QANTIC

QANTIC

Economic Optimization and Control of Distributed Energy Resources

Deep Reinforcement Learning as an Option for Automated Real-Time Operation

Thomas Kalitzky, Managing Director, Qantic, Berlin

Electrify Europe, 19 - 21 June 2018, Vienna

Using AI to optimize energy systems

Situation

The integration of large amounts of Renewable Energy Sources (RES) and other Distributed Energy Resources (DER) as well as technological and regulatory developments increase the requirements to

- React to **real-time signals** (e.g. of intraday markets, balancing systems or peer-to-peer-trading)
- Extract the valuable information from **high-frequency measurement data** (e.g. from smart meters)
- Deal with **complex and distributed systems** (e.g. hybrid microgrids)

Challenge

Conventional methods for optimization and dispatching can have problems to

- Model systems that exceed a certain level of complexity
- Ensure sufficient speed of calculation
- Fully extract the value of available data

How can Machine Learning help?

Agenda

- Overview
- Deep Reinforcement Learning
- Case Study

Machine Learning can help to reduce costs in dispatching complex energy systems

What is our task?	Dispatching all controllable energy resources of an energy system to serve a time-varying demand at least costs
What makes it difficult?	Complex system dynamics and intertemporal effects
How would we normally solve it?	Heuristics or techniques for deterministic or stochastic optimization (e.g. dynamic programming)
What can AI do better?	Optimizing energy systems with complex system dynamics at high computation speed using large amounts of input data

AI can yield benefits if energy systems exceed a certain level of **complexity**, **computation speed** is crucial or **large amounts of input data** (e.g. measurement data, forecasts) have to be processed

Agenda

- Overview
- Deep Reinforcement Learning
- Case Study

Combining Deep Learning with Reinforcement Learning

- Mixes two techniques from the domain of Machine Learning: Deep Learning and Reinforcement Learning
- A lot of recent AI-applications are based on Deep Reinforcement Learning (DRL)
- Reinforcement learning algorithms that incorporate deep learning can for example beat world champions at the game of Go

The idea is to train an agent to optimally control a system by making use of deep artificial neural networks

Agenda

- Overview
- Deep Reinforcement Learning
  - Reinforcement Learning
  - Deep Learning
- Case Study

An agent is trained to take actions on an environment in order to maximize a cumulative reward (or to minimize costs)

