimization. The appendix' target is to explain the main ideas to the technical user. It is focused on the mathematical framework rather than the code of the Python implementation, which is explained in the corresponding documentation[8].

We believe, that our framework helps in providing a guidance and a starting point for solving many problems based on this code basis. It can easily be integrated into an existing machinery for trading and asset dispatch. Where needed, it can be easily extended with one's own building blocks such as complex structured assets.

A.1 Starting point: A portfolio of assets

A portfolio is a collection of assets that are analyzed in a group. Typically, grouping follows ownership and as the asset owner you would typically collectively optimize your assets to maximize the overall cash flow of your portfolio.

Assets: Assets may be of very different types. In energy or commodity trading, they may be physical assets or contracts. What they have in common is that they may produce or source a commodity as well as store, consume or sell it. They are a source or sink of the respective commodity in each point in time. When the commodity is sourced or consumed, a corresponding cash flow may be created.

A simple example: Assume you own a battery storage and you optimize it against the intraday power market. In this case you would add two assets to your portfolio the storage and the market.

The storage asset implements physical properties of your battery such as capacity, size, efficiency, own power consumption, etc. The market allows you to buy and sell power as you charge or discharge your battery. The terminology of calling the market 'your asset' is, of course, misleading. However, let us look at it from another perspective: As your trade in the market, you enter contracts with third parties. In that sense, the market is an asset that allows you to close specific contracts with third parties. Choosing the right asset class, you can implement anything from a simple approximation with given price curve (assuming you may be able to buy or sell at a forecasted price) and bid-ask spread to a more complex asset with a given order book (e.g. when optimizing in the intraday market).

In the portfolio, the storage is optimized against the market, following the restriction that any power that enters or leaves the battery, needs to come from or go to the market. When adding more assets such as wind or PV, the portfolio

[8] The documentation is available along with some samples at https://energyassetoptimization.github.io/EAO/

becomes more complex, but the mechanism remains the same. You simply add those assets to your portfolio.

Interaction only via dispatch. The starting point for the optimization framework is that single assets in a portfolio only interact via their dispatch, as the sum of all assets' dispatch at a given point in time must be zero for every location. A commodity is preserved and can only be changed by inflow (production or sourcing) or outflow (consumption or sale) of an asset.

When setting up the optimization problem as an LP, we can exploit this property and only need to translate an asset's characteristics into its own LP. The assets may then be combined into a portfolio by combining the single LPs. The restriction, that the sum of all assets' dispatches must be zero at each location (called node hence forward) is added.

Structured assets: Where complex assets, such as staged hydro reservoirs with several water reservoirs and water inflows, appear to be interlinked, they can be implemented as one asset with inner-asset restrictions. Alternatively, they can be implemented via several simple assets located in one or more nodes with directed connections. In our framework, we allow for this by packaging portfolios in one 'structured asset'. Such a structured asset may have a complex inner logic, but may be used just as other assets. An example is a green supply contract as described in section 3.

Multi-commodity problems. In the following we formulate the problem as a one-commodity problem. The extension to the multi commodity case is straightforward. Preservation of volumes is formulated per commodity and assets are defined to have inflow and outflow for one or several commodities. Only where an asset links one commodity to the other (e.g. by converting gas to power or power to heat), commodities are actually interlinked and we obtain what is often called sector coupling.

A.2 Formulation as an LP

A.2.1 LP for single assets

As above, the formulation of the single assets' LP is the starting point of the optimization. As described in section 2.2, we added several prototype assets to the EAO package, such as simple or complex contracts, storages and transport. Power plants are easily handled in LPs as long as their cost structure is convex. Start-up costs may require the usage of binary variables, which also fits into the described framework, but turning the LP into an MIP.

