Carl von Ossietzky Universität Oldenburg

Analyzing and Resiliently Operating Cyber-Physical Systems with Extended Deep Reinforcement Learning Architectures

Eric MSP Veith <eric.veith@uol.de> , 2023-09-27

% whoami

- Eric MSP Veith `<eric.veith@uol.de>`
- Currently head of a junior research group at University of Oldenburg, Germany
- Computer scientist by heart: First ICT, then distributed heuristics, then Multi-Agent Systems, now advanced Deep Reinforcement Learning
- PhD in 2017: "Universal Smart Grid Agent for Distributed Power Generation Management."
- Creator of the Adversarial Resilience Learning methodology (advanced DRL in CNIs)

The Origin of Artifical Intelligence

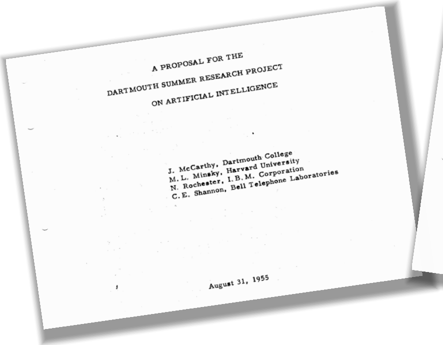

A PROPOSAL FOR THE
DARTMOUTH SUMMER RESEARCH PROJECT
ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

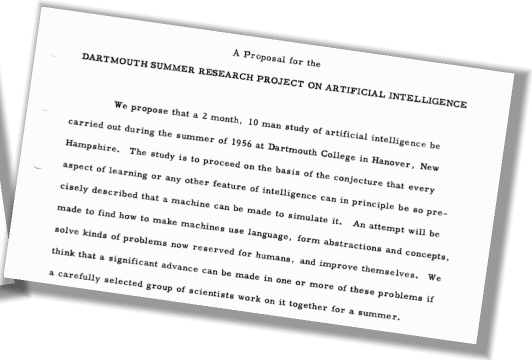

A Proposal for the
DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

Carl von Ossietzky Universität Oldenburg

The Origin of Artifical Intelligence

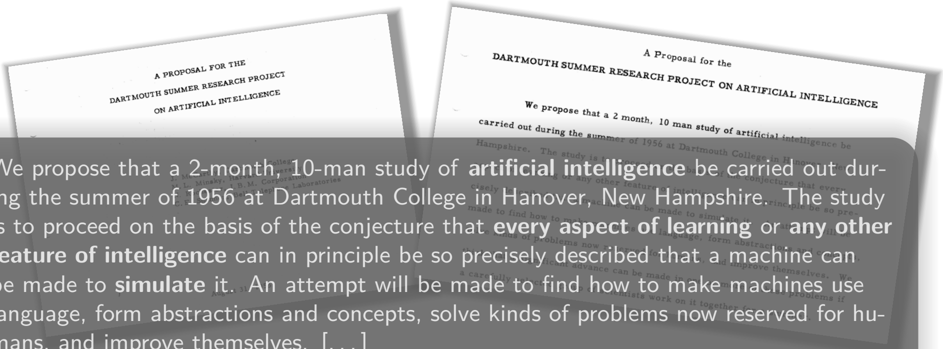
A PROPOSAL FOR THE
DARTMOUTH SUMMER RESEARCH PROJECT
ON ARTIFICIAL INTELLIGENCE

A Proposal for the
DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be
carried out during the summer of 1956 at Dartmouth College in Hanover, New
Hampshire.

We propose that a 2-month, 10-man study of **artificial intelligence** be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that **every aspect of learning** or **any other feature of intelligence** can in principle be so precisely described that a machine can be made to **simulate** it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. [. . .]

March 2016: **AlphaGo beats Lee Sedol**, Go world champion of 18 matches, in in **4 von 5 matches**.

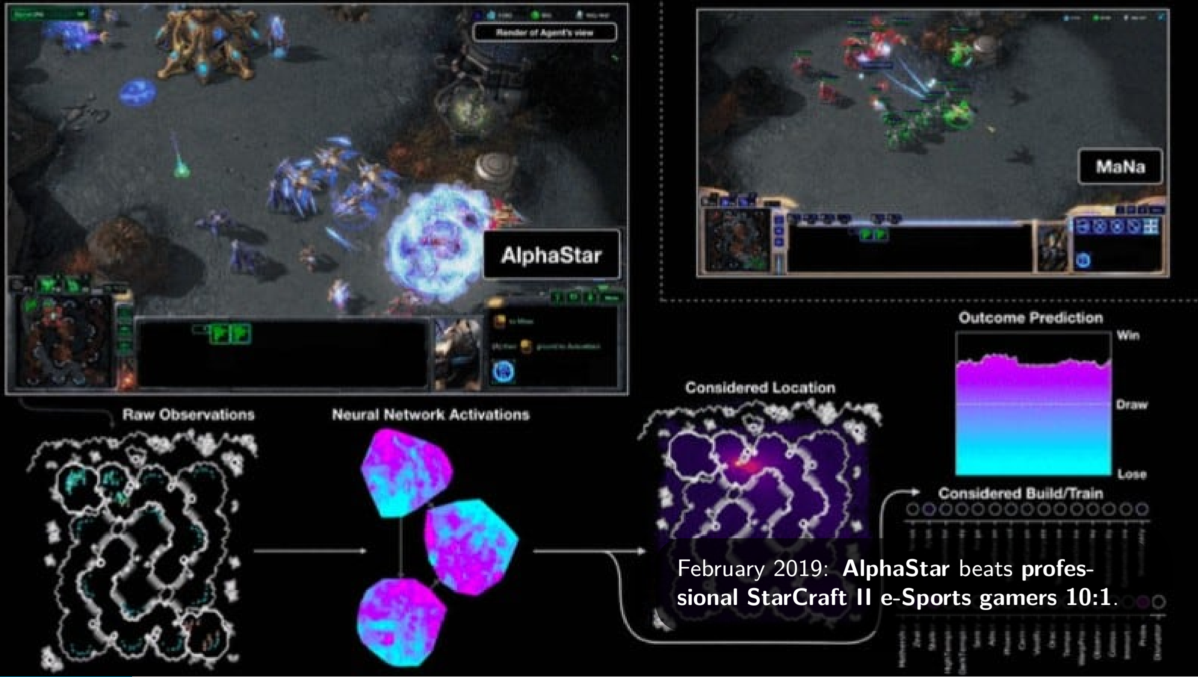

February 2019: **AlphaStar** beats **professional StarCraft II e-Sports gamers 10:1**.

On October 25th, 2018, Christie's auctions the painting "**Edmond de Belamy**" for **432.500 $**. It was created by Generative Adversarial Networks.

Worum geht es in dem Buch "O je, ein Zaubernieseln" von Anke de Vries?

Das Buch "O je, ein Zaubernieseln" von Anke de Vries handelt von einer Familie namens van der Steen, die in einer kleinen Stadt in den Niederlanden lebt. Die Geschichte wird aus der Sicht von Lotta, der jüngsten Tochter, erzählt.

Eines Tages findet Lotta eine kleine, seltsame Kreatur, die sie Nieseln nennt. Nieseln ist ein Zauberwesen, das Lotta und ihrer Familie Glück bringt und ihnen hilft, ihre Probleme zu lösen. Doch als Nieseln plötzlich verschwindet, müssen Lotta und ihre Geschwister alles daran setzen, ihn zu finden, bevor es zu spät ist.

Das Buch behandelt Themen wie Familie, Freundschaft und das Überwinden von Schwierigkeiten. Es ist für junge Leserinnen und Leser im Alter von 8 bis 12 Jahren geeignet und wurde mit dem niederländischen Kinderbuchpreis ausgezeichnet.

November 2022: OpenAI presents **Chat-GPT**. The model can generate natural language texts. New models are, as of now, developed by the week.

Carl von Ossietzky Universität Oldenburg

We like to diskuss the Dangers of AI

- “AI is the High-Interest Card of Technical Debt.” — Google, 2014.

Carl von Ossietzky Universität Oldenburg

We like to diskuss the Dangers of AI

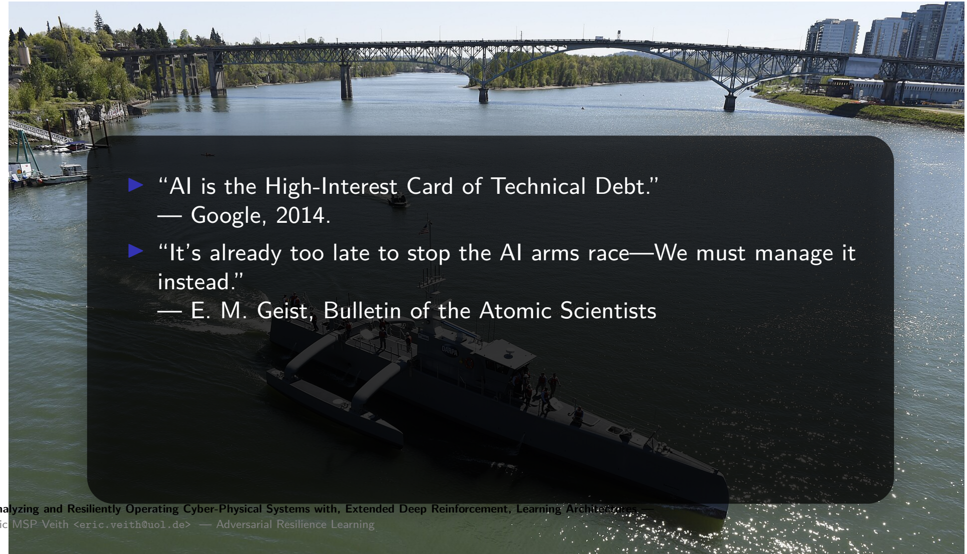

- “AI is the High-Interest Card of Technical Debt.”
  — Google, 2014.
- “It’s already too late to stop the AI arms race—We must manage it instead.”
  — E. M. Geist, Bulletin of the Atomic Scientists

We like to diskuss the Dangers of AI

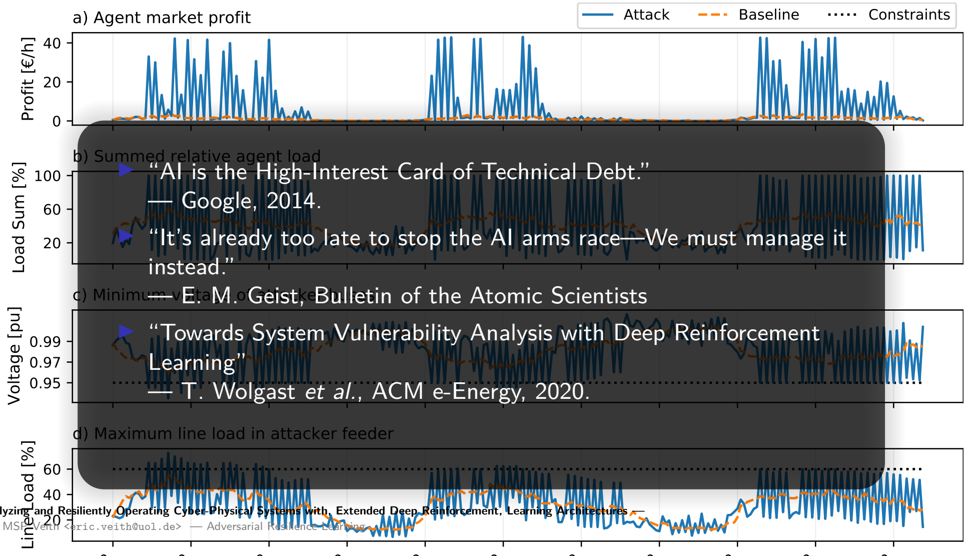

- “AI is the High-Interest Card of Technical Debt.”
  — Google, 2014.
- “It’s already too late to stop the AI arms race—We must manage it instead.”
  — E. M. Geist, Bulletin of the Atomic Scientists
- “Towards System Vulnerability Analysis with Deep Reinforcement Learning”
  — T. Wolgast *et al.*, ACM e-Energy, 2020.