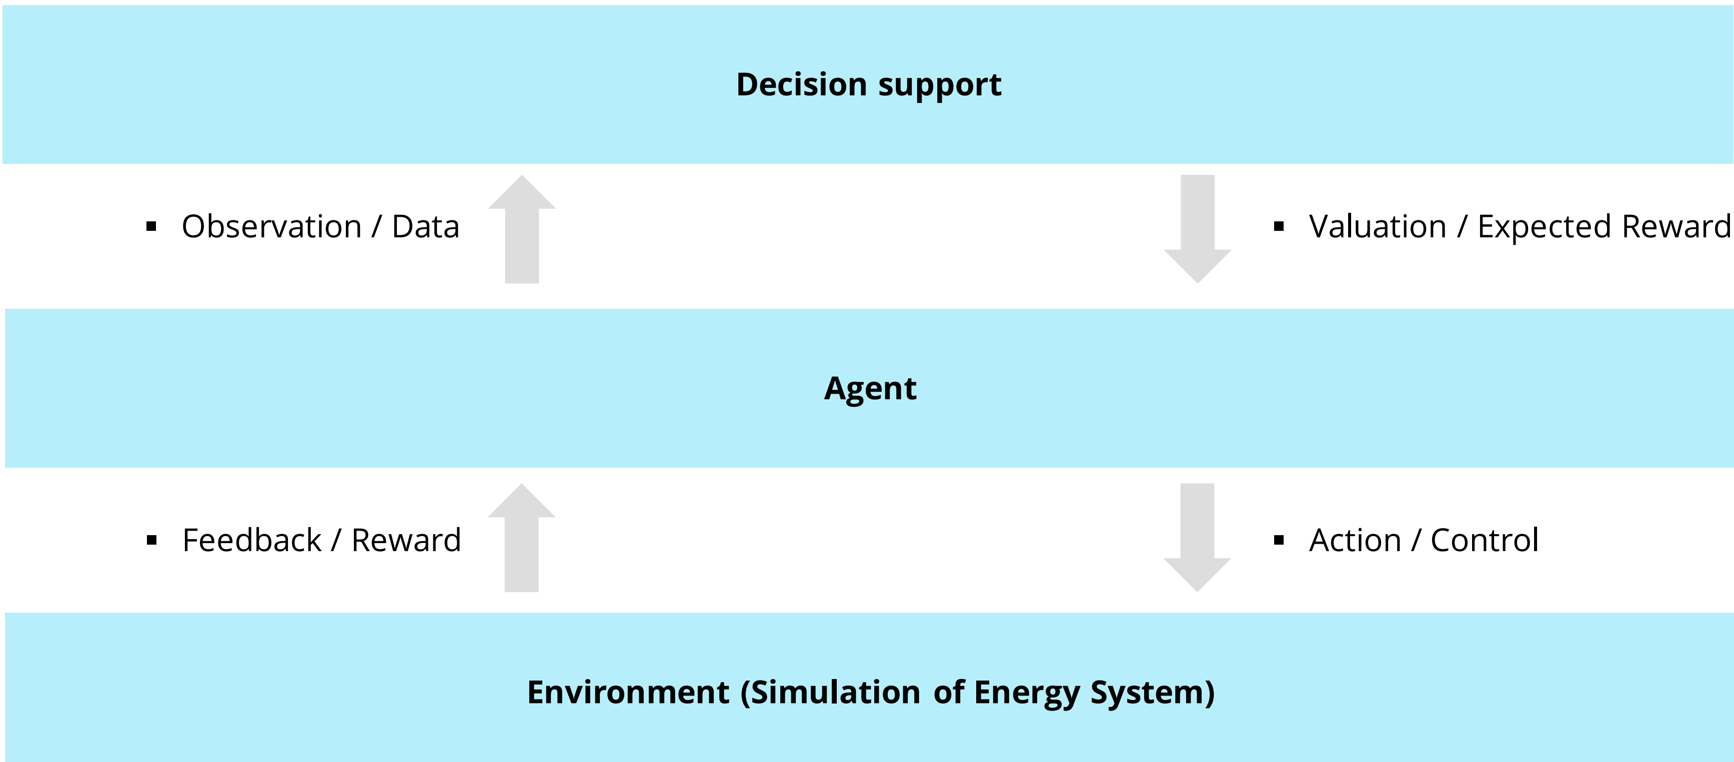

Agenda

- Overview
- Deep Reinforcement Learning
  - Reinforcement Learning
  - Deep Learning
- Case Study

Deep Learning is a class of Machine Learning algorithms

- Cascade of multiple layers of nonlinear processing units (Artificial Neural Network)
- Ability to learn hierarchy of concepts
- In our application supervised training on a large training data-set
- In DRL Deep Learning is used to provide the expected value for a certain action (e.g. start a plant) in a given state of the system (e.g. represented by all measurement data)

