AutoDC - Autonomous “Smart” Datacenters for long term deployment

WP 3 Machine Learning and Modelling

The aim of AutoDC is to provide an innovative design framework for autonomous data centers.

An autonomous datacenter should be able to, without any human intervention, from a best effort perspective continue its operation independent of contextual interference, such as intermittent power failure, failing components, overheating etc.

ITEA3

Agenda

- Overview of WP
- Exploitation Related Achievements (overview)
- High-level status of ongoing activities
- Demo
- Exploitation Related Achievements (details)

Overview of Work Package

- Develop data-driven concepts, methods, algorithms and tools for DC management and operations
- Innovate on data analysis and machine learning technologies using SW and HW statistics to support automated control, management, and sustainability of DCs and edge clouds
- Use cases include fault management, performance optimization, resource optimization, energy efficiency, and business strategies
- Expected results from this work package include:
  - Novel and scalable data-driven/machine learning technology for data center automation with respect to both infrastructure and services
  - A Prescriptive Decision Support System that optimize data center operations with business objectives demonstrators showing key concepts
  - Intellectual property rights in the form of patent applications
  - Research papers at top venues describing and evaluating concepts and approaches

AutoDC WP3 tasks

PROJECT IDEA NUMBER: 37

T3.1 – Novel DCM and PDSS solutions

- ML models for strategic business support

T3.2 – Predictive and prescriptive modelling

- Development of ML models supporting automated DC operations
- Energy efficiency, anomaly detection, optimization, bottleneck mitigation, ...

T3.3 – Nonstationary learning and scalability

- Techniques for sharing and adapting learned knowledge encoded in ML models
- Distributed learning and feature reduction for scalability and increased learning efficiency

T3.4 – Integration of ML and model-based techniques

- Control theoretical techniques or physical models to overcome challenges related to e.g. lack of data

Partners

PROJECT IDEA NUMBER: 37

- Aalto
- Clavister
- Ericsson
- KTH
- Luleå University
- Lund University
- Mariner
- MLT
- RISE North
- RISE

Ongoing activities

PROJECT IDEA NUMBER: 37

#	Activities	T3.1	T3.2	T3.3	T3.4	Partner(s)
A1	Development of DCM and PDSS	x				MLT
A2	Security Breach Detection		x	x		Clavister, RISE North
A3	Prediction of power consumption		x	x		RISE North, Mariner, Ericsson, KTH
A4	Risk aware control for micro-grid connected edge DC		x		x	RISE North, LTU
A5	Transfer learning for edge DC thermal models		x	x	x	RISE North, LTU
A6	Self-driving systems		x	x		KTH, RISE, Ericsson
A7	Source selection in transfer learning for dynamic edge DC environments			x		Ericsson, KTH
A8	Feature selection in DC environments for scalability			x		Ericsson, KTH, Mariner
A9	Root cause analysis based on log files		x			Ericsson
A10	Data exploration and description		x			Aalto, RISE North
A11	Observe and predict future operational states		x			Mariner

Demo

Exploitation Related Achievements

- 3 Utility (patent)
- 2 Human capital
- 7 Academic exams
- 9 Publications (dissemination)
- 1 Collaboration (exploitation)
- 3 Internal (dissemination)
- 3 Conference (dissemination)

A3 Prediction of Power consumption

PROJECT IDEA NUMBER: 37

- Use case: Prediction and forecasting of power consumption
  - Resource planning of cooling, electricity, and workload orchestration
  - Identify anomalies at component level and isolate the cause
  - Determining mitigation actions, e.g. load balancing, re-orchestration, re-routing, chiller control, ...

- Experimental platform/data sets
  - RISE North edge DC

- Developed ML model
  - LSTM model to forecast future CPU utilization and power consumption

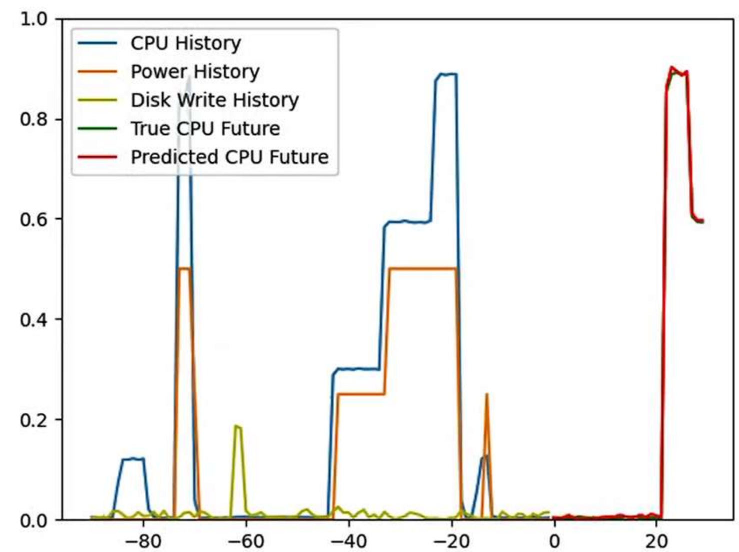

- Ongoing work
  - NN models for prediction of power consumption and detection of anomalies
  - Transfer of NNs to reduce training time when environment changes
  - Reinforcement learning to decide the number of powered-on servers in a cluster