Notation: Let us denote the LP for an asset a as follows:

- The dispatch $\mathbf{x}^a \in \mathbb{R}^T$ (inflow or outflow of a commodity of the asset) is given as
$$\mathbf{x}^a = (x_1^a, x_2^a \dots x_T^a) \text{ for time steps } i = 1 \dots T. \tag{1a}$$
In some cases, formulating the objective function of the LP is simpler if we distinguish positive and negative dispatch at each point in time, i.e. we split x_i^a into two new new variables x_i^{a+} and x_i^{a-}, where $x_i^{a+}, x_i^{a-} \in \mathbb{R}$, $x_i^{a+} \geq 0$ and $x_i^{a-} \leq 0$ doubling the dimension of the original dispatch vector. For notational simplicity, we omit this fact and write $\mathbf{x}^a$ as a vector from $\mathbb{R}^T$. In the multi-commodity case, the dispatch may be defined per commodity.

- The asset internal variables are denoted by
$$\mathbf{y}^a = (y_1^a, y_2^a \dots y_{i^a}^a), \tag{1b}$$
i.e. $\mathbf{y}^a \in \mathbb{R}^{i_a}$. Such variables describe the internal state of an asset where necessary. For example, to model start-up processes, binary variables in addition to the dispatch may be required. The number of internal variables is denoted by i^a and will differ by asset type.

- Using the above definitions, the LP (or MIP) is formulated as
$$\max \mathbf{c}^{aT} \cdot \begin{bmatrix} \mathbf{x}^a \\ \mathbf{y}^a \end{bmatrix}^T \tag{1c}$$
s.t.
$$A^a \cdot \begin{bmatrix} \mathbf{x}^a \\ \mathbf{y}^a \end{bmatrix} \quad \leq \quad \mathbf{b}^a \text{ with } A^a \in \mathbb{R}^{(T+i^a)\times r} \text{ and } \mathbf{b}^a \in \mathbb{R}^r \tag{1d}$$
where A^a encodes the asset's structure and r denotes the number of asset restrictions. The vector $\mathbf{c} \in \mathbb{R}^{T+i^a}$ typically reflects costs from sourcing or revenues from selling the commodity as well as handling costs.

A.2.2 Combination into a portfolio

As outlined in section A.1, assets in energy or commodity trading only interact via their dispatch. Therefore, the optimization problem can be set up in four separate steps:

1. Build LPs for each single asset and combine them into one large LP
2. Assign each asset to one or several locations (nodes)
3. Add connections (transport) between nodes (implemented the same way as other assets)
4. Enforce dispatches to sum up to zero in each node for all time steps

Step 1: Build LP Specific LPs, describing different asset types, may differ substantially. In the presented framework, it is important to consistently separate dispatch variables $\mathbf{x}$ from internal variables $\mathbf{y}$, so dispatch variables may be tracked and joined in restrictions across assets.

The joint LP for a portfolio consisting of assets $1 \dots n$ is given by combining the LPs (1d) and (1c) for all assets. If written in matrix form, it reads:

$$\max \mathbf{c} \cdot \mathbf{x}^T \tag{2a}$$

s.t.

$$A \cdot \mathbf{x} \quad \leq \quad \mathbf{b} \tag{2b}$$

where

$$A = \begin{bmatrix} A^1 & 0 & \dots & 0 \\ 0 & A^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A^n \end{bmatrix} \tag{2c}$$

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{y}^1 \\ \vdots \\ \mathbf{x}^n \\ \mathbf{y}^n \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} \mathbf{b}^1 \\ \vdots \\ \mathbf{b}^n \end{bmatrix} \quad \text{and } \mathbf{c} = \begin{bmatrix} \mathbf{c}^1 \\ \vdots \\ \mathbf{c}^n \end{bmatrix} \tag{2d}$$

As you may note, the single assets' LPs are completely disjoint at this stage. The restriction matrix A only has (sub matrix) entries in its diagonal and the solution of the overall LP is the solution of every single LP.

Step 2: Assigning assets to nodes. Each asset must be located in one or several nodes. Power and gas are prominent examples, where those locations or nodes may be virtual trading hubs. We treat transport such as pipelines or connections as specific assets that are assigned to two nodes. Similar assets are used for sector coupling, e.g. converting gas to power or power to heat.