We like to diskuss the Dangers of AI

- “AI is the High-Interest Card of Technical Debt.”
  — Google, 2014.
- “It’s already too late to stop the AI arms race—We must manage it instead.”
  — E. M. Geist, Bulletin of the Atomic Scientists
- “Towards System Vulnerability Analysis with Deep Reinforcement Learning”
  — T. Wolgast *et al.*, ACM e-Energy, 2020.
- “Pause Giant AI Experiments: An Open Letter”
  — Future of Life, März 2023.

We like to discuss the dangers of AI

Carl von Ossietzky Universität Oldenburg

Electricity Demand Rising

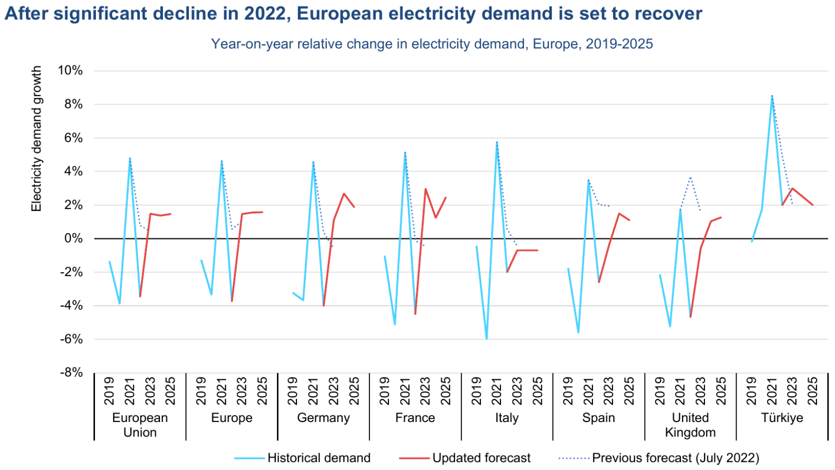

Renewables Are Replacing Fossil Fuels

Following two years of increases, CO_2 intensity starts to decline again from 2023 onward

Year-on-year change in electricity generation, European Union, 2019-2025

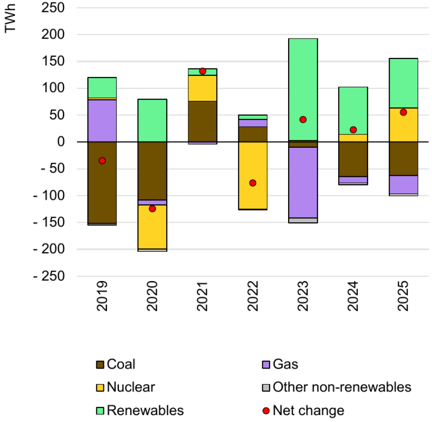

Development of average CO_2 intensity, Europe, 2019-2025

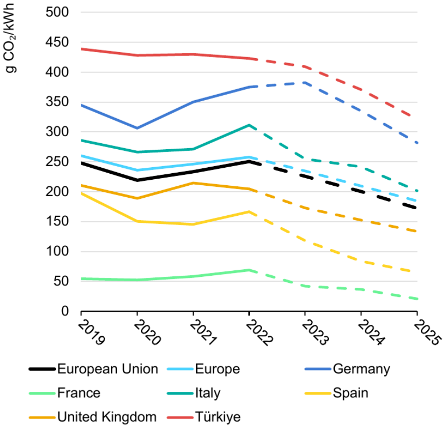

Carl von Ossietzky Universität Oldenburg

Electricity Demand + DERs = Grid Expansion

- DERs are volatile
- Consumer behavior becomes multi-faceted, no longer easy to average (direct market access, battery swarms, . . .)
- “Typical” grid usage largely atypical

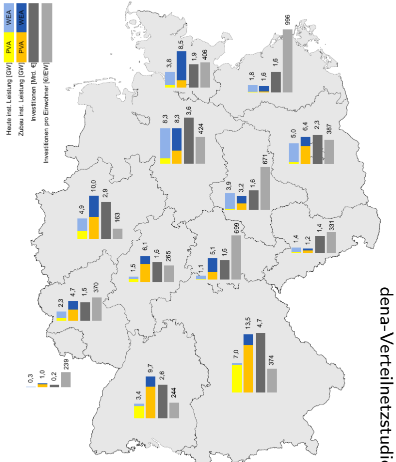

Copyright Kartengrundlage: Lutum+Tappert

dena-Verteilnetzstudie

Grid Expansion Planning Today

Step 1: Data Acquisition

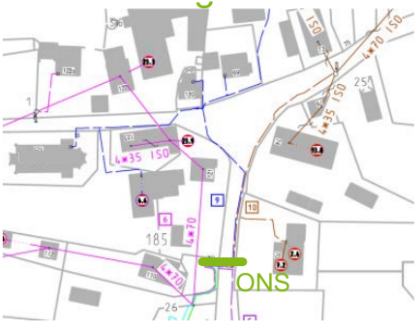

dena-Verteilnetzstudie

- Acquisition of GIS data
- List of local network stations
- Most probably load-flow calculation not possible;
- . . . so manual labor to create digital grid plan

Grid Expansion Planning Today

Step 2: Modelling

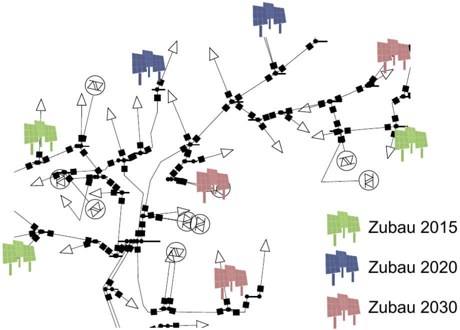

- Local network stations get prognosis assigned
- Stochastic distribution of DERs according to prognosis

dena-Verteilnetzstudie

Grid Expansion Planning Today

Step 3: Limits

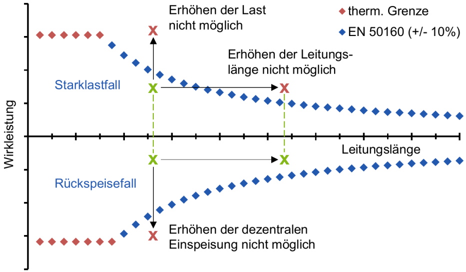

dena-Verteilnetzstudie

- Simulate, calculate
- Document limits & violations
- Mostly thermic limits, voltage band limits; also wear & tear of tap transformers

Grid Expansion Planning Today

Step 4: Expansion

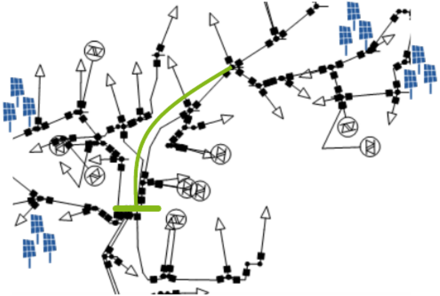

dena-Verteilnetzstudie

- Split subgrids: add local stations
- Split subgrids: add parallel lines
- Rinse & repeat

Grid Expansion Planning

- Modern grid expansion based on scenarios
- Data prognoses
- Modern tools allow to automatically propose expansions
- Calculate fault conditions, potential overloads, etc.
- But what scenarios are possible, which ones are unrealistic?
- **Expansion vs. efficienct operation** (?)

AI as Promise of an Alternative

- Multi-Agent Systems promise local, more more efficient grid operation
- Each node (subgrid, . . .) an agent
- Nodes (agents) forecast local power generation/consumption
- On disequilibrium, match forecasts to achieve equilibrium

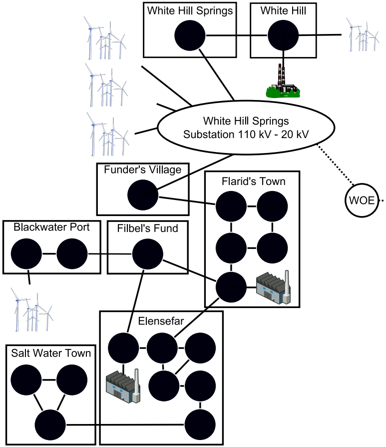

Agent Design

Input

- Grid
- Local Env.
- Training

Agent

- Demand—Supply
- Messaging
- Reserve
- Constraint Calculation
- Learner
- Forecaster
- Data Extraction
- Logging
- Hardware Interface

Automatic Hardware Control
(e.g., failsafe, emergency shutdown)