A4 Transfer learning for edge DC thermal models

PROJECT IDEA NUMBER: 37

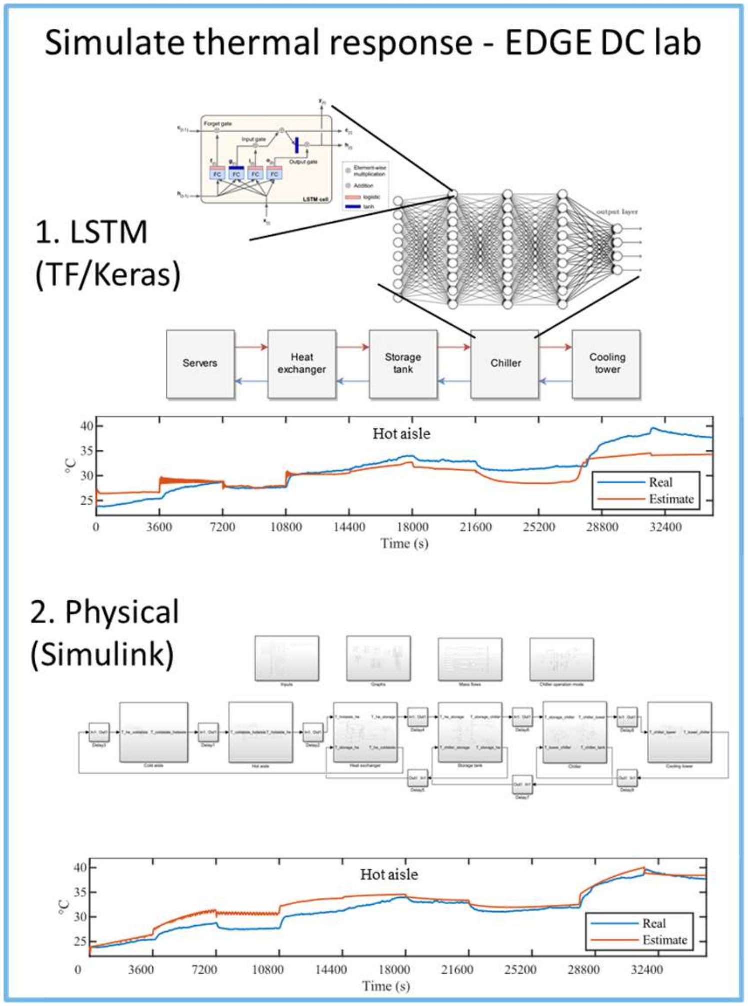

Comparison of *first principle* and Purely *data driven* models

- Development of grey-box ML based controller for small cluster
- Train grey-box model on small server cluster set-up and deploy in new context
  - Model predicts temperature of CPU/hot isle
  - Build controller for thermal management
  - Use modern open source machine learning libraries TensorFlow/Keras
  - Compare models e.g. with vanilla RNN
  - Transfer to other set-up

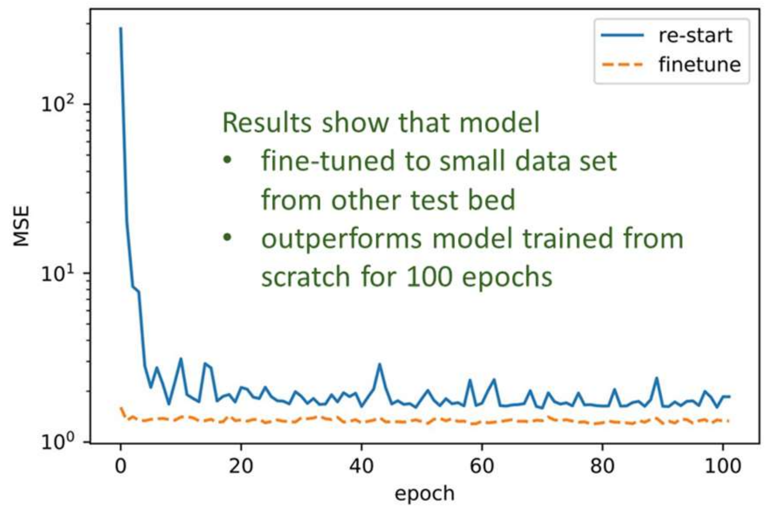

A5 Risk aware control for micro-grid connected edge DC

PROJECT IDEA NUMBER: 37

- Development of micro-grid control model for edge DCs

Build model	Simulate	Condition on	Plan
for edge DC test **bed from** previous work (WP2) based on discrete Markov **process**	microgrid assets - battery **charge** - **solar** panel yield with conf intervals	public **weather forecast (probabilities for scenarios)** **electricity** grid **price tariff** load forecast	battery use and load commitment under risk-aware constraints, e.g. island mode and to meet SLAs

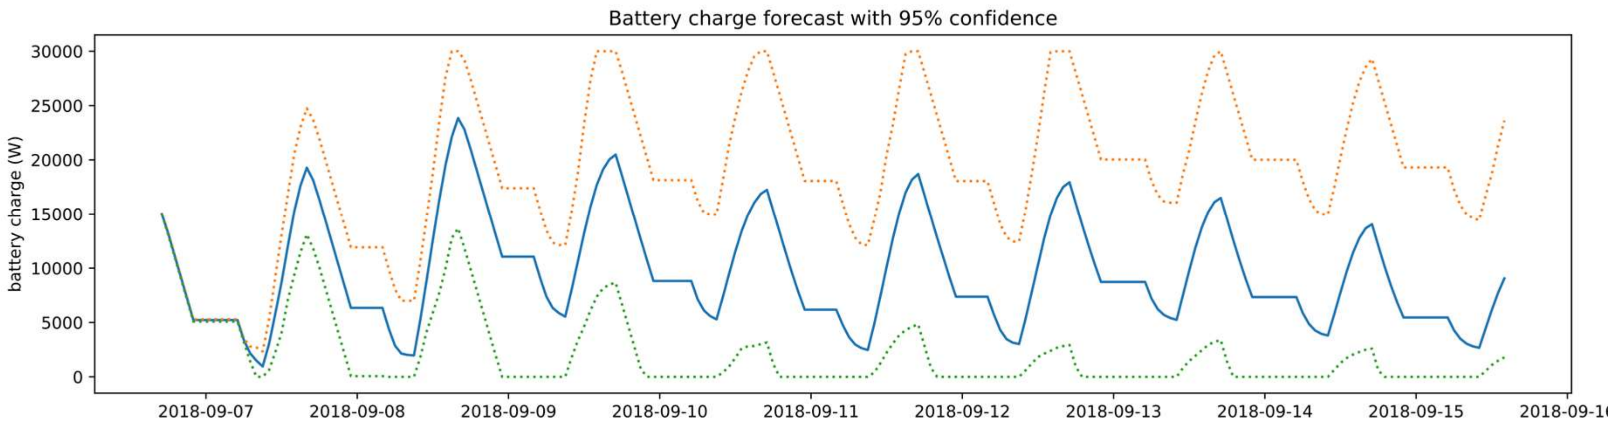

A6 Self-driving systems

PROJECT IDEA NUMBER: 37

- Use case:
  - KPI-based resource control of K8 services
  - System dynamically configures to meet management objectives
  - System dynamically adjusts configuration parameters to changes in the environment

- Approach
  - Learn regression models for service-level KPIs prediction and forecasting from monitored data
  - Apply reinforcement learning to control DC resource allocation, e.g., horizontal and vertical scaling
  - Scalability though dimensionality reduction

- Status
  - Testbed setup (K8, MongoDB)
  - Literature review on autonomic systems, self-adapting systems, ...