In the implementation, assignment to nodes consists in careful mapping, but is mathematically trivial. In order to be able to perform the next step, we need to be able to track type (dispatch or internal), point in time and node for each asset.

Step 3: Add nodal restrictions to LP Let there be nodes $N_1 \dots N_m$. The index collections $I_1 \dots I_m$ contain the indices of assets assigned to nodes $1 \dots m$. The portfolio LP is then extended by the following restrictions:

$$\sum_{a \in I_i} x_t^a = 0 \;\; \forall \text{ times } t = 1 \dots T \text{ and nodes } i = 1 \dots m \tag{2e}$$

Scaled assets: Up to now we assumed assets were given. In addition to the prototype assets introduced in section 2.2, we found that scaling assets may be very useful for many optimization tasks. In the example of a green portfolio we made heavy use of them.

Assume we can invest in or rent a flexible size or share of a given asset type. Investment or rent for a given time results in fix costs that are independent of the actual dispatch. The 'scaled asset' extends prototype assets by scaling them with an additional size variable.

Let S be the normalization of the base asset, s be the size variable and f be the fix costs, M (m) be the maximum (minimum) size. We define the constraints of the scaled asset, using the base asset's constraints by:

$$
\begin{aligned}
Ax - b\frac{s}{S} &\leq 0 && \text{(3a)}\\
x - \frac{u}{S}s &\leq 0 \text{ dispatch variables only} && \text{(3b)}\\
x - \frac{l}{S}s &\geq 0 \text{ dispatch variables only} && \text{(3c)}\\
\mathrm{m} \leq s &\leq \mathrm{M} && \text{(3d)}
\end{aligned}
$$

In this formulation we added u and l as upper and lower bounds for the base asset's dispatch variables, since they are often given. The cost vector needs to be appended with fix costs for the size variable. Scaling applies only to dispatch variables. How scaling applies to internal variables may differ from asset to asset, but for the implemented prototype assets we can simply scale the restrictions (here storage size and minimum or maximum offtake).

A.2.3 Optimizing and retrieving results

The portfolio's LP is defined by equations (2). The main contributions are coming from the specific LPs of each asset, which are glued together with the restriction that quantities at nodes must be preserved. This enables us to use a modular approach, where the definition of single assets can be separated from the setup of the overall structure of the portfolio including its locational characteristics.

As it turned out in many practical use cases, few asset classes (storage, contracts, production & demand) in combination with a specific nodal structure and transport are sufficient to cover a large variety of real life problems – from applications as different as decentral renewable power, sector coupling and global fuel sourcing. Once the generic framework is set up, portfolios can be changed flexibly without the need to touch code.

Once the solution of the overall LP is available, the resulting optimal values for asset variables $[\mathbf{x}^a, \mathbf{y}^a]$ may be passed back to assets to create a detailed view on their optimal dispatch or other behavior within the portfolio.

B Dealing with uncertainty using Stochastic Linear Programming

B.1 Theoretical background

The structure of the optimization problem in (2) is highly specific and uncertainty will enter the problem via costs or revenues (e.g. via market prices) or restrictions (e.g. reflecting demand or supply capacity).

Besides an efficient approach to manage, create and solve the LP itself, taking into account such uncertainties is a major challenge. The problem is naturally described as a stochastic linear program (SLP). We loosely follow the notation from [12]. We start by reformulating the energy asset optimization problem as a two stage SLP at time d. All steps up to time d need to be fixed now without uncertainty, beyond which they are uncertain. We write

$$\mathbf{x}^d = x_1, x_2, \ldots x_d \tag{4}$$
$$\hat{\mathbf{x}}^d = x_{d+1}, x_{d+2}, \ldots x_T \tag{5}$$

referring to the full set of variables of the portfolio LP. The original LP (without uncertainty) reads

$$\max \left[\mathbf{c}^T \mathbf{x}\right] \tag{6}$$
$$\text{with } A\mathbf{x} \leq \mathbf{b} \tag{7}$$