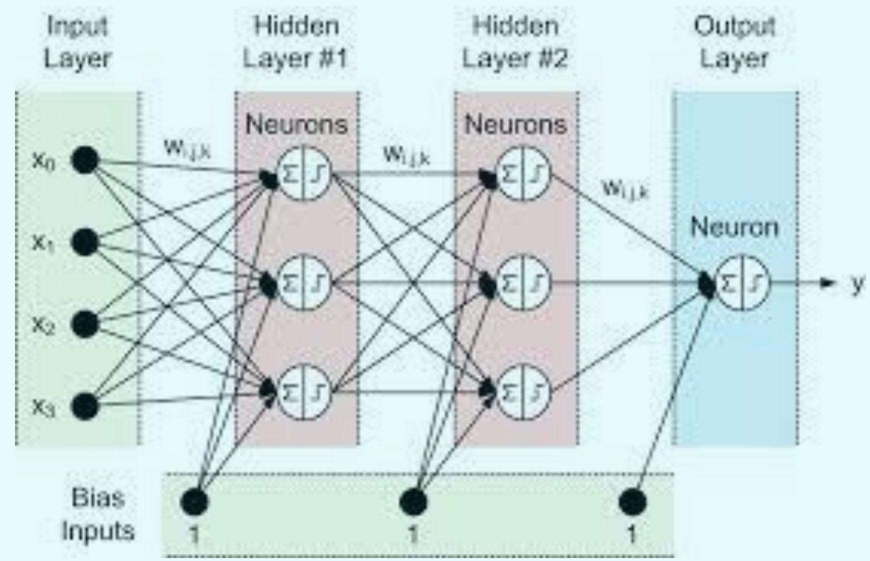

QANTIC

Agenda

- Overview
- Deep Reinforcement Learning
  - Reinforcement Learning
  - Deep Learning
- Case Study

A Simulator is used to dispatch the energy systems in a 50 Hz frequency (every 20 milliseconds) and faster