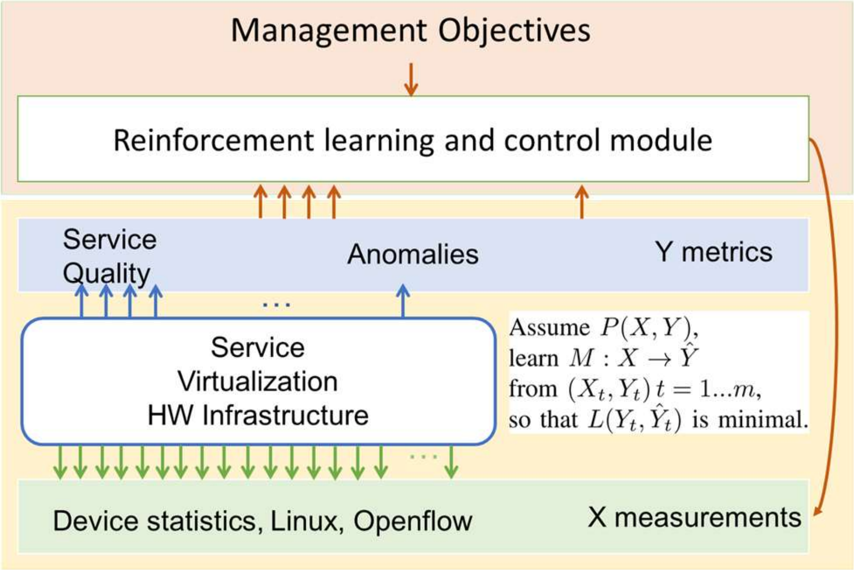

- Example
  - Objective: minimize energy usage while 95 % of request served below 20ms
  - Environmental change example: system adjusts to changes in load pattern, query profile, available resources

A2 Finding Security Breaches

- Use case: Identify communication anomalies wrt NW control traffic in DC to find potential breaches
  - Example: Can we suspect a host to be breached if it change its communication behavior and start using new protocols/applications, change data volume/rate for applications or communicate with new network destinations?

- Approach
  - Use RISE DC as an experimental platform and gather network statistics for the control traffic
  - Use ML to learn behaviour of NW control traffic in DC in order to identify security breaches and other potential anomalies and determine mitigation actions (e.g. NW segmentation)

- Ongoing work
  - Data logging has just started, and we are working on gathering statistics from the RISE DC to have data we can use for learning

CONNECT . PROTECT

A11 Observe and Predict Future Operational States – Energy & Temperature

PROJECT IDEA NUMBER: 37

- Use case:
  - Dynamically propose new sets of operational parameters to minimize facility’s energy consumption while maintaining normal performance.
  - Assess proposed parameters and predict future facility states.
  - (WP4) Adjust the operational parameters (providing predictions indicated acceptable facility state).
- Example
  - Objective: minimize energy usage by 15% while maintaining optimal ambient temperature.
  - Adjusts to changes external (weather) and internal (load, heat dissemination, etc).
- Approach
  - Train/learn RNN/LSTM models for each zone/space within facility from thousands input parameters (HVAC, sensors, internal temperature, external weather, etc).
  - Apply reinforcement learning to propose changes to operational parameters (set points).
  - Use Live RNN Models to predict the future states.
  - Execute changes and observe actual states.
- Status
  - Primary use case testing complete.
  - Software Implementation in progress.
  - Models finetuning in progress.

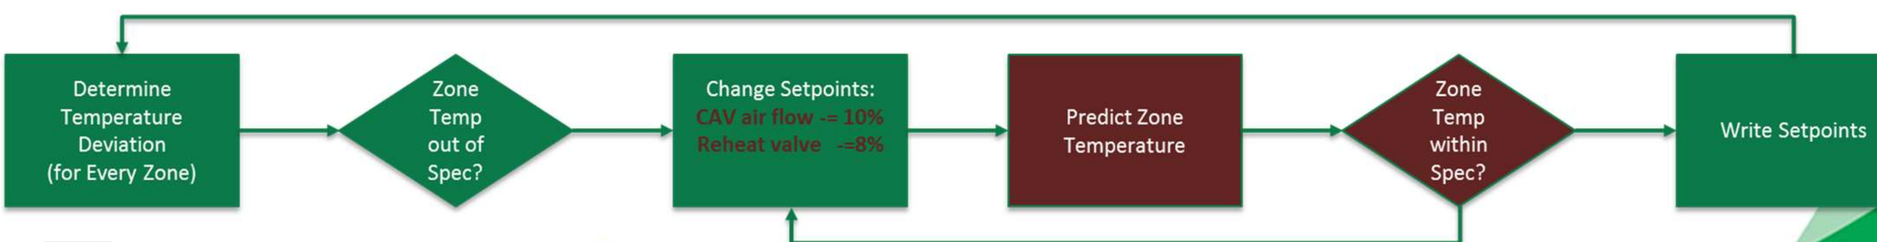

A11 Observe and Predict Future Operational States – Anomaly Detection

PROJECT IDEA NUMBER: 37

- Use case:
  - Use ML algorithms to detect, anticipate and predict failures within DC network / systems / services (elements).
- Example:
  - Monitor the SysLog messages collected from multiple/varied vms/servers within data center.
  - Detect and alert on the changes in SysLog behavior for the given elements in real-time (e.g. increased log levels, changed log types, etc).
  - (WP4) Issue actions to correct/prevent failures in the detected elements.
- Approach:
  - Learn and generate models for every element based on its logs.
  - Using the generated models, predict expected behavioral parameters or each element.
  - Continuously observe reported logs from each element.
  - Use standard deviation classification to identify anomalous elements.
- Status:
  - Software development complete.
  - PoC implementation complete.

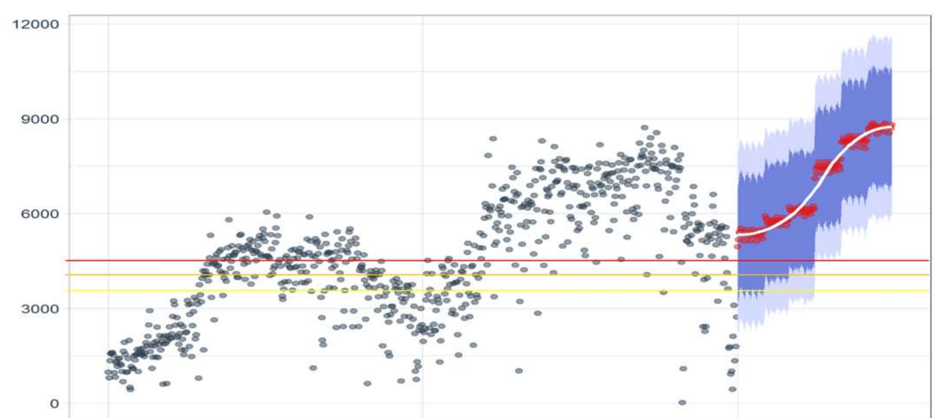

A few words about the other activities

PROJECT IDEA NUMBER: 37

- A1 Development of DCM and PDSS (MLT)
  - Focus on data engine portion of the PDSS
  - Implementation of query/algorithm insertion methods and a DaaS interface
  - Started working on integrating ML to the PDSS

- A8 Feature selection in DC environments for scalability (Ericsson, KTH, Mariner)
  - Exploring multiple strategies for feature reduction and selection
  - E.g. “must-have policies” (**patented**), supervised/unsupervised, and online feature selection

- A9 Root cause analysis based on log files (Ericsson)
  - Automation of root-cause analysis using log files from an operational DC at Ericsson
  - Investment in new HW for logfile collection and analysis pipeline
  - A new logfile parser and ML-based classifier has been prototyped