Priority

- Power Grid
- Micro Grid
- Local Unit
  (Power plant)
- Device Layer

Forecasting

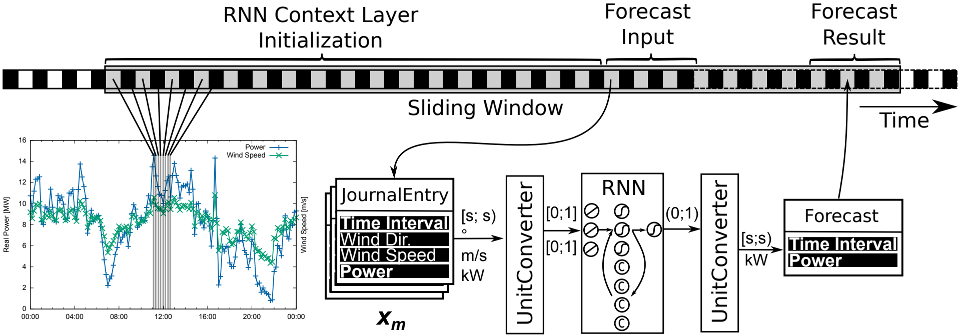

Carl von Ossietzky Universität Oldenburg

Forecasting

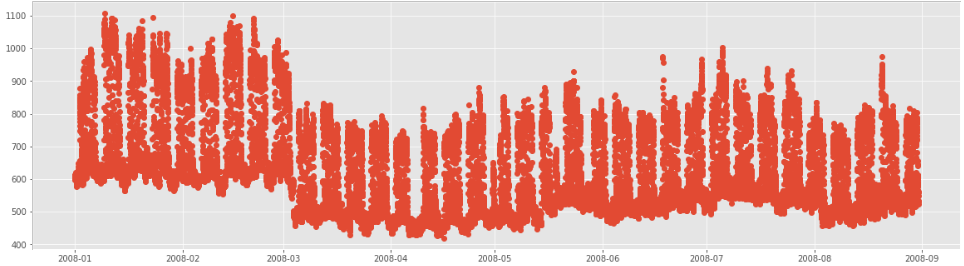

Forecasting

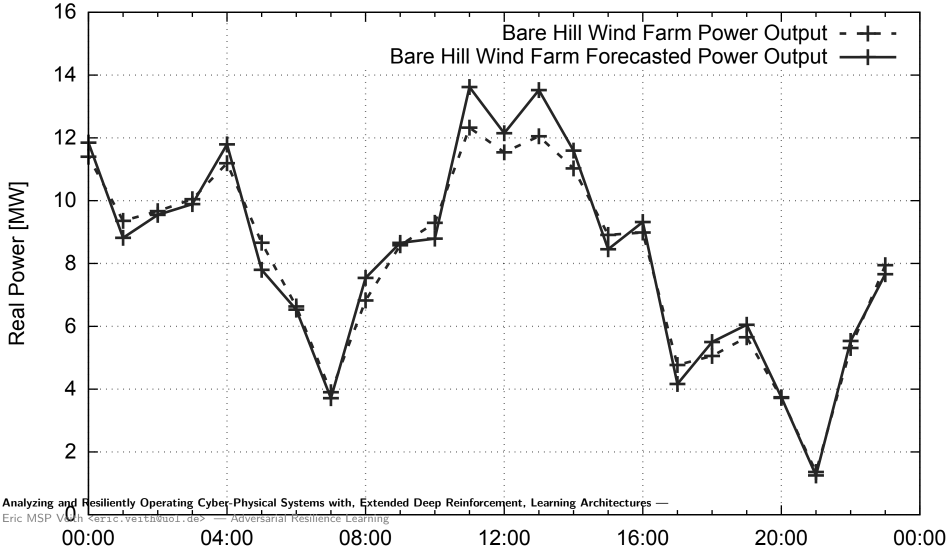

Training with Evolutionary Algorithms

- Evolutionary training algorithm
  - Each individual an **ANN candidate**
  - **Fitness** of an individual: the **cost function**
  - Moving through the search space by **mutation and crossover**
  - Advantage: possibly better to escape **local minima**
- A variant: **REvol**
  - Implicit gradient information
  - Dynamic reproduction probability density function

Carl von Ossietzky Universität Oldenburg

Properties of an Individual

- Parameter Vector: the **genome**
- Scatter Vector: limits **parameter modification** $p_{t,i} = [-s_i \cdot p_{t-1,i}, s_i \cdot p_{t-1,i}]$
- Time to Live (TTL)
- Fitness

