



Destroy Your Power Grid as a Service (DYPGaaS)

How Novel Grid Analysis Improves Grid Expansion Planning

Eric MSP Veith <eric.veith@uol.de> , 2023-03-13



In A Nutshell

- ▶ Journey from grid expansion to Multi-Agent Systems to Agents destroying the power grid to grid expansion.
- ▶ Grid expansion suffers from run-time resilience analysis.
- ▶ Too many actors, grid becomes too intelligent to calculate robustness beforehand.
- ▶ AI for grid testing AND resilient operation — but needs safeguards, online learning, explainability
- ▶ There will be no “flag day,” so we need to develop the hybrid scenario all the way.



% 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)



Electricity Demand Rising

After significant decline in 2022, European electricity demand is set to recover

Year-on-year relative change in electricity demand, Europe, 2019-2025




IEA. CC BY 4.0.




Renewables Are Replacing Fossil Fuels

Following two years of increases, CO₂ intensity starts to decline again from 2023 onward

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




Development of average CO₂ intensity, Europe, 2019-2025





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





Grid Expansion Planning Today

Step 1: Data Acquisition



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

dena-Verteilnetzstudie


Destroy Your Power Grid as a Service (DYPGaaS) — How Novel Grid Analysis Improves Grid Expansion Planning

Eric MSP Veith <eric.veith@uol.de> — Adversarial Resilience Learning



Grid Expansion Planning Today

Step 2: Modelling




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



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



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

dena-Verteilnetzstudie




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. efficient operation (?)**




AI as Promise of an Alternative

- ▶ Multi-Agent Systems promise local, 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






Forecasting






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




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






Algorithm Behavior






Algorithm Behavior






Algorithm Behavior






Algorithm Behavior






Four-Way Handshake





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
- ▶ $\text{distance}(l_{i,k})$: 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
- ▶ $\text{distance}(l_{i,k})$: Distance metric
- ▶ m_j : j -th message
- ▶ M_i : Message Journal of the i -th agent

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



Forwarding

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

$$M_i = \{m_1 \mapsto \{(l_{i,1}, m_{1,\text{distance}(l_{i,1})}), \dots, (l_{i,n}, m'_{1,\text{distance}(l_{i,n})})\},$$

$\dots,$

$$m_n \mapsto \{(l_{i,1}, m_{n,\text{distance}(l_{i,1})}), \dots, (l_{i,n}, m'_{n,\text{distance}(l_{i,n})})\}\}$$

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



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 forwarding
 - 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





Problem Statement

‘Power Balance Algebra’:

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

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



Problem Statement

‘Power Balance Algebra’:

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

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

Distance Function:

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



Problem Statement

'Power Balance Algebra':

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

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

Distance Function: $d(r_i) : r_i \mapsto \mathbb{R} \quad (3)$

Problem Statement:

$$\sum_i b_i r_i \subseteq r_0 , \quad i \neq 0, b_i \in \{0, 1\} , \quad (4)$$

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



Atomization

$$\mathbf{P} = (|P_0|, |P_1|, \dots, |P_i|, |P_C|) ,$$

$$\mathbf{t} = (t_{2,0} - t_{1,0}, t_{2,1} - t_{1,1}, \dots, t_{2,i} - t_{1,i}) ,$$

$$\Delta \mathbf{P} = \text{ggT}(\mathbf{P}) ,$$

$$\Delta \mathbf{t} = \text{ggT}(\mathbf{t}) ,$$



Atomization

$$\mathbf{P} = (|P_0|, |P_1|, \dots, |P_i|, |P_C|) ,$$

$$\mathbf{t} = (t_{2,0} - t_{1,0}, t_{2,1} - t_{1,1}, \dots, t_{2,i} - t_{1,i}) ,$$


$$\Delta P = \text{ggT}(\mathbf{P}) ,$$

$$\Delta t = \text{ggT}(\mathbf{t}) ,$$

$$x_{i,\tilde{t},\tilde{P}} = \begin{cases} 1 & \text{if agent } i \text{ influences the grid in time-subinterval } \tilde{t} \text{ with} \\ & \text{power from the power-subinterval } \tilde{P}, \\ 0 & \text{else.} \end{cases}$$



Atomization Illustrated





Model of the Disequilibrium

A symmetric function for each time-subinterval:

$$S_k^n(\mathbf{x}_{i,\tilde{t}=k},\tilde{\mathbf{p}}) = \begin{cases} 1 & \text{if } n \text{ variables in } \mathbf{x}_{i,\tilde{t}=k},\tilde{\mathbf{p}} \text{ equal } 1, \\ 0 & \text{else;} \end{cases}$$

Full Disequilibrium:

$$S = \bigcap_{k=1}^m S_k^n(\mathbf{x}_{i,\tilde{t}=k},\tilde{\mathbf{p}})$$




Modelling Responses

Acceptance Function:

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

$$\begin{aligned} r_2(\mathbf{x}_{i,\tilde{t},\tilde{p}}) &= \bar{x}_{2,3,1} \wedge \bar{x}_{2,3,2} \wedge \bar{x}_{2,4,1} \wedge \bar{x}_{2,4,2} \\ &\vee x_{2,3,1} \wedge x_{2,3,2} \wedge \bar{x}_{2,4,1} \wedge \bar{x}_{2,4,2} \\ &\vee x_{2,3,1} \wedge x_{2,3,2} \wedge x_{2,4,1} \wedge x_{2,4,2} \end{aligned}$$