- A10 Data exploration and description (Aalto and RISE North)
  - Received data from RISE North DC
  - Recruitment of post-doctoral researcher at Aalto, 2020
  - Initial work on exploratory data analysis an dfeature reduction techniques

DEMO

A7 Source selection in transfer learning for dynamic edge DC environments

PROJECT IDEA NUMBER: 37

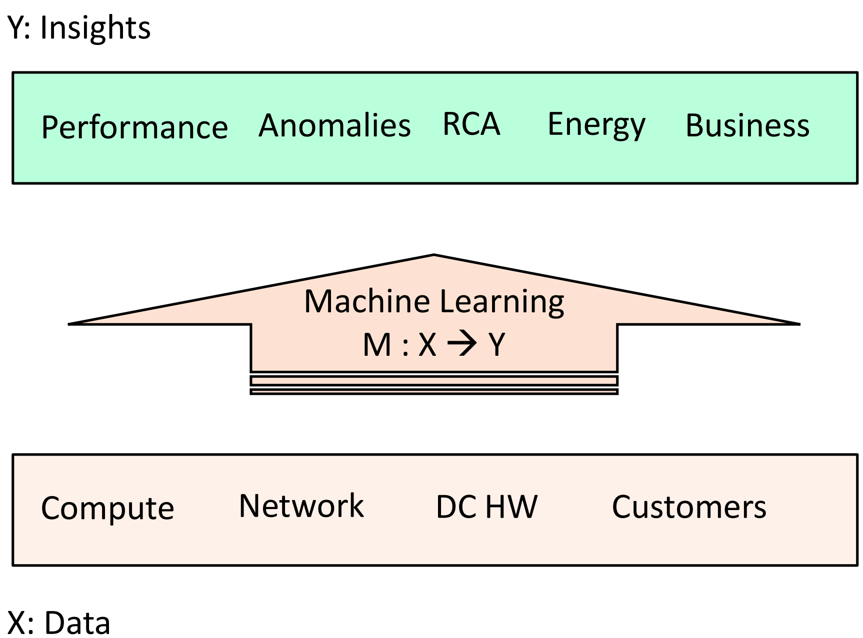

- Multiple examples of training ML models aiming at automating aspects of DC ops
  - Power consumption forecasting
  - Security breaches
  - Thermal response
  - RL for self-driving systems
  - KPI prediction
  - ...
- Long-term deployments render changes in the DC environment (e.g. $X \rightarrow \hat{X}$)

How to obtain new model $\widehat{M}: \hat{X} \rightarrow Y$?

T3.3: Non-stationary learning using transfer learning

PROJECT IDEA NUMBER: 37

- Challenge: Changes in execution environment renders existing ML models obsolete
  - Gradual deterioration of DC components, e.g. measurement tools or compute resources
  - Changes in load pattern and co-location of services
  - Changes in HW configurations, e.g. power states
  - Few samples available for training new ML model after a change

- Solution
  - Reuse ML model knowledge using transfer learning
  - Source selection to find the most fitted model based on
    - KL-divergence
    - Entropy (**patented**)

- Demonstrator key messages
  - Need for source selection in transfer learning as it can either boost or reduce the performance
  - Transfer learning critical when target-domain samples are limited
  - Transfer gain (and penalty) is most visible where there are limited data in the target domain

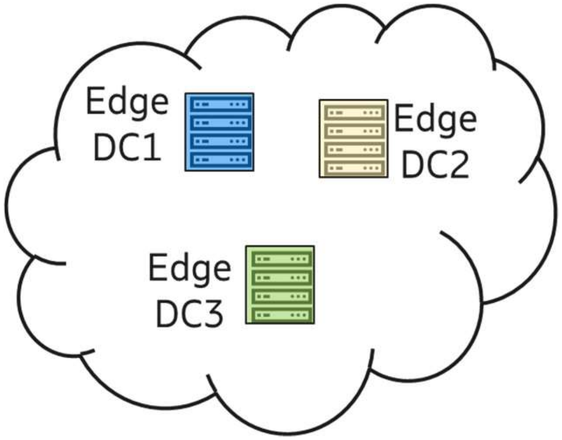

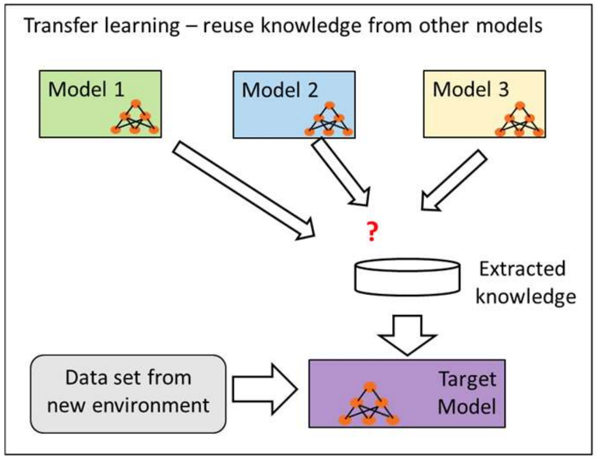

Exploitation Related Achievements

- Utility (patent)
  - Measurement Reporting and Configuration in Communication Networks, PCT/EP2020/065127, 2020
  - Methods and systems for dynamic service performance prediction using transfer learning, PCT/SE2019/050672
  - Improving performance modeling in dynamic clouds, US 62/770,330

- Human capital
  - Recruitment of post-doctoral researcher at Aalto, 2020
  - Employment of Mikko Siltala, 2020

Exploitation Related Achievements

- Academic exams
  - Mikko Siltala, Simulating data center cooling systems: data-driven and physical modeling methods, Msc thesis, Aalto University, 2020
  - X. Wang, “Dimensionality reduction for performance prediction in networked systems,” master thesis, KTH Royal Institute of Technology, Stockholm, 2020.
  - C. Teng, “Forecasting service metrics for network services,” master thesis, KTH Royal Institute of Technology, Stockholm, 2020.
  - **Rickard Brännvall, Machine learning based control of small-scale autonomous data centers, Licentiate dissertation, Luleå University of Technology, 2020**
  - Vamshi Pulluri, Hardware Power Optimization of Base Station Using Reinforcement Learning, MSc thesis, KTH, 2019
  - Hongyi Zhang, Efficient learning on high-dimensional operational data, MSc thesis, KTH, 2019
  - Service Metric Prediction in Clouds using Transfer Learning, MSc thesis, 2019