Algorithm Behavior

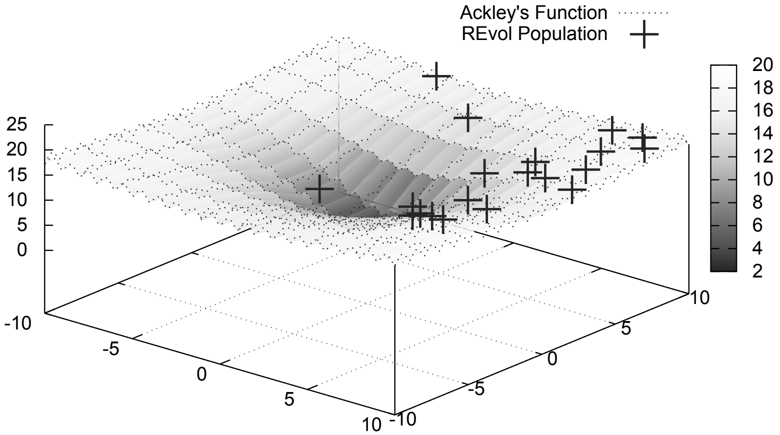

Algorithm Behavior

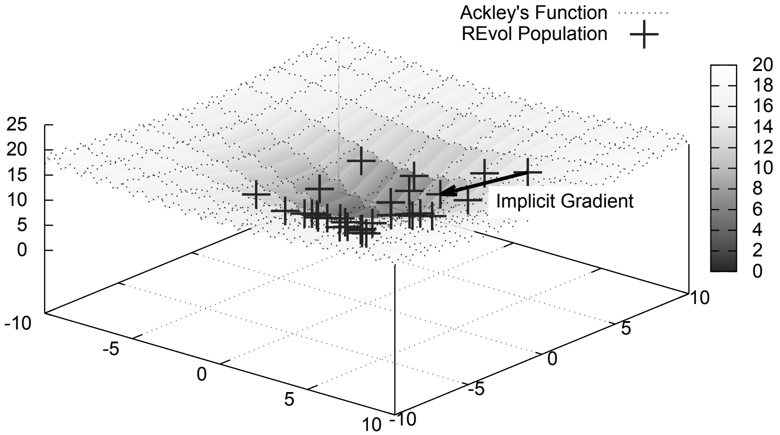

Algorithm Behavior

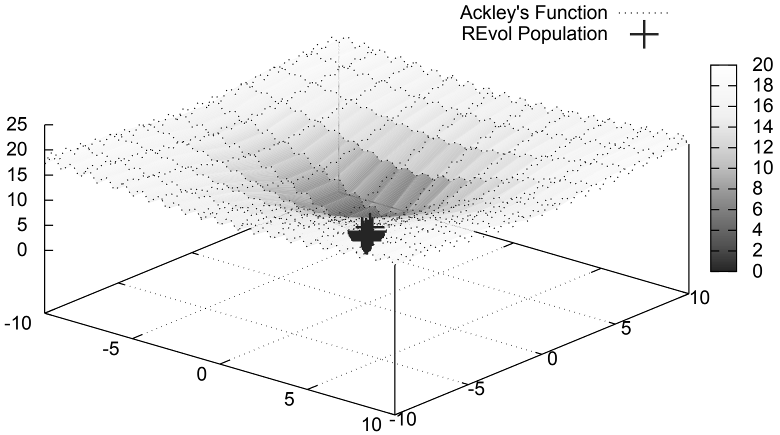

Carl von Ossietzky Universität Oldenburg

Algorithm Behavior

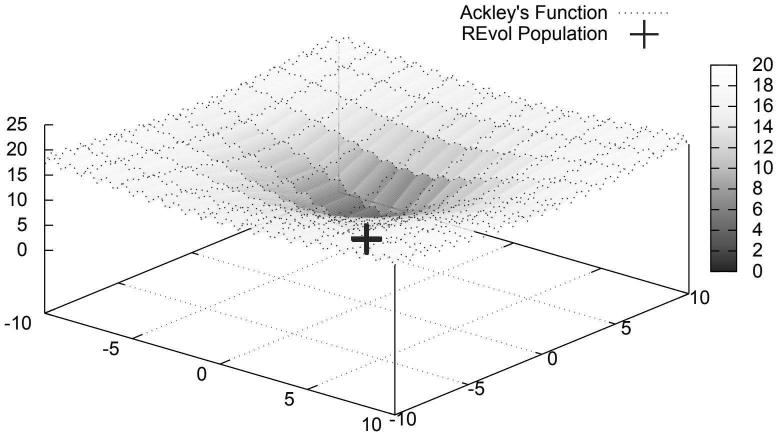

Algorithm Behavior

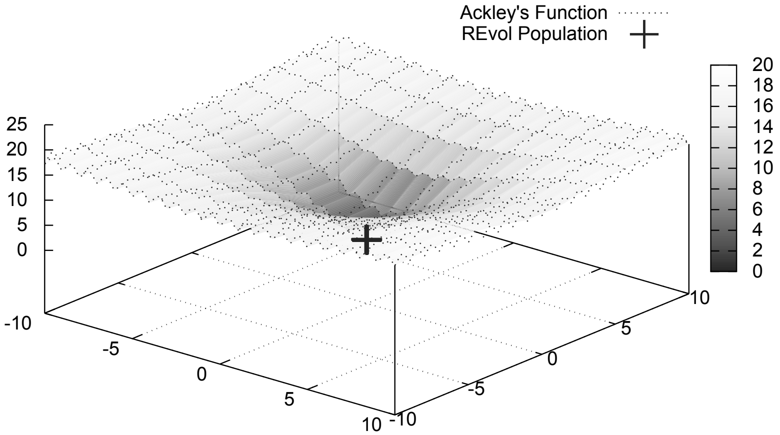

Algorithm Behavior

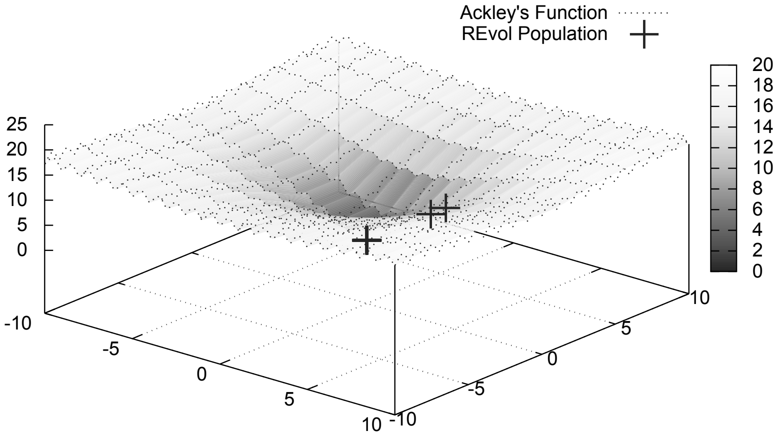

Algorithm Behavior

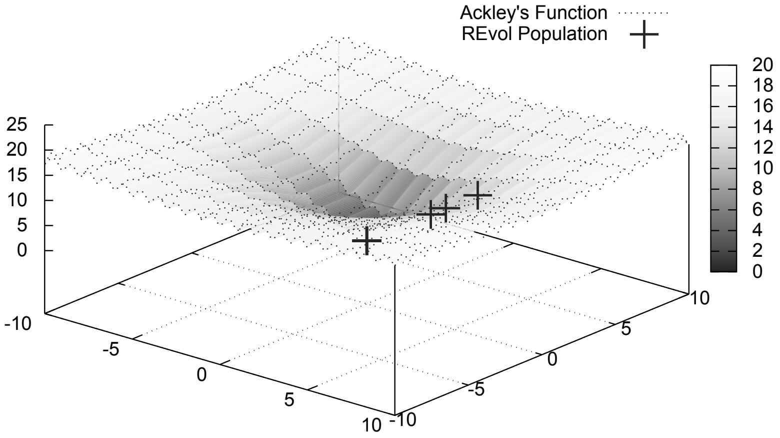

Four-Way Handshake

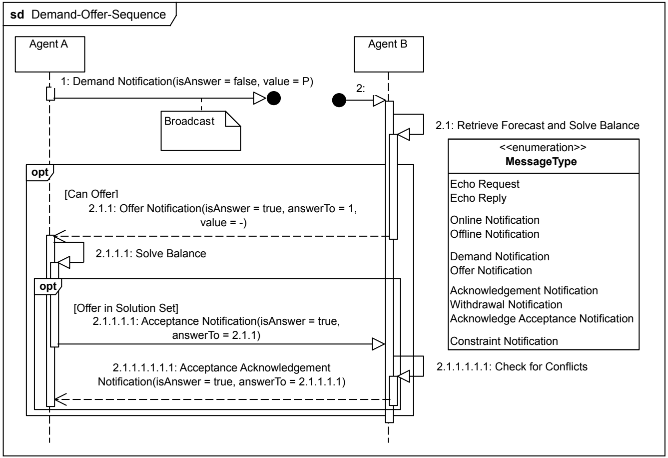

Carl von Ossietzky Universität Oldenburg

LPEP Routing

1. No Zombies
2. Match or Forward
3. Forwarding Ruleset

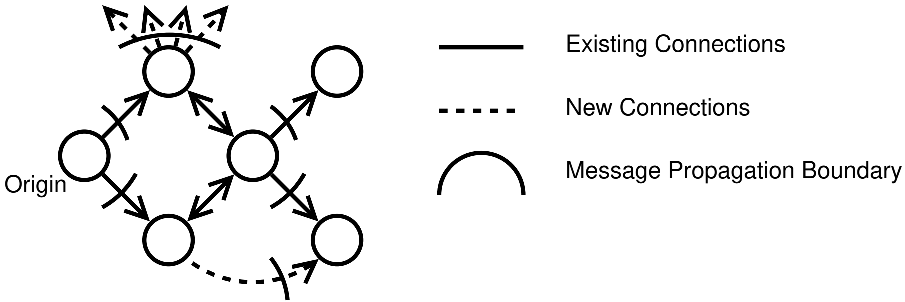

Forwarding

- L_i: *Links* of the i-th agent
- $l_{i,k}$: k-th link of i-th agent
- $\mathrm{distance}(l_{i,k})$: Distance metric
- m_j: j-th message
- M_i: Message Journal of the i-th agent

Forwarding

- L_i: *Links* of the i-th agent
- $l_{i,k}$: k-th link of i-th agent
- $\mathrm{distance}(l_{i,k})$: Distance metric
- m_j: j-th message
- M_i: Message Journal of the i-th agent

$$
\begin{aligned}
M_i = &\{m_1 \mapsto \{(l_{i,1}, m_{1,\mathrm{distance}(l_{i,1})}), \ldots, (l_{i,n}, m'_{1,\mathrm{distance}(l_{i,n})})\}, \\
&\ldots, \\
&m_n \mapsto \{(l_{i,1}, m_{n,\mathrm{distance}(l_{i,1})}), \ldots, (l_{i,n}, m'_{n,\mathrm{distance}(l_{i,n})})\}\}
\end{aligned}
$$

Forwarding

- L_i: *Links* of the i-th agent
- $l_{i,k}$: k-th link of i-th agent
- $\mathrm{distance}(l_{i,k})$: Distance metric
- m_j: j-th message
- M_i: Message Journal of the i-th agent

$$\begin{aligned}M_i =& \{m_1 \mapsto \{(l_{i,1}, m_{1,\mathrm{distance}(l_{i,1})}), \ldots, (l_{i,n}, {m'}_{1,\mathrm{distance}(l_{i,n})})\},\\ & \ldots,\\ & m_n \mapsto \{(l_{i,1}, m_{n,\mathrm{distance}(l_{i,1})}), \ldots, (l_{i,n}, {m'}_{n,\mathrm{distance}(l_{i,n})})\}\}\end{aligned}$$

$$l_{i,1}(t) \leq l_{i,2}(t) \quad \Leftrightarrow \quad l_{i,1,\mathrm{distance}}(t) \leq l_{i,2,\mathrm{distance}}(t)$$

Forwarding

1. Respect *Constraint Notifications*:
   1.1 No answer if $\min(M(m))$ a constraint notification to m, additionally
   1.2 send *Withdrawal Notification* iff already answered

Forwarding

1. Respect *Constraint Notifications*:
   1.1 No answer if $\min(M(m))$ a constraint notification to m, additionally
   1.2 send *Withdrawal Notification* iff already answered
2. $m_{isAnswer}$: forward on best connect $(\min(M(m_{answerTo})))$

Forwarding

1. Respect *Constraint Notifications*:
   1.1 No answer if $\min(M(m))$ a constraint notification to m, additionally
   1.2 send *Withdrawal Notification* iff already answered
2. $m_{isAnswer}$: forward on best connect $(\min(M(m_{answerTo})))$
3. *Selective Broadcast* for requests:
   3.1 Replace request with *Constraint Notification*, if necessary
   3.2 $M(m) = \emptyset$: forward on $|L| - 1$ links
   3.3 $m' = \min(M(m'))$: Update by fowarding
   3.4 Otherwise: no forwarding

How to Decide. . . ?

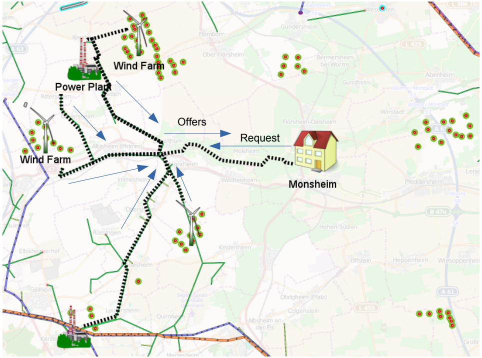

1. Local forecasting shows demand or oversupply of energy
2. Requests are sent
3. Other nodes make offers
4. Offers reach requestor
5. **Decision about offers?**

Power Balance Concept

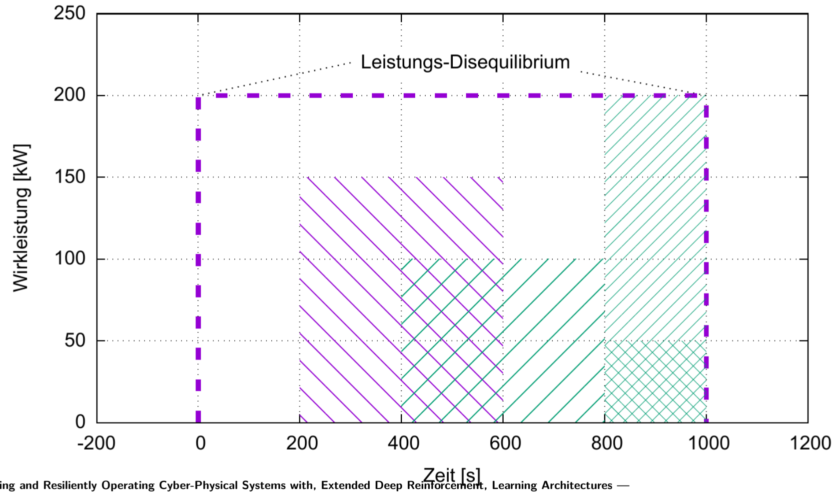

Set of Mappings $[t_1; t_2) \mapsto P$.

Problem Statement

'Power Balance Algebra':

$$\{[t_1; t_3) \mapsto P_1\} \cup \{[t_2; t_4) \mapsto P_2\} = \{[t_1; t_2) \mapsto P_1, [t_2; t_3) \mapsto P_1 + P_2, [t_3; t_4) \mapsto P_2\}\ , \quad (1)$$

$$[t_1; t_2) \mapsto P_1 \ \subseteq \ [t_3; t_4) \mapsto P_2 \quad \Leftrightarrow \quad t_1 \geq t_3 \ \wedge \ t_2 \leq t_4 \ \wedge \ P_1 \leq P_2\ ; \quad (2)$$

Problem Statement

'Power Balance Algebra':

$$\{[t_1;t_3) \mapsto P_1\} \cup \{[t_2;t_4) \mapsto P_2\} = \{[t_1;t_2) \mapsto P_1, [t_2;t_3) \mapsto P_1 + P_2, [t_3;t_4) \mapsto P_2\} \,, \tag{1}$$

$$[t_1;t_2) \mapsto P_1 \;\subseteq\; [t_3;t_4) \mapsto P_2 \quad \Leftrightarrow \quad t_1 \geq t_3 \;\wedge\; t_2 \leq t_4 \;\wedge\; P_1 \leq P_2 \,; \tag{2}$$

Distance Function:

$$\mathrm{d}(r_i) : r_i \mapsto \mathbb{R} \tag{3}$$

Problem Statement

'Power Balance Algebra':

$$\{[t_1; t_3) \mapsto P_1\} \cup \{[t_2; t_4) \mapsto P_2\} = \\ \{[t_1; t_2) \mapsto P_1, [t_2; t_3) \mapsto P_1 + P_2, [t_3; t_4) \mapsto P_2\} \; , \tag{1}$$

$$[t_1; t_2) \mapsto P_1 \;\subseteq\; [t_3; t_4) \mapsto P_2 \\ \Leftrightarrow \quad t_1 \geq t_3 \;\wedge\; t_2 \leq t_4 \;\wedge\; P_1 \leq P_2 \; ; \tag{2}$$

Distance Function:

$$\mathrm{d}(r_i) : r_i \mapsto \mathbb{R} \tag{3}$$

Problem Statement:

$$\sum_i b_i r_i \subseteq r_0 \;,\; i \neq 0, b_i \in \{0, 1\} \;, \tag{4}$$

$$\text{Subject to: } \min \sum_i b_i \mathrm{d}(r_i),\; i \neq 0, b_i \in \{0, 1\} \;. \tag{5}$$

Atomization

$$
\begin{aligned}
\boldsymbol{P} &= \left(|P_0|, |P_1|, \ldots, |P_i|, |P_C|\right) \ , \\
\boldsymbol{t} &= \left(t_{2,0} - t_{1,0}, t_{2,1} - t_{1,1}, \ldots, t_{2,i} - t_{1,i}\right) \ , \\
\Delta P &= \mathrm{ggT}(\boldsymbol{P}) \ , \\
\Delta t &= \mathrm{ggT}(\boldsymbol{t}) \ ,
\end{aligned}
$$