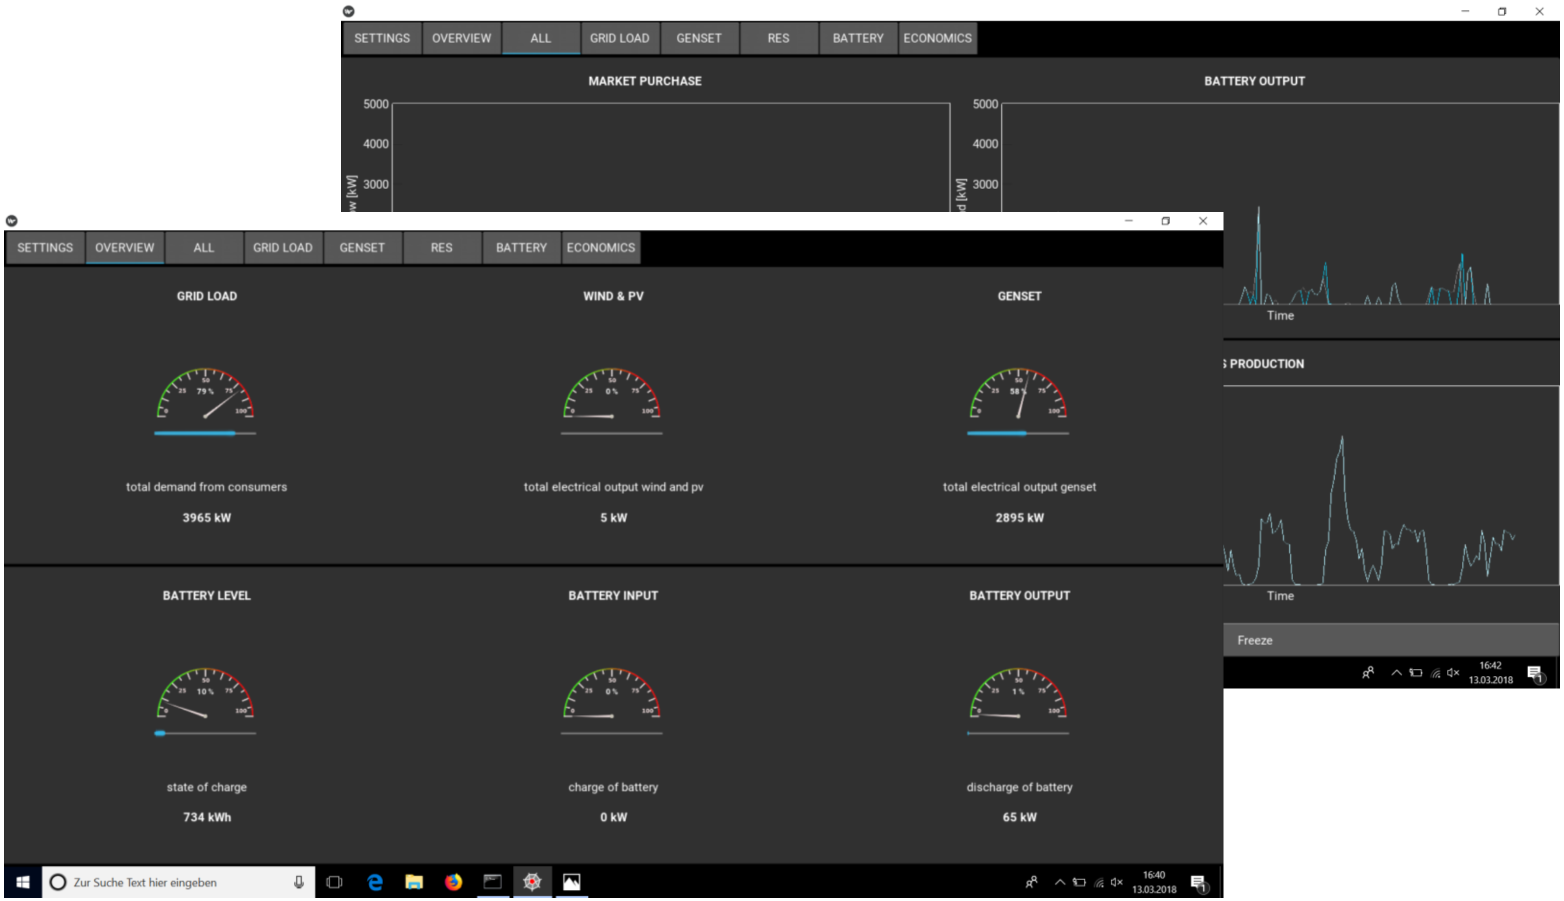

*) Dispatching at 50 Hz frequency is only possible for a heuristic approach or DRL. Conventional stochastic optimization techniques (e.g. Dynamic Programming) are too slow to fulfill that task at the same level of detail.

QANTIC

The system is dispatched by DRL and a heuristic and the results are compared: Costs are reduced by 16.9 % using DRL

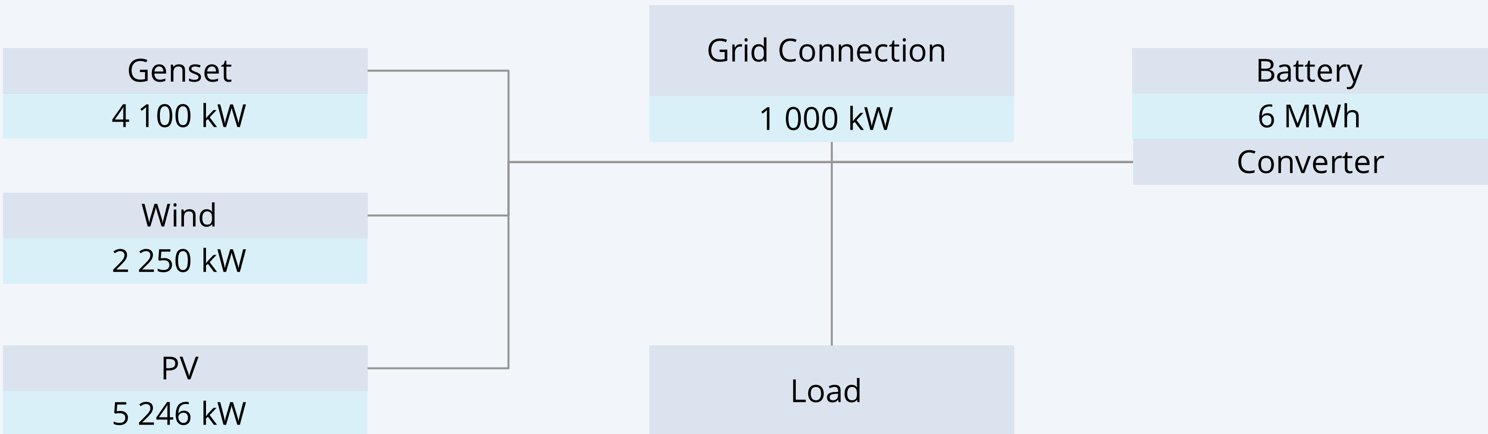

- Using input data for one representative year (load, weather, ...) the system is dispatched by a heuristic used in a standard microgrid planning software ("load following") as well as with a Deep Reinforcement Learning (DRL) model *)
- Total costs are calculated considering fuel costs, O&M and costs for replacement of genset and battery depending on operation regime
- Costs can be reduced by **16.9 %** using DRL