Exploitation Related Achievements

- Publications (dissemination)
  - X. Wang, F. Shahab Samani, R. Stadler: “Online feature selection for rapid, low-overhead learning in networked systems,” 6th IFIP/IEEE International Conference on Network and Service Management, 2-6 November 2020.
  - Chemouil, P., Hui, P., Kellerer, W., Limam, N., Stadler, R. and Wen, Y., Guest Editorial Special Issue on Advances in Artificial Intelligence and Machine Learning for Networking. IEEE Journal on Selected Areas in Communications, 38(10), pp.2229-2233. 2020.
  - Chemouil, P., Hui, P., Kellerer, W., Li, Y., Stadler, R., Tao, D., Wen, Y. and Zhang, Y., Special issue on artificial intelligence and machine learning for networking and communications. IEEE Journal on Selected Areas in Communications, 37(6), pp.1185-1191.2019.
  - Brännvall, R., Siltala, M., Gustafsson, J., Sarkinen, J., Vesterlund, M. & Summers, J. (2020). EDGE: Microgrid Data Center with Mixed Energy Storage. In: e-Energy '20: Proceedings of the Eleventh ACM International Conference on Future Energy Systems. Paper presented at 11th ACM International Conference on Future Energy Systems (ACM e-Energy 2020), 22-26 June, 2020, Virtual Event, Australia (pp. 466-473). Association for Computing Machinery (ACM)
  - C. Flinta, W. Yan, and A. Johnsson, "Predicting Round-Trip Time Distributions in IoT Systems using Histogram Estimators", NOMS 2020-2020 IEEE/IFIP Network Operations and Management Symposium. IEEE, 2020.
  - Efficient Learning on High-dimensional Operational Data, International Conference on Network and Service Management (CNSM), 2019
  - Digital Twin for Tuning of Server Fan Controllers, IEEE International Conference on Industrial Informatics (INDNI), 2019
  - Performance Prediction in Dynamic Clouds using Transfer Learning, IFIP/IEEE Integrated Network Management (IM), 2019 **(best paper award)**
  - Predicting Distributions of Service Metrics, IFIP/IEEE Integrated Network Management (IM), 2019

Exploitation Related Achievements, cont.

- Collaboration (exploitation)
  - Established collaboration RISE North, Luleå University of Technology, and Ericsson
- Internal (dissemination)
  - Demonstration of transfer learning, 2020
  - Seminar on CDE using MDNs, 2019
  - Seminar on transfer learning, 2019
- Conference (dissemination)
  - IEEE JSAC Issue on AI and Networking, 2020
  - KTH organized and moderated the Distinguished Experts Panel (DEP) at NOMS 2020 on "Management in the Age of Softwarization, Artificial Intelligence, and Cybersecurity", 2020
  - Introduction to data-driven engineering of networked systems, Keynote at BlackSeaCom, 2019

Backup, old, deselected, and templates

WP3 is developing technology enablers contributing towards AutoDC goals

1. Decrease in the total cost of ownership due to the autonomous operation of the data center

2. Decrease in the amount and the frequency of maintenance operations required for the operations of the autonomous data center
   - Failure will happen, but the project will optimize the mean-time-to-repair.

3. Level of recovery of the data center from unexpected failure
   - Short outages per data center will happen but the mean time to recover need to be shortened and the overall availability target must be very high.

4. Energy-efficiency of the autonomous data center

5. Environmental impact of the operations and construction of the autonomous data center
   - Target is to minimize site travel and on-site construction.

6. Level of autonomy in the subsystems (power supplies, cooling, computing infrastructure) to enable improved mean-time-to-recover

Coping with large feature spaces and massive data

PROJECT IDEA NUMBER: 37

Scalability challenge

- Impractical to collect and maintain all data features in a DC
- Large data storage
- Cost associated to data transfer, network capacity and energy

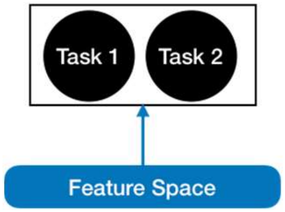

Feature Selection for improved ML:

- Reduces data dimensionality
- Maintains interpretability (e.g., root-cause analysis)

Supervised **Feature Selection**:

- Use-case (task) specific

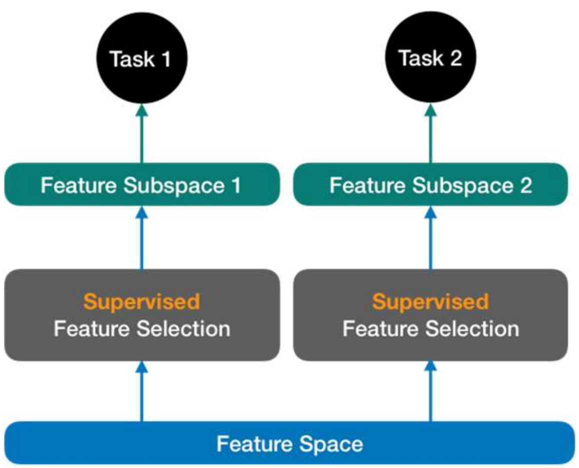

Unsupervised **Feature Selection**:

- Use-case (task) specific

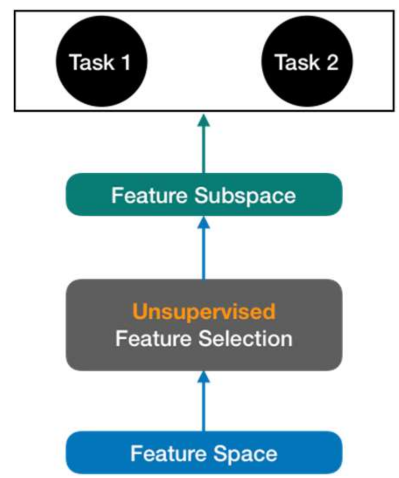

Coping with large feature spaces and massive data

PROJECT IDEA NUMBER: 37

- Models for predicting service-level metrics Y from infrastructure metrics X

- Approach for dimensionality reduction
  - Feature selection
  - Principle component analysis (PCA)
  - Autoencoders