Carl von Ossietzky Universität Oldenburg

Atomization

$$
\begin{aligned}
\boldsymbol{P} &= \left(|P_0|, |P_1|, \ldots, |P_i|, |P_C|\right) \;, \\
\boldsymbol{t} &= \left(t_{2,0} - t_{1,0}, t_{2,1} - t_{1,1}, \ldots, t_{2,i} - t_{1,i}\right) \;, \\
\Delta P &= \mathrm{ggT}(\boldsymbol{P}) \;, \\
\Delta t &= \mathrm{ggT}(\boldsymbol{t}) \;, \\
x_{i,\tilde{t},\tilde{P}} &= \begin{cases} 1 & \text{if agent } i \text{ influences the grid in time-subinterval } \tilde{t} \text{ with} \\ & \text{power from the power-subinterval } \tilde{P}, \\ 0 & \text{else.} \end{cases}
\end{aligned}
$$

Atomization Illustrated

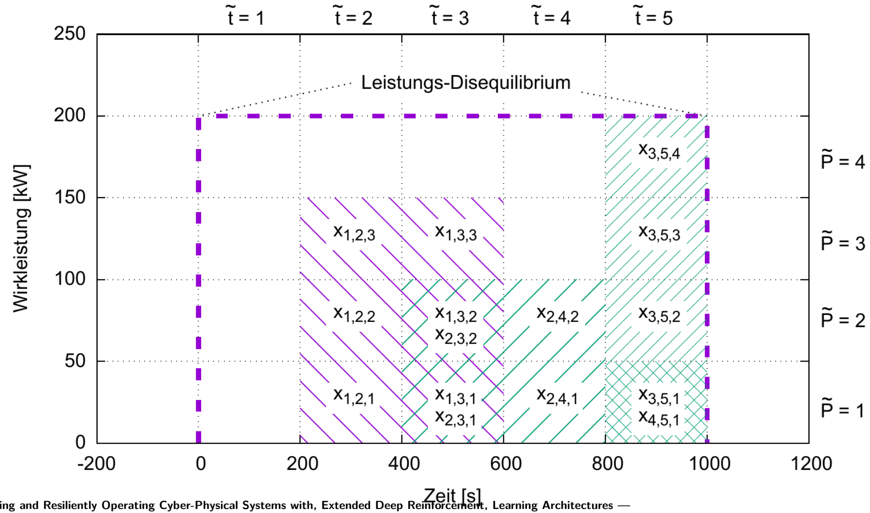

Model of the Disequilibrium

A symmetric function for each time-subinterval:

$$\mathrm{S}_k^n(\boldsymbol{x}_{i,\tilde{t}=k,\tilde{P}}) = \begin{cases} 1 & \text{if } n \text{ variables in } \boldsymbol{x}_{i,\tilde{t}=k,\tilde{P}} \text{ equal 1,} \\ 0 & \text{else;} \end{cases}$$

Full Disequilibrium:

$$S = \bigcap_{k=1}^{m} \mathrm{S}_k^n(\boldsymbol{x}_{i,\tilde{t}=k,\tilde{P}})$$

Modelling Responses

Acceptance Function:

$$\mathrm{r}_i(\boldsymbol{x}_{i,\tilde{t},\tilde{P}}) = \begin{cases} 1 & \text{if } \boldsymbol{x}_{i,\tilde{t},\tilde{P}} \text{ describes a valid interval for accepting the response of } i, \\ 0 & \text{else.} \end{cases}$$

$$\begin{aligned} \mathrm{r}_2(\boldsymbol{x}_{i,\tilde{t},\tilde{P}}) = {} & \bar{x}_{2,3,1} \wedge \bar{x}_{2,3,2} \wedge \bar{x}_{2,4,1} \wedge \bar{x}_{2,4,2} \\ & \vee x_{2,3,1} \wedge x_{2,3,2} \wedge \bar{x}_{2,4,1} \wedge \bar{x}_{2,4,2} \\ & \vee x_{2,3,1} \wedge x_{2,3,2} \wedge x_{2,4,1} \wedge x_{2,4,2} \end{aligned}$$

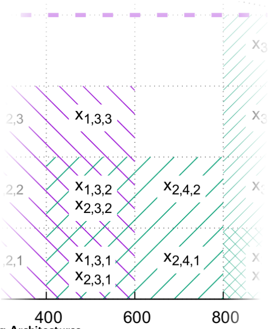

Equilibrium

$$S = \bigcap_{k=1}^{m} \mathrm{S}_k^n(\boldsymbol{x}_{\boldsymbol{i},\tilde{\boldsymbol{t}}=\boldsymbol{k},\tilde{\boldsymbol{P}}})$$

$$R = \bigcap_{i \in I', \tilde{t}, \tilde{P}} r_i(\boldsymbol{x}_{\boldsymbol{i},\tilde{\boldsymbol{t}},\tilde{\boldsymbol{P}}}) \ ,$$

$$C = S \cap R \ .$$

Equilibrium

$$S = \bigcap_{k=1}^{m} \mathrm{S}_k^n(\boldsymbol{x}_{\boldsymbol{i},\tilde{\boldsymbol{t}}=\boldsymbol{k},\tilde{\boldsymbol{P}}})$$

$$R = \bigcap_{i \in I', \tilde{t}, \tilde{P}} r_i(\boldsymbol{x}_{\boldsymbol{i},\tilde{\boldsymbol{t}},\tilde{\boldsymbol{P}}}) \ ,$$

$$C = S \cap R \ .$$

- Best solution through ordering: $r_i \leq r_{i'} \quad \Leftrightarrow \quad \mathrm{d}(r_i) \leq \mathrm{d}(r_{i'})$
- Generating next vector in S through permutation
- Exploiting the commutative property of the intersection operator: $R_n \cap (\ldots \cap (R_2 \cap (R_1 \cap S)))$

Efficiency

Data Effect $\kappa = \frac{W}{D} \left[\frac{\mathrm{kWh}}{\mathrm{kB}}\right]$

Data Efficiency $\xi = \frac{\Delta P}{D} \left[\frac{\mathrm{kW}}{\mathrm{kB}}\right]$

Carl von Ossietzky Universität Oldenburg

Comparison

Comparison with BDD approach by Inoue *et al.* (2014):

	BDD	Universal Agent
Loss Avoided (ΔP)	17 208 kW	17 208 kW
Runtime	$>$ 16 min	$<$ 11 min (simulated)
D	100 MB	28.9 MB
ξ	0.168 kW/kB	0.581 kW/kB

Universal Agent Efficiency

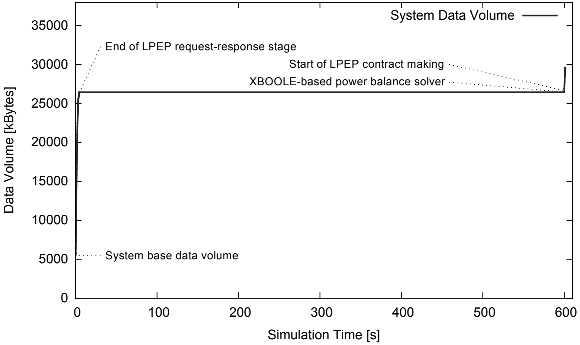

- BDD approach in low-load situation: 100 kB
- *Universal Agent* concept especially useful in complex load situations

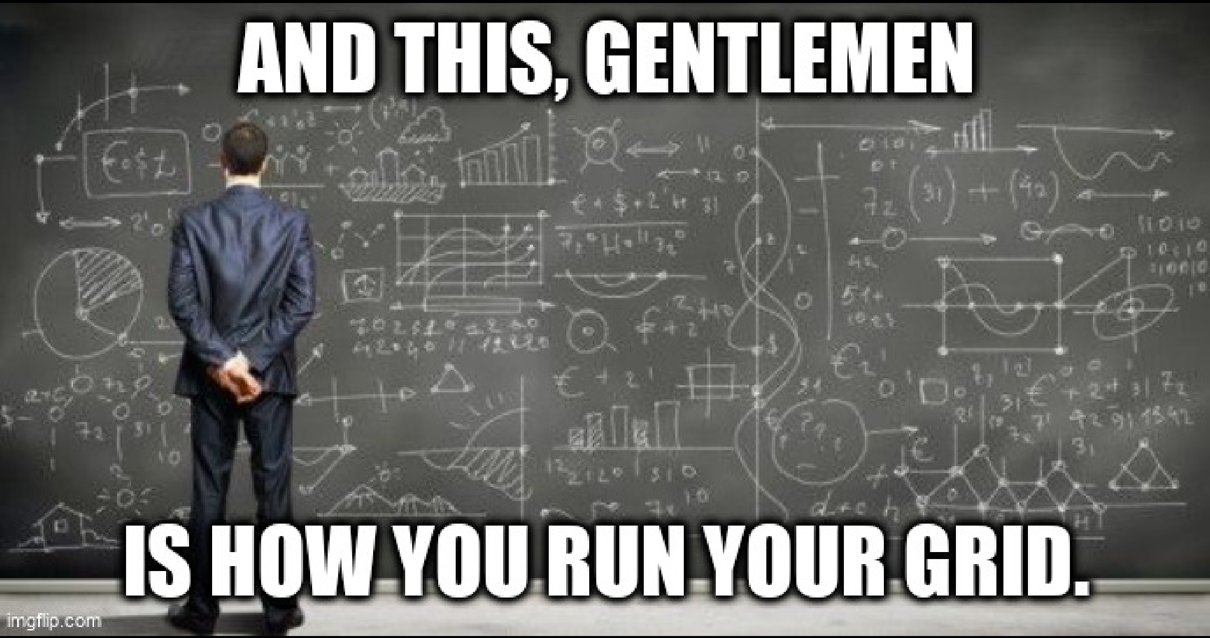
AND THIS, GENTLEMEN
IS HOW YOU RUN YOUR GRID.
imgflip.com