Equilibrium

$$S = \bigcap_{k=1}^m S_k^n(\mathbf{x}_{i,\tilde{t}=k}, \tilde{\mathbf{p}})$$

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

$$C = S \cap R .$$



Equilibrium

$$S = \bigcap_{k=1}^m S_k^n(\mathbf{x}_{i, \tilde{t}=k}, \tilde{\mathbf{p}})$$

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

$$C = S \cap R .$$

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



Efficiency

Data Effect

$$\kappa = \frac{W}{D} \left[\frac{\text{kWh}}{\text{kB}} \right]$$

Data Efficiency

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




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



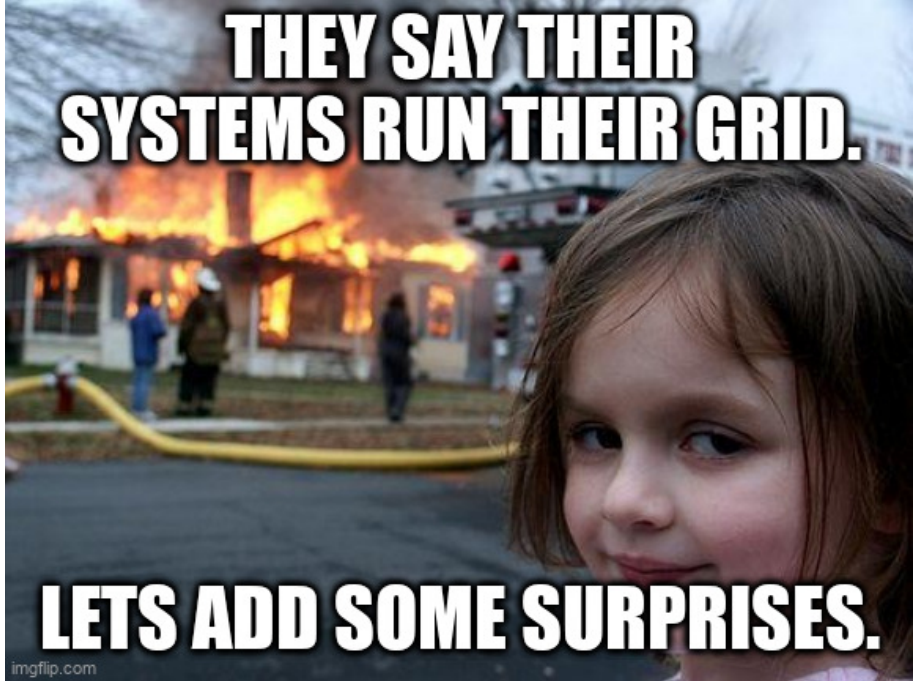
- ▶ BDD approach in low-load situation: 100 kB
- ▶ *Universal Agent* concept especially useful in complex load situations

AND THIS, GENTLEMEN

IS HOW YOU RUN YOUR GRID.



**THEY SAY THEIR
SYSTEMS RUN THEIR GRID.**



LETS ADD SOME SURPRISES.




**“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







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





Systemstabilität - Minutenreserve

- Abgerufene Menge (+)
- Abgerufene Menge (-)
- Arbeitspreis (+)
- Arbeitspreis (-)
- Vorgehaltene Menge (+)
- Vorgehaltene Menge (-)
- Leistungspreis (+)
- Leistungspreis (-)




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 some 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
- ▶ **Businesspeople**
 - ▶ Goal: (short-term) profit maximization
 - ▶ Damage to CNIs unexpected (or simply “don’t care”)
 - ▶ Uses loopholes and grey areas in codes



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
- ▶ Critical Services as SPOF (DNS, BGP, SCADA, SDL)
- ▶ **Learning Strategies for automatic issue management**
- ▶ “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

Loads Connected



Generators Connected



Buses Connected




Transformers Connected...



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
Changed Scaling from MFCG Windm...	2018-01-01	

Map



Time Left (Coins Left in %)



23%

Attacker Points

7344


Constraint Violations

ConstraintGeneratorVoltageChange 1026

ConstraintLoadVoltageChange 6017

ConstraintReactivePower 371

Malfunctions



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



ARL Agent Can Discover Attacks


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





Multi-Agent Autocurricula

- ▶ ARL is an autocurriculum setup
- ▶ Independently 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





Behavior Cloning




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

- Behavior 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





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!?




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.**



HRES Optimization Loop





HRES Optimization Metrics

- ▶ **Common optimization goals:**

- Economic** Cost of Energy generation (COE)

- Technical** Loss of Power Supply Probability (LPSP)

- Environmental** CO₂ 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

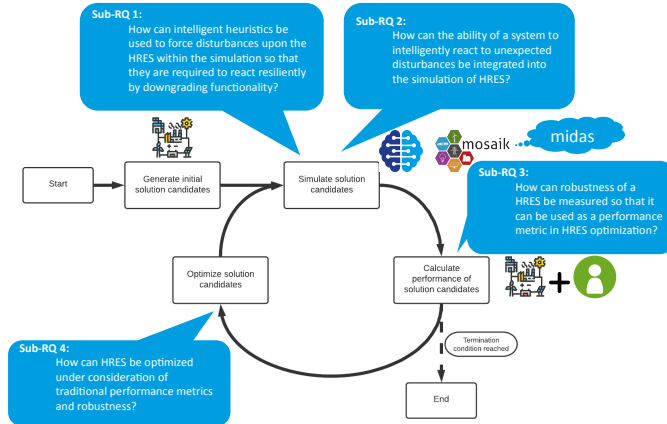
HRES optimization does currently not take resilience into account.

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!