- Results

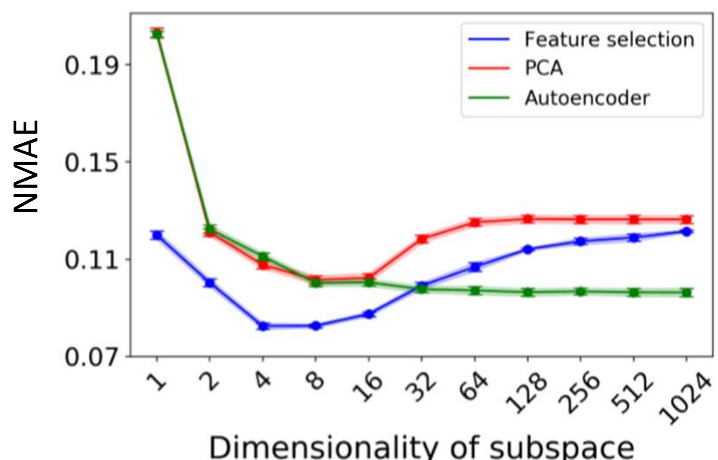

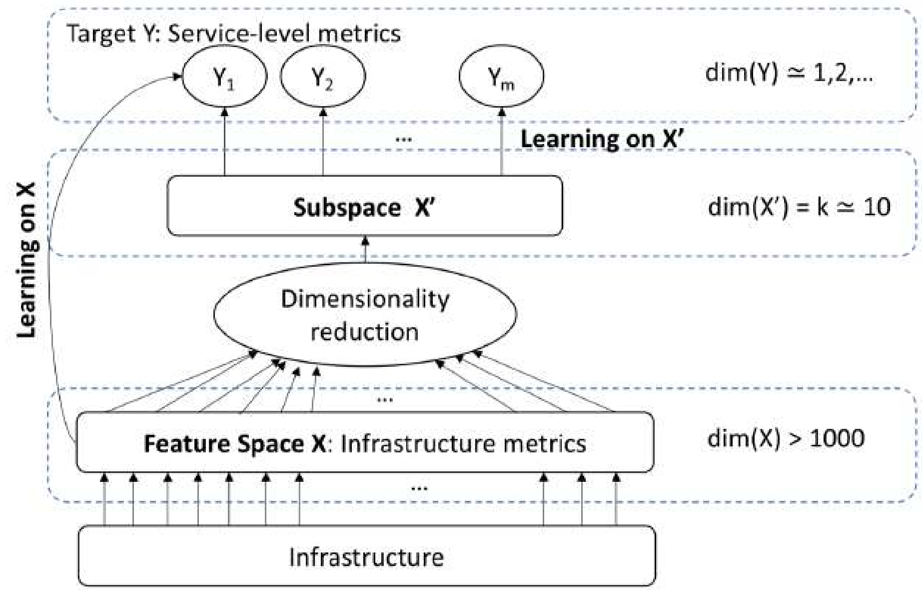

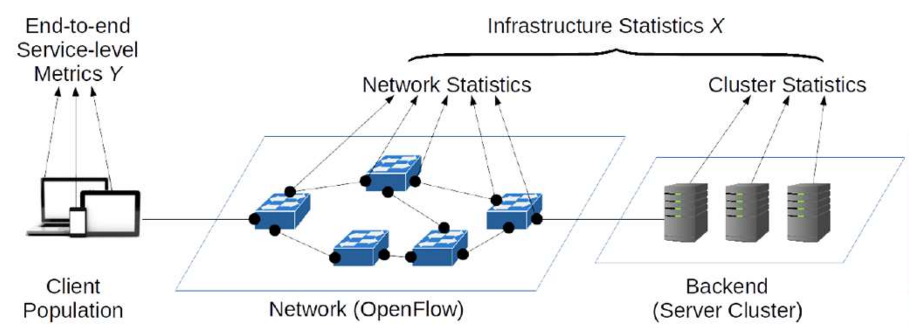

A Feature selection in DC environments for scalability

PROJECT IDEA NUMBER: 37

- Challenges
  - The number of metrics possible to measure in a DC is very high
  - Monitoring and data maintenance suffers from scalability challenges
  - ML modeling suffers from curse of dimensionality

- Exploring multiple strategies for feature reduction and selection
  - Must-have policies (**patented**)
  - Supervised and unsupervised feature selection
  - Online feature selection

- Important for reducing energy consumption
  - Data sampling, ML training, ML inference

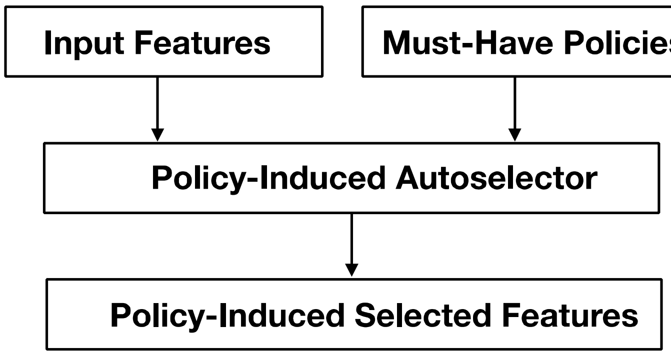

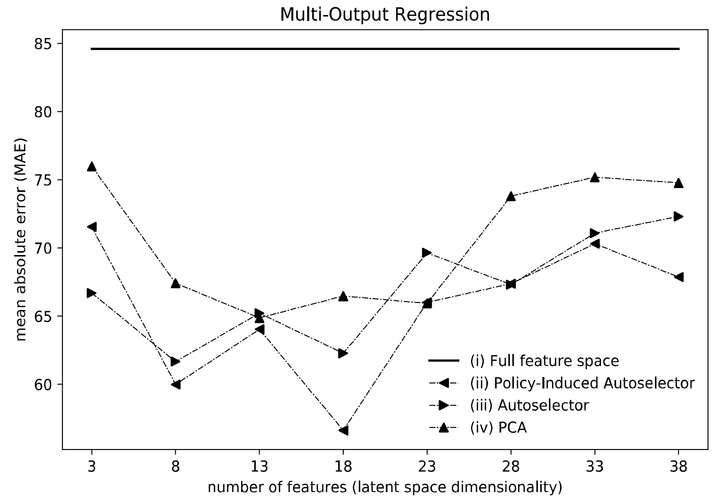

Simulation models for EDGE DC cooling management

PROJECT IDEA NUMBER: 37

- Towards energy efficient autonomous DCs
- Predicting DC temperatures using ML
  - Data-driven model using LSTM neural network
  - Implemented using Keras/TensorFlow
  - The model was trained with 90 hours of data from the edge DC at RISE I.C.E.
  - The model was tested by using 1 hour of data to initialize the LSTM network, and then predicting 9 hours into the future, while changing the setpoints hourly.
- Model is an enabler for improved control strategies and fault detection