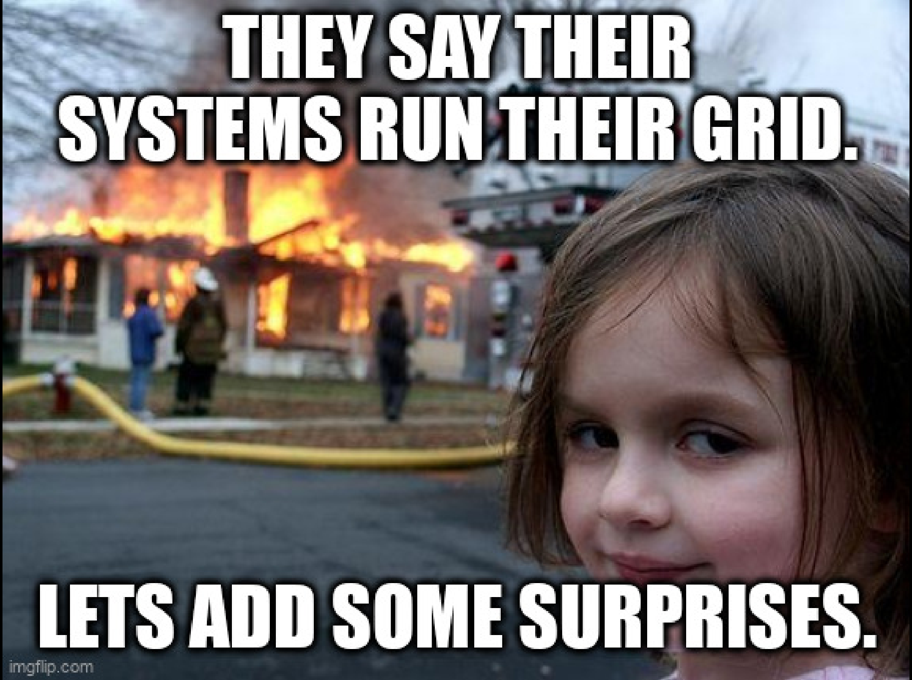
THEY SAY THEIR SYSTEMS RUN THEIR GRID.
LETS ADD SOME SURPRISES.
imgflip.com


"There are only two types of companies:
those who have been hacked,
and those who don't yet know
they have been hacked."
— John T. Chambers

Energy Systems Fit The Bill Just As Well

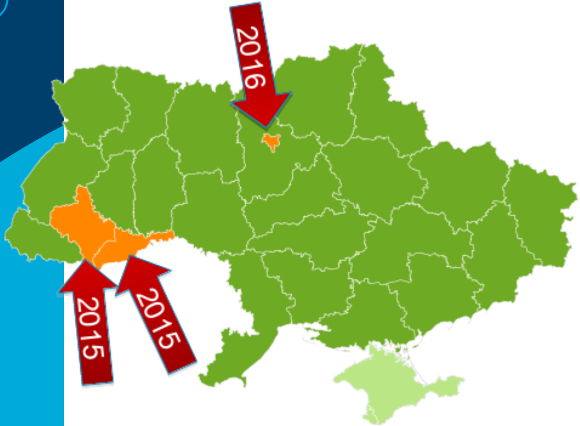

Dec 23rd, 2005

- **Cyber attack** causes blackout in the Ukraine
- **3 DSOs** targeted
- **High level of automation** helps attackers
- Operative intrusion in **OT**; disconnection of **several substations**
- Several months in preparation

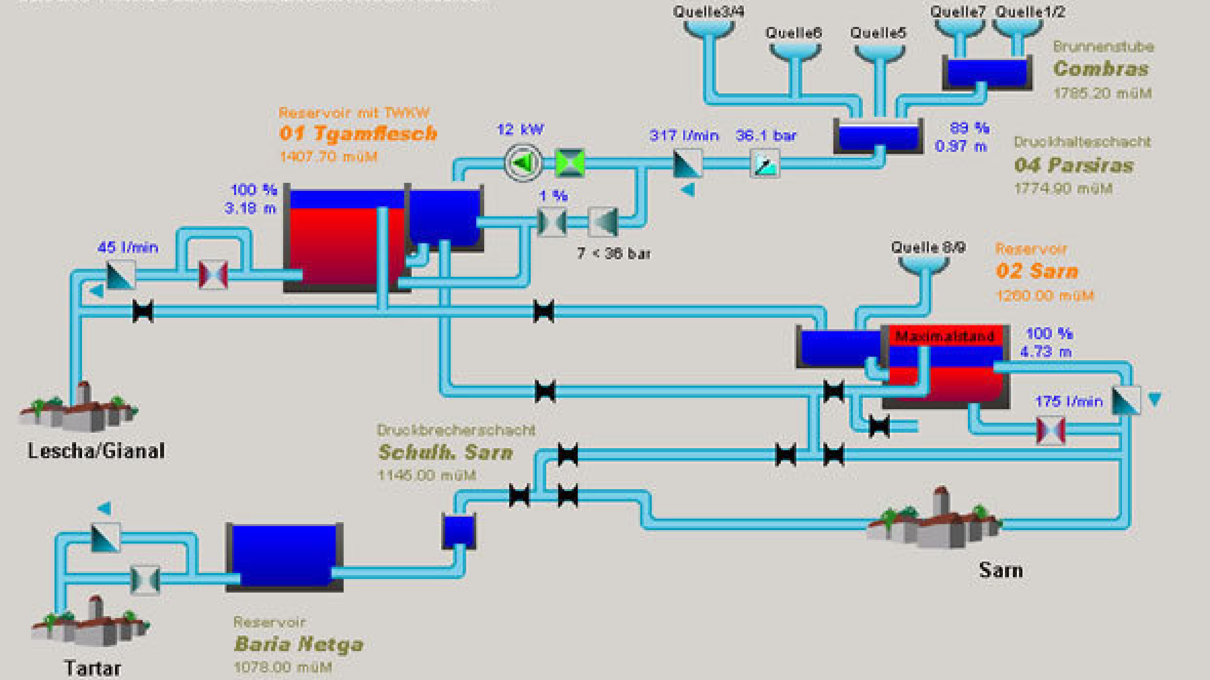

Quelle3/4
Quelle6
Quelle5
Quelle7
Quelle1/2
Brunnenstube
Combras
1785.20 müM
Reservoir mit TWKW
01 Tgamflesch
1407.70 müM
12 kW
317 l/min
36.1 bar
89 %
0.97 m
Druckhalteschacht
04 Parsiras
1774.90 müM
100 %
3.18 m
1 %
7 < 36 bar
45 l/min
Quelle 8/9
Reservoir
02 Sarn
1260.00 müM
Maximalstand
100 %
4.73 m
175 l/min
Lescha/Gianal
Druckbrecherschacht
Schulh. Sarn
1145.00 müM
Sarn
Tartar
Reservoir
Baria Netga
1078.00 müM

Energieversorgung in Deutschland

Stromhändler zocken fast bis zum Blackout

Der deutsche Strommarkt stand in den vergangenen Tagen mehrfach vor dem Zusammenbruch. Laut Bundesnetzagentur waren dafür aber nicht die Kälte oder der Atomausstieg verantwortlich, sondern Energiehändler - die offenbar ihre Profite maximieren wollten. Die Aufsichtsbehörde ist alarmiert.

dapd

Stromnetz in Deutschland: Starke Preisschwankungen durch hohe Nachfrage

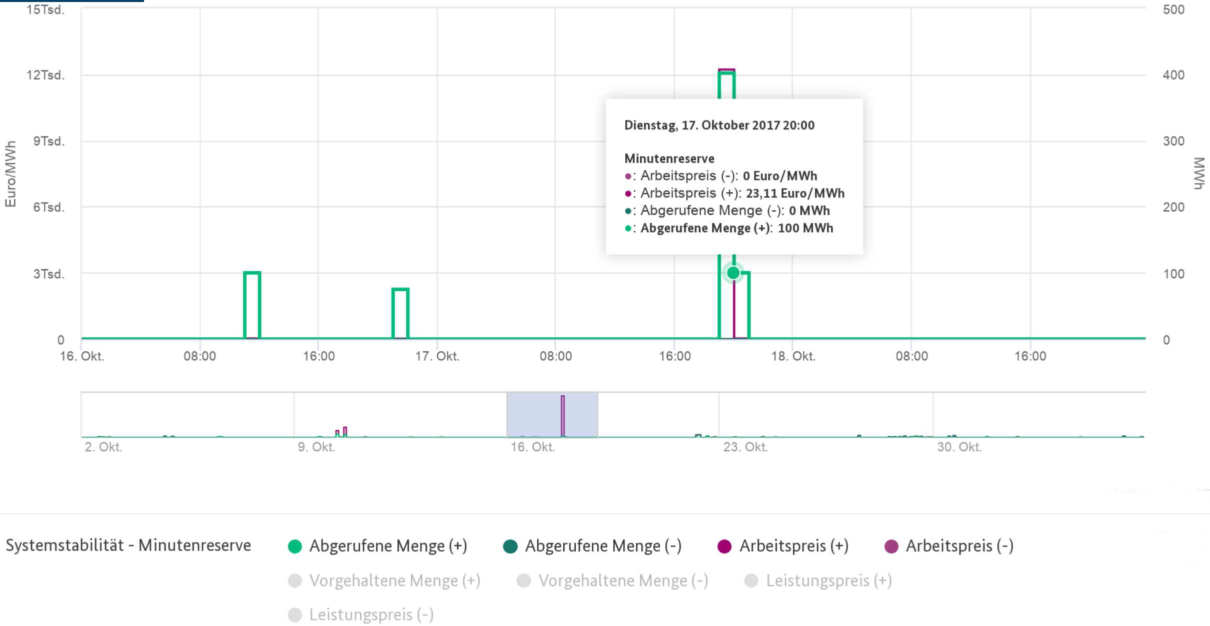

Dienstag, 17. Oktober 2017 20:00
Minutenreserve
●: Arbeitspreis (-): 0 Euro/MWh
●: Arbeitspreis (+): 23,11 Euro/MWh
●: Abgerufene Menge (-): 0 MWh
●: Abgerufene Menge (+): 100 MWh
Euro/MWh
MWh
15Tsd.
12Tsd.
9Tsd.
6Tsd.
3Tsd.
0
500
400
300
200
100
0
16. Okt.
08:00
16:00
17. Okt.
08:00
16:00
18. Okt.
08:00
16:00
2. Okt.
9. Okt.
16. Okt.
23. Okt.
30. Okt.
Systemstabilität - Minutenreserve
Abgerufene Menge (+)
Abgerufene Menge (-)
Arbeitspreis (+)
Arbeitspreis (-)
Vorgehaltene Menge (+)
Vorgehaltene Menge (-)
Leistungspreis (+)
Leistungspreis (-)

Carl von Ossietzky Universität Oldenburg

The Adversary

- Consumer behavior (prosumers), VPP, outages, weather effects: probabilistic modelling
- **DERs**
  - Prognosis deviations
  - VPP, direct marketing: highly non-deterministic
- **Terrorist**
  - Goal: Demolition
  - No route back needed in som cases
  - No sophisticated tactics necessary
- **Military**
  - Goal: destruction & takeover
  - Damage to CNIs is mostly collateral damage (or explicitly wanted)
  - Usually, CNIs are "don't care," but should be usable afterwards
- **Businnesspeople**
  - Goal: (short-term) profit maximization
  - Damage to CNIs unexpected (or simply "don't care")
  - Usues loopholes and grey areas in codices

Carl von Ossietzky Universität Oldenburg

Learning Resilient Control

- **Interconnected CPS have always attack surface due to their inherent complexity**
- Low latency of ICT and OT
- High interdependence
- Complexity in breadth and depth
- Cricital Services as SPOF (DNS, BGP, SCADA, SDL)
- **Learning Stratgies for automatic issue mangement**
- "Adversarial Resilience Learning"

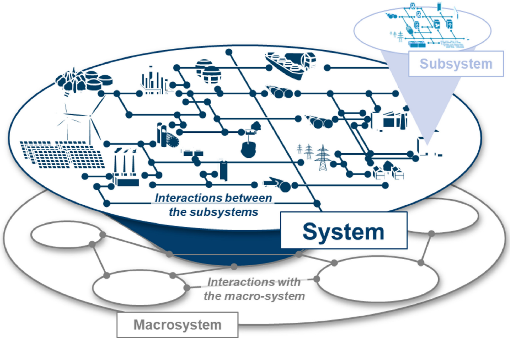

Kotzur, Leander, et al. "A modeler's guide to handle complexity in energy systems optimization." Advances in Applied Energy 4 (2021): 100063.