Let us further assume we approximate the uncertainty using a set of samples for $\hat{\mathbf{c}}^{ds}$, $\hat{A}^{ds}$ and $\hat{b}^{ds}$ with $s = 1 \ldots D$. Note that samples only differ in future values beyond time d. The we obtain

$$\max \left[\mathbf{c}^{dT} \mathbf{x}^d + \frac{1}{S} \sum_s \hat{\mathbf{c}}^{dsT} \hat{\mathbf{x}}^{ds}\right] \tag{8}$$
$$\text{s.t. } \hat{A}^s \begin{pmatrix} \mathbf{x}^d \\ \hat{\mathbf{x}}^{ds} \end{pmatrix} \leq \begin{pmatrix} \mathbf{b}^d \\ \hat{\mathbf{b}}^{ds} \end{pmatrix} \quad \forall s = 1 \ldots S \tag{9}$$

On the basis of sampling, the problem is exact. However, we observe the following:

- If we have a smart way of building the LP, we can easily build up the full SLP recursively from the two stage formulation
- For other than toy models, the problem quickly explodes in size. One way would be to use only a small number of *decision steps* that is not necessarily equal to the number of time steps in the discretization
- Samples need to come from a random process, where uncertainty in the future part of each step in the recursion has one sample of the present part in common. This is no restriction in our case

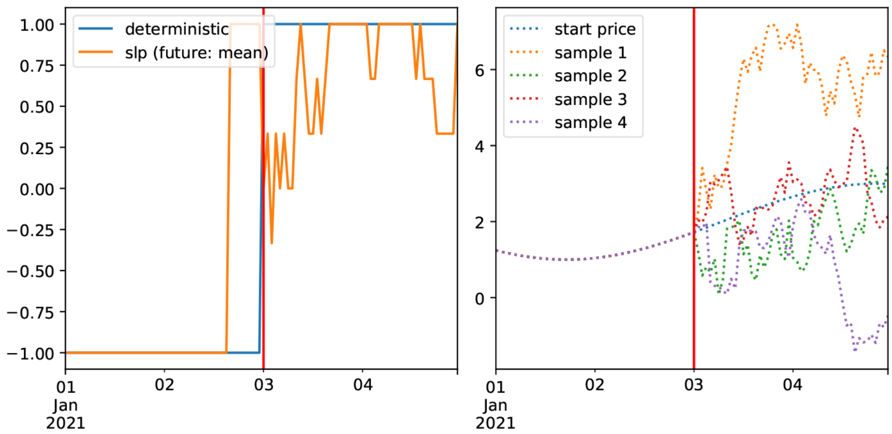

Figure 7: Simple example for a two-stage optimization. We optimize a storage against a given price curve. Until Jan 3 prices are given, for the future we have the 'start price' plus four price samples. While the deterministic optimization does a hard optimization against prices given today, SLP optimizes the dispatch until Jan 3 so the expected value in the future is maximized. Note that the future dispatch is optimized separately for each price sample.

B.2 Implementation and example

In figure 7 we show an example of a two-stage setup for a storage that is optimized in a two-stage approach. The red line indicates the start of the uncertain future (Jan 3). While for the pending decision (stage 1) prices are certain, for the future (stage 2) we have five price samples that encode the uncertainty. Solutions in this example are relatively similar – but note that the deterministic optimization charges the storage maximally until Jan 03 as future prices are higher than present prices. Since uncertainty is high, the SLP solution starts discharging earlier, as future prices may be forecasted today to be higher but remain uncertain.

As shown in the above mathematical formulation, uncertain future is optimized by modeling the future dispatch for each price sample. They have one present dispatch in common, that they are linked to via LP restrictions. In the figure we plot the mean future dispatch.

In our implementation we provide a formulation of a two-stage SLP generically for any portfolio structure. For the two-stage case, problems can mostly be handled and often already provide a useful approximation in practice. For the multi-stage case we believe that it will be necessary to resort to numerical approximations or strategies that are specific to the given portfolio. Machine learning methods could also be a promising approach to the problem. Providing a solution to the SLP for a relevant problem class could be a valuable research thread for further work.