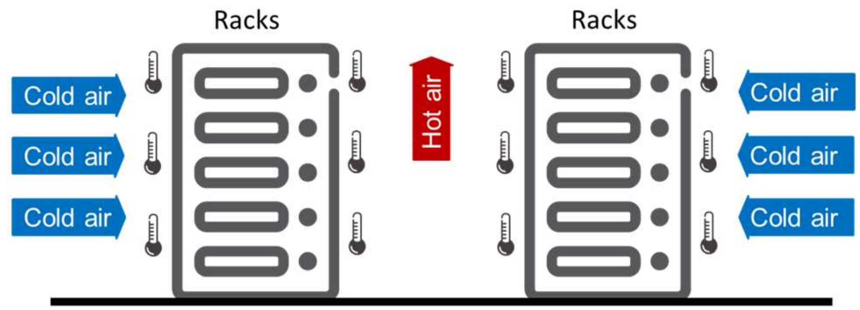

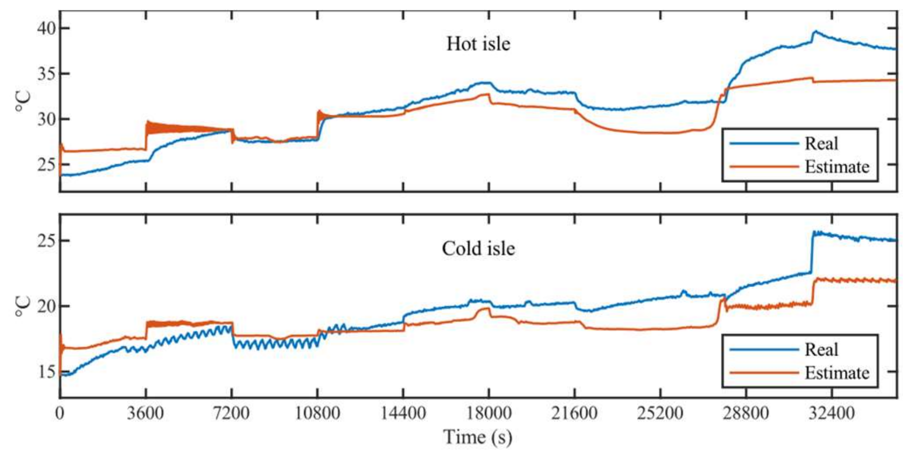

Figure: Hot and cold isle temperature predictions from the simulation model test.

MAE = 1.84 on hot isle
MAE = 1.44 on the cold isle

ICE
RI.SE SICS NORTH

Anomaly detection and root-cause analysis

PROJECT IDEA NUMBER: 37

- Anomaly detection and RCA is an enabler for autonomous, self-driving DC operations

- The process

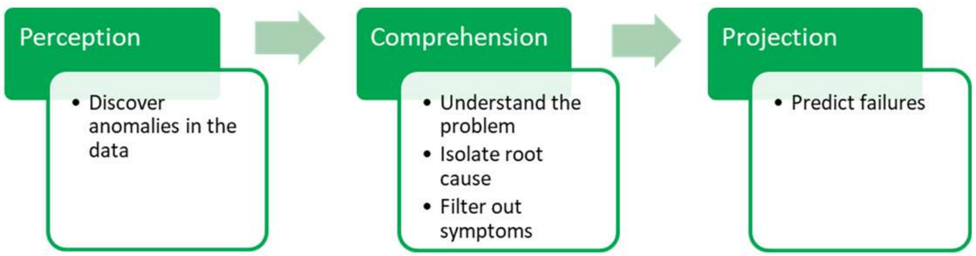

- Ongoing work
  - Evaluation of anomaly detection models based on ML for streaming event-based DC data
  - Investigation of methods for determining causal relationships among nodes in multi-dimensional graphs

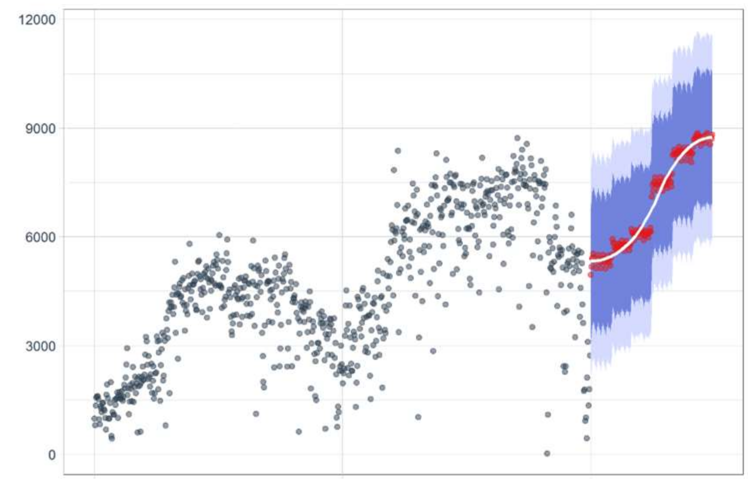

Example:

- Model data behaviour
- Use model to forecast next value
- Detect if new observed values represent anomalies

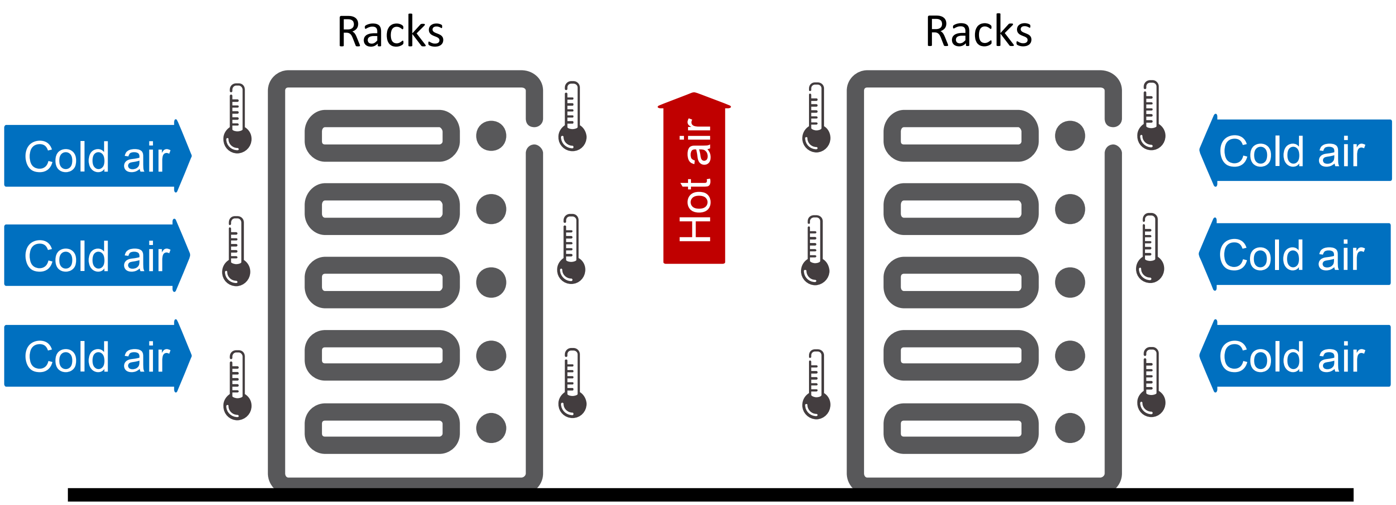
Racks
Racks
Cold air
Cold air
Cold air
Hot air
Cold air
Cold air
Cold air