Adversarial Resilience Learning

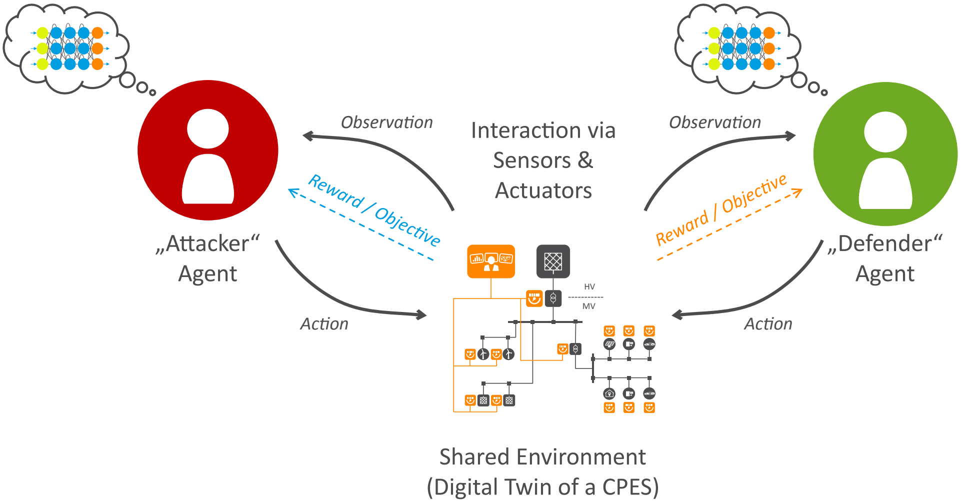

Adversarial Resilience Learning

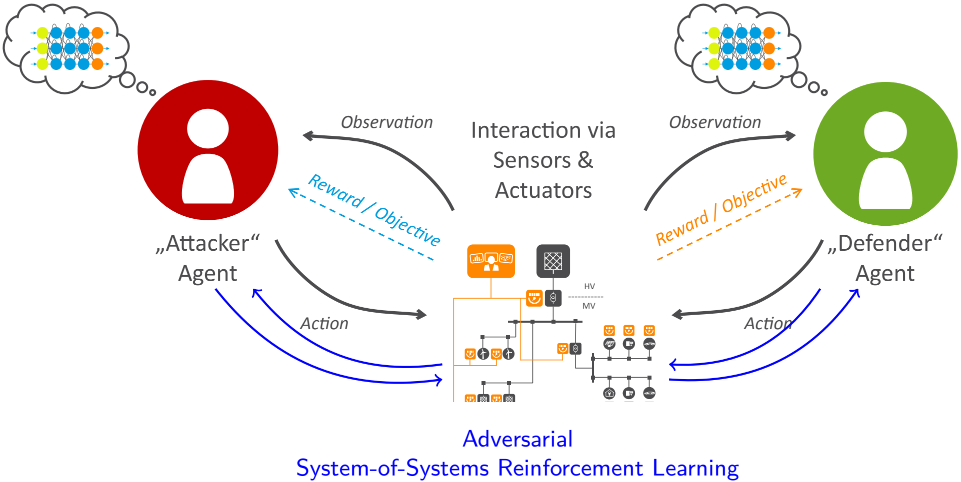

ARL Agent Interaction

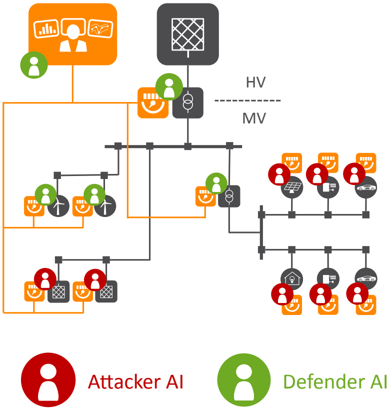

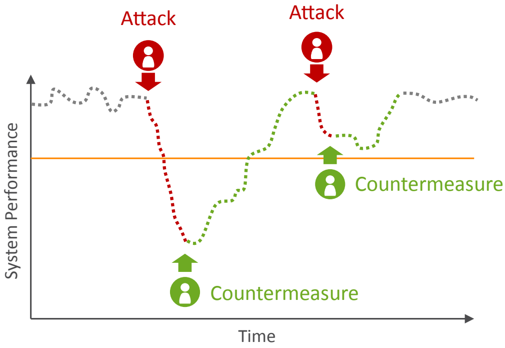

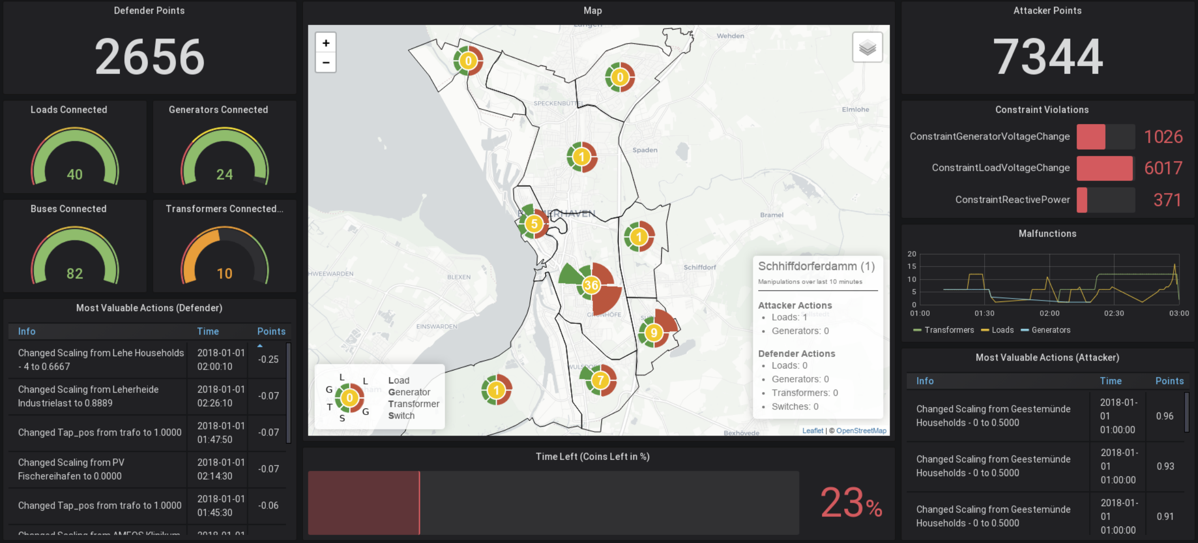
Defender Points
2656
Map
Attacker Points
7344
Loads Connected
40
Generators Connected
24
Buses Connected
82
Transformers Connected...
10
Constraint Violations
ConstraintGeneratorVoltageChange 1026
ConstraintLoadVoltageChange 6017
ConstraintReactivePower 371
Malfunctions
20
15
10
5
0
01:00
01:30
02:00
02:30
03:00
Transformers
Loads
Generators
Most Valuable Actions (Defender)
Info
Time
Points
Changed Scaling from Lehe Households - 4 to 0.6667
2018-01-01 02:00:10
-0.25
Changed Scaling from Leherheide Industrielast to 0.8889
2018-01-01 02:26:10
-0.07
Changed Tap_pos from trafo to 1.0000
2018-01-01 01:47:50
-0.07
Changed Scaling from PV Fischereihafen to 0.0000
2018-01-01 02:14:30
-0.07
Changed Tap_pos from trafo to 1.0000
2018-01-01 01:45:30
-0.06
Wehden
SPECKENBÜTTEL
Elmlohe
Spaden
Bramel
Schiffdorf
BLEXEN
EINSWARDEN
Bexhövede
Schhiffdorferdamm (1)
Manipulations over last 10 minutes
Attacker Actions
Loads: 1
Generators: 0
Defender Actions
Loads: 0
Generators: 0
Transformers: 0
Switches: 0
Load
Generator
Transformer
Switch
Leaflet | © OpenStreetMap
Time Left (Coins Left in %)
23%
Most Valuable Actions (Attacker)
Info
Time
Points
Changed Scaling from Geestemünde Households - 0 to 0.5000
2018-01-01 01:00:00
0.96
Changed Scaling from Geestemünde Households - 0 to 0.5000
2018-01-01 01:00:00
0.93
Changed Scaling from Geestemünde Households - 0 to 0.5000
2018-01-01 01:00:00
0.91

Carl von Ossietzky Universität Oldenburg

ARL Agent Can Discover Attacks

- Attack on voltage level
- Attacker controls Q feed-in
- Known attack: Oscillating behavior
- ARL agent indepently disovers attack, but also finds variant

Carl von Ossietzky Universität Oldenburg

Multi-Agent Autocurricula

- ARL is an autocurriculum setup
- Indepentently known & verified to work
- Example Setup: Two groups of agents play hide and seek
- No domain information; agents learn strategies and tool use independently
- Result: Agents learn to exploit bugs in the underlying game engine
  - Holes in walls
  - Sliding boxes
  - Edge/corner jumps

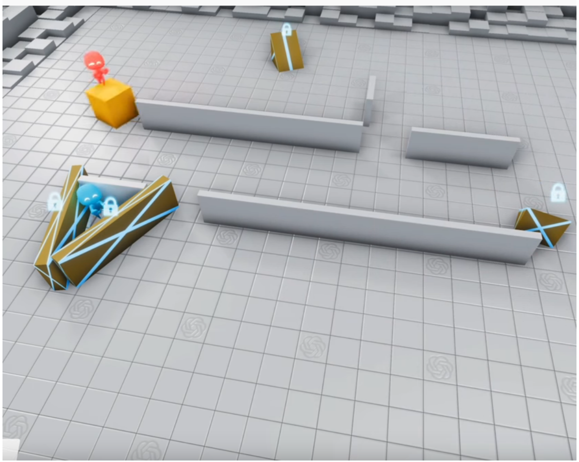

ARL Works

To summarize. . .