*) To be able to calculate the results for a representative period (a full year) we switched to a lower resolution than 20 ms

DRL has some properties that allow it to outperform simpler approaches

Foresighted system control	Not only considering present but also future costs → e.g. avoid having to start a plant in the future or having to run it at low efficiency
Extracting value from data	Recognizing typical patterns in the data → Get better foresight e.g. on demand and production and expected future costs
Dealing with uncertainty	Preparing for different possible future developments → Also considering events with low probability but high costs (e.g. events leading to fall-out)

Additional benefits can also result from using DRL in system planning and design

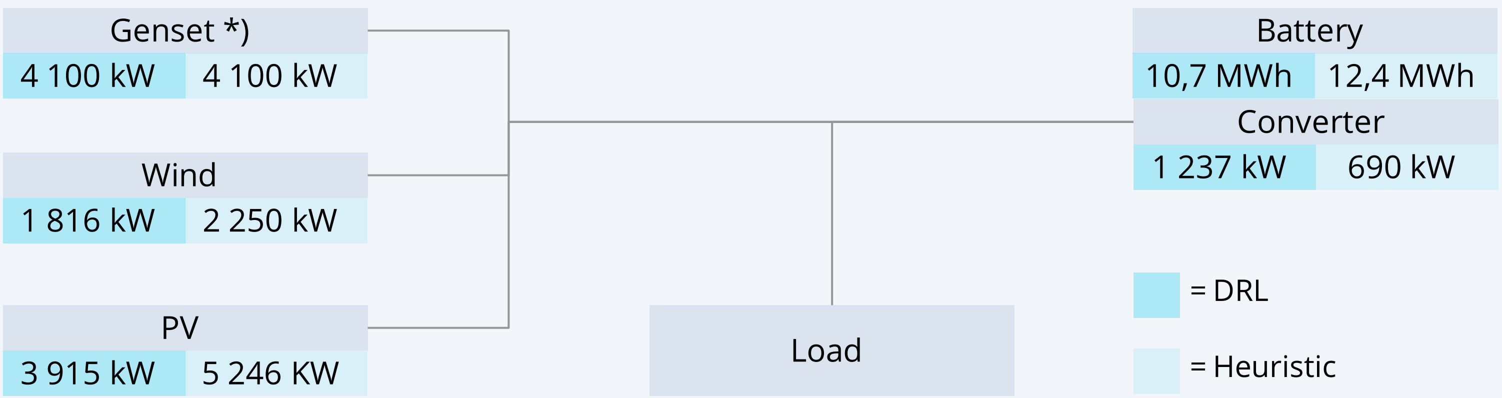

- Results for dimensioning of components differ significantly between heuristic and DRL
- Planning and operating using DRL reduces fuel costs and emissions **(-30%** compared to heuristic)
- Additionally the upfront investments can be reduced significantly by using DRL to support system design **(-20%** compared to heuristic)

*) Kept at the same level to ensure same level of peak production and security

DRL is a well-suited technique for optimization and control of energy systems

- Deep Reinforcement Learning (DRL) can significantly outperform simpler approaches (heuristics) and lead to lower total costs of electricity supply
- The use of DRL should already be anticipated in system planning, leading to other system designs and lower investments
- DRL is particularly a good choice when:
  - Real-time-applicability is needed
  - Systems exceed a certain degree of complexity
- This makes DRL well suited for
  - Smart Microgrids
  - Virtual Power Plants
  - Intraday optimization (high frequency optimization)
  - Peer-to-Peer-Trading and Blockchain-Applications