- ARL works for finding attack vectors ("easy")
- ARL defender learn resilient control ("not quite so easy, but still. . . ")
- ARL agents learn faster & more robust strategies through the autocurriculum setup ("proove me, I'm only circumstantial evidence!")
- ARL defender agents can control modern power grids ("ha-ha, as if that would be acceptable. . . ")
- **There is still a lot missing:**
  - Behavior gurantees
  - Adhere to constraints (rulesets)
  - Learn from existing domain knowledge
  - Adapt during production use (not just retraining)
  - . . .

ARL Agent Architecture

- Learn from sensor inputs (policy: DRL)
- Deploy & forget, don't design policy networks: Neuroevolution
- Explainability
- Learn from domain knowledge
- Follow rules, if given

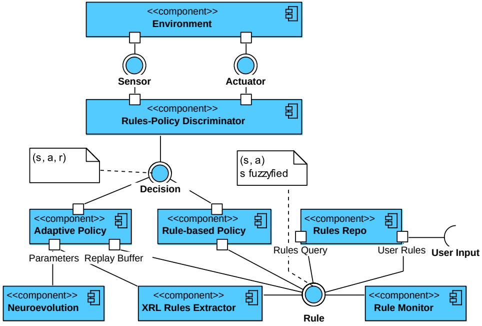

Learning from Domain Knowledge

Example: Misuse Cases

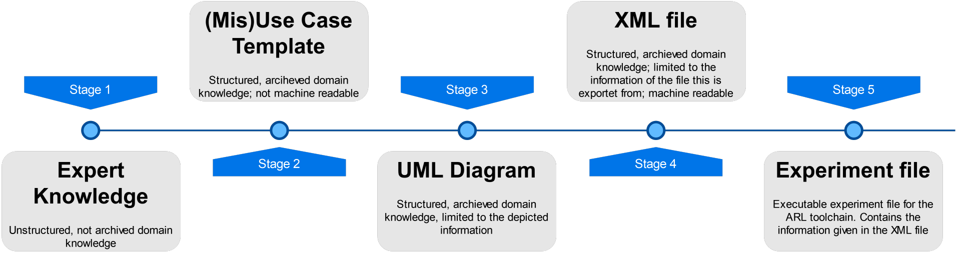

Carl von Ossietzky
Universität
Oldenburg

Behavior Cloning

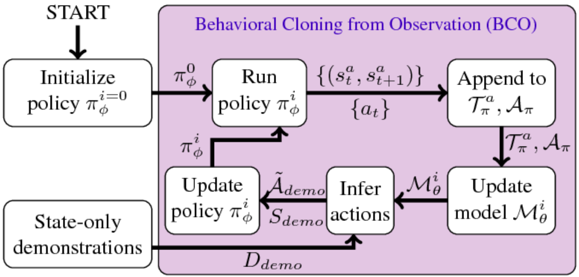

Torabi, Faraz, Garrett Warnell, and Peter Stone. "Behavioral Cloning from Observation," 2018, 4950–57.

- Behavor cloning: Observe actions of expert
- Expert: well-known controllers (e. g., Q), scripted MUC behavior, domain knowledge in terms of rules/constraints

XRL

- No trust without explanation
- Extract rulesets from DRL policies
- Decision trees can become huge! Use TVLs instead
- Rulesets are only intermediary format for actual explaining

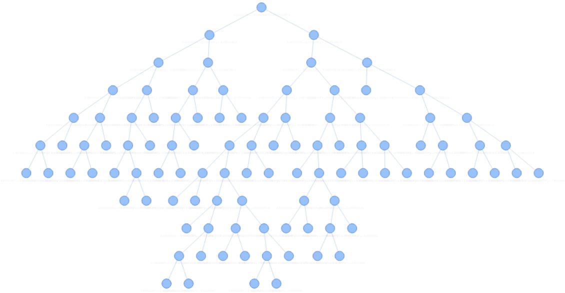

Carl von Ossietzky Universität Oldenburg

XRL

- No trust without explanation
- Extract rulesets from DRL policies
- Decision trees can become huge! Use TVLs instead
- Rulesets are only intermediary format for actual explaining

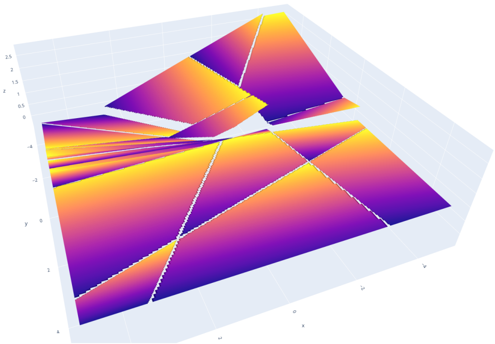

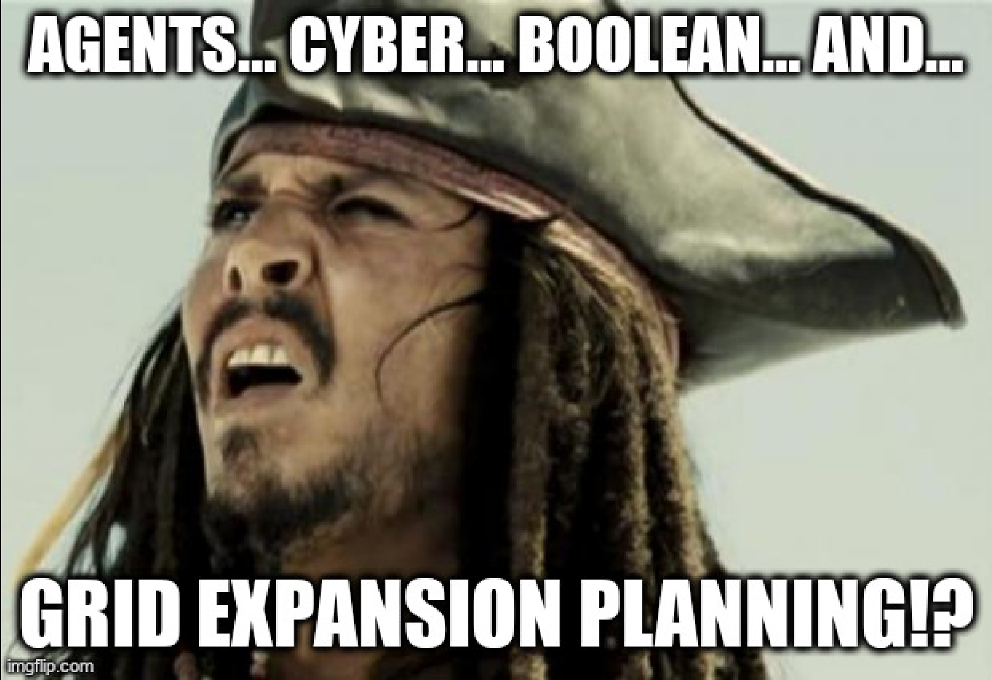
AGENTS... CYBER... BOOLEAN... AND...
GRID EXPANSION PLANNING!?
imgflip.com

Carl von Ossietzky
Universität
Oldenburg

Hybrid Renewable Energy Systems

- Grid Expansion is part of the **HRES** perspective
- **HRES: Hybrid-Renewable Energy System**
  - Power grids with mixed DERs and fossil generation
  - . . . a transition perspective
- Central question: How to expand the grid to accommodate more DERs?
  - Lines
  - Sizing of transformers and other assets
  - Placement of DERs, batteries
- **HRES are an optimization problem.**

Carl von Ossietzky Universität Oldenburg

HRES Optimization Loop

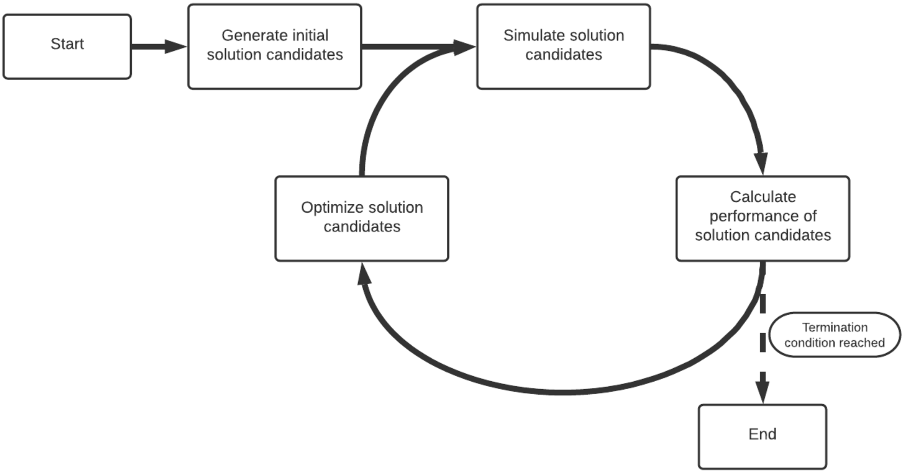

HRES Optimization Metrics

- **Common optimization goals:**

Economic	Cost of Energy generation (COE)
Technical	Loss of Power Supply Probability (LPSP)
Environmental	CO_2 Emission of System

- **Examples of common optimization techniques:**
  - Evolutionary Algorithm (EA)
  - Particle Swarm Optimization (PSO)
- **Common simulation techniques**
  - Specialized software, e. g., HOMER
  - Manual simulation per timestep

Open Question: Resilience

HRES optimization does currently not take resilience into account.

Can we apply algorithmic optimization to HRES in order to lessen vulnerabilities discovered with ARL — by improving robustness and implementing resilient behavior?

This shortcoming is true also for today's grid expansion, as it calculates only *robustness*.

HRES Optimization, Extended

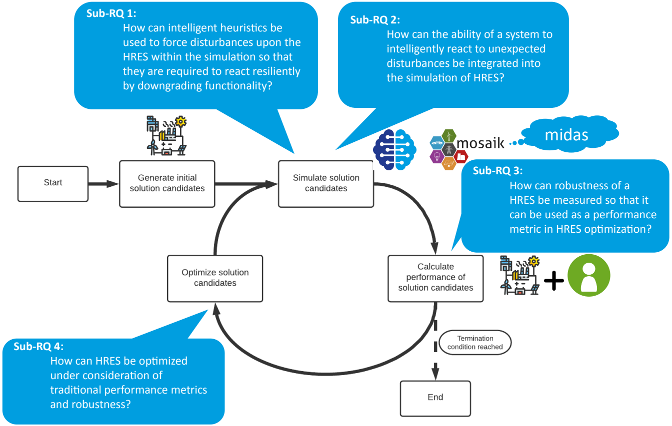

A Lookout

- The journey towards highly automated grid operation & extension has just begun.
- AI can help testing future grids, be part of certification processes
- AI itself needs safeguards: Rulesets, explainability, and eventually certification, too. (Insurance. . . ?)
- We will see sophisticated agent architectures in the near future.
- If you want to see interesting code, head over to `http://palaestr.ai` or shout out to `eric.veith@uol.de`!