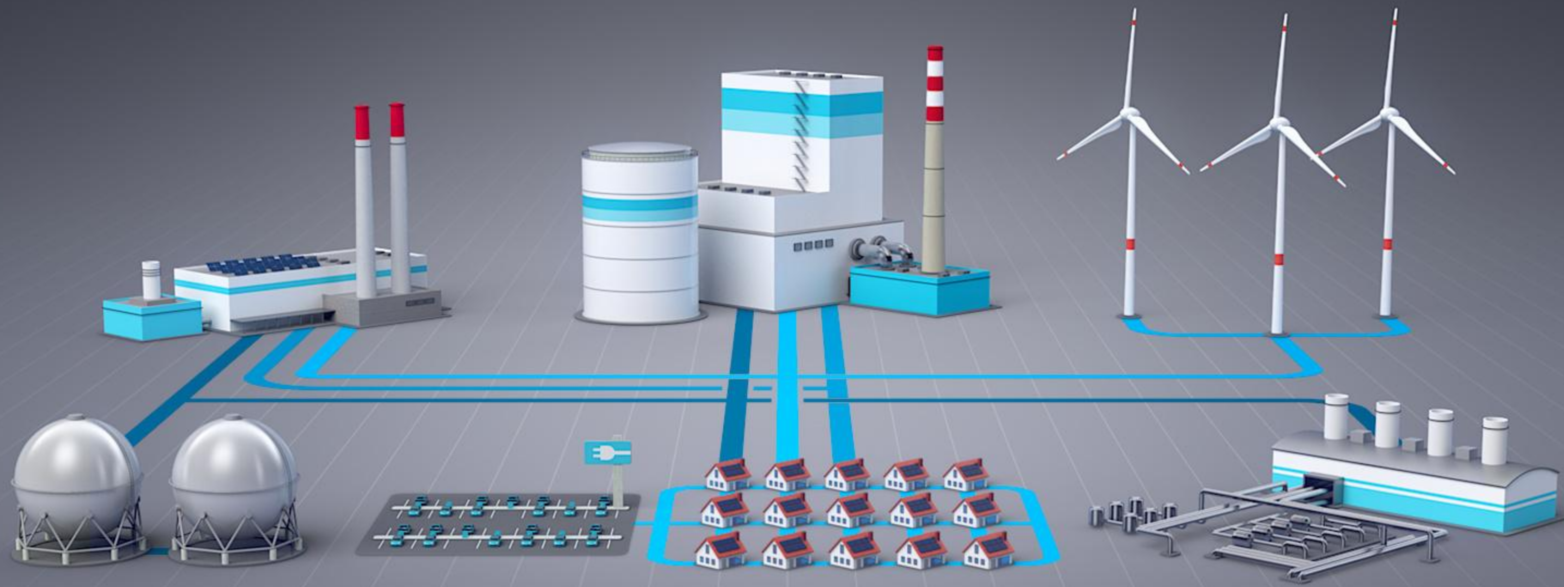

QANTIC

Boxhagener Straße 28
10245 Berlin
Phone +49 (0)30 892 099 03
Mobile +49 (0)163 611 9123
t.kalitzky@qantic.com
www.qantic.com

Agenda

- Overview
- Deep Reinforcement Learning
  - Reinforcement Learning
  - Deep Learning
- Case Study
- Appendix

Machine Learning combines the advantages of stochastic optimization and fast heuristic approaches

Heuristics	Deterministic Optimization	Stochastic Optimization (conventional)	Machine Learning (Deep Reinforcement Learning)
Straightforward approaches to control system by following a set of easy implementable rules →Suboptimal results →Easy to implement →High calculation speed	Full optimization of the system control, e.g. Dynamic Programming (DP) or Mixed Integer Linear Programming (MILP). Using singular forecasts →Neglecting uncertainty (suboptimal results) →Representation of system can be challenging	Same as deterministic optimization but using scenarios to represent uncertainties, e.g. stochastic DP or MILP approaches →Theoretically optimal →Usually simplifications needed to keep model computationally tractable (suboptimal results) →Low calculation speed	Stochastic optimization using AI with a freely choosable simulation model of the system →Lowest restrictions in the representation of the energy system (detailed modeling possible) →Very good coverage of stochastic aspects →High calculation speed

*) For more details see the manuscript provided with the conference proceedings and contact us to get more detailed information in the full appendix

"Q" names the function that returns the expected reward and stands for the "quality" of an action taken in a given state

Q (state, action)

Observations:

- State (e.g. load, production, forecasts, ...)
- Action (e.g. Discharging a battery, starting a plant, ...)

Valuation:

- Expected value of taking an action for the given state of the system

→In DRL the function Q (state, action) is represented by a deep artificial neural network

In Deep Reinforcement Learning an Artificial Neural Network is used to guide the decisions of the agent