hermal response - EDGE DC lab

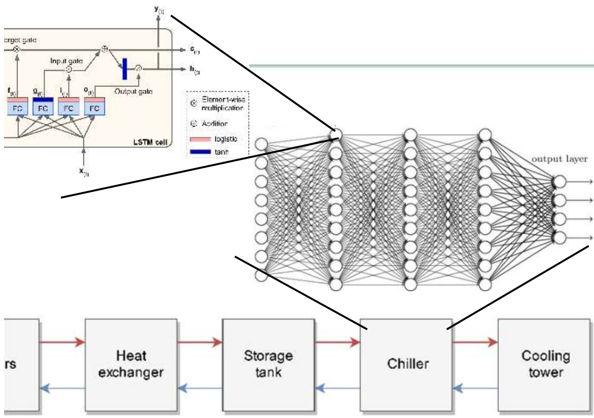

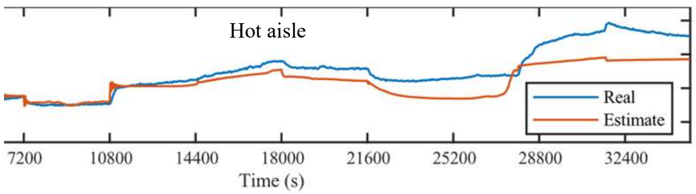

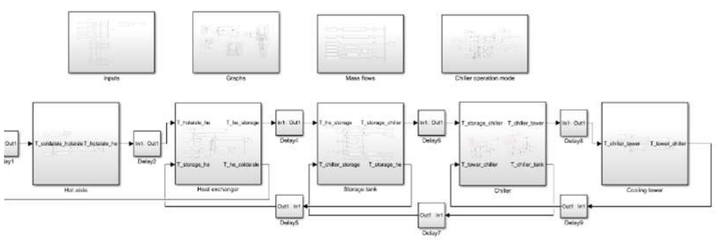

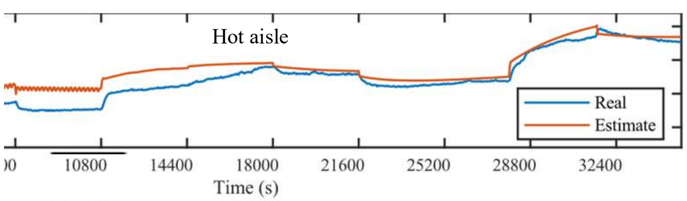

Risk-aware planning for PV /
- Markov chain for solar irrad
- take in public weather forec
- also electricity price forecas
predict tolerable workload at

PROJECT IDEA NUMBER: 37

ML based controller for small cluster
- hybrid = data driven physical model
- TF permits flexible cost functions
- set-point for CPU and/or hot-aisle
- also thermal-aware load balancing

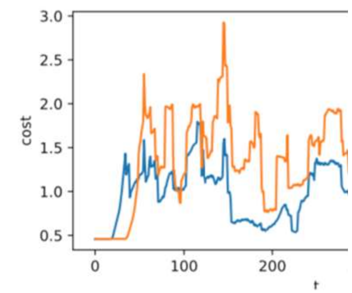

Ideas / next steps
- constrain ML/NN by physics e.g. fan laws, monotonicity,
- transfer learning by leveraging hierarchical priors
- generalizable vector embeddings predict DC metrics
- jointly model electrical and thermal parts for EDGE
- reinforcement learning using CFD to design control

T3.2: Observe and predict future operational states

PROJECT IDEA NUMBER: 37

In task T3.2 Mariner progressed on machine learning models for Electrical, Chilled Water and Natural Gas consumption. Further, Mariner has continued work on applying an AI-based anomaly detection approach to a new set of use cases aiming at detecting behavioral changes in the large groups of network elements.

Goal:

- Utilize Machine Learning algorithms to
  - Observe and predict future "operational states" of the target building facility (main inputs in the execution of the optimization strategies.
- Determine the accuracy of the "reward" functions (needed to determine the efficacies of the executed optimizations).

Activities:

- Feature Selection:
  - Research approach to automatically analyze provided input data (building energy consumption and operations readings) and determine a subset of candidate variables that should be used in the prediction of energy consumption.
- Reward Function Calculations (Energy Cost Functions):

Contribution from:
Aleksandar Petrovic <alex@marinerpartners.com>

hermal response - EDGE DC lab

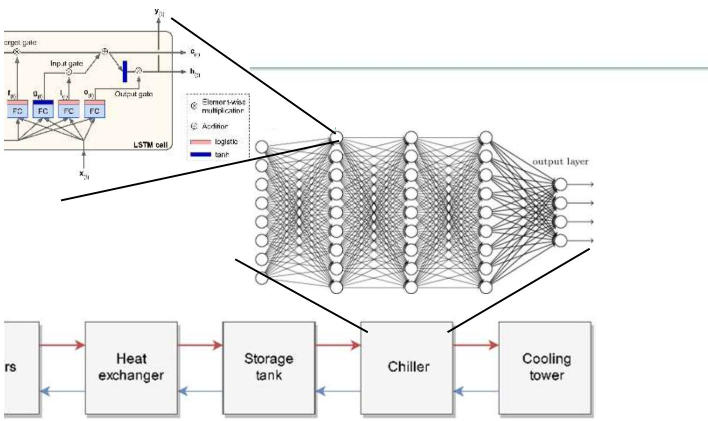

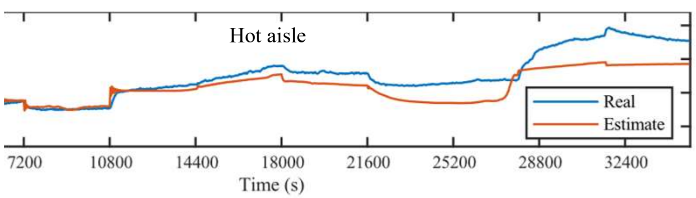

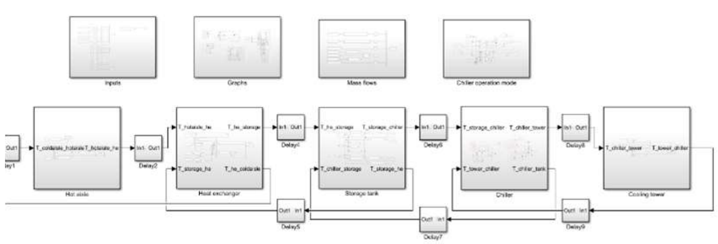

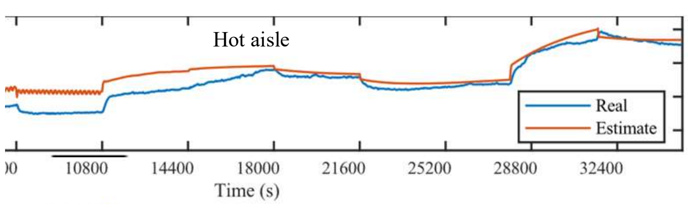

Risk-aware planning for PV /
- Markov chain for solar irrad
- take in public weather forec
- also electricity price forecas
predict tolerable workload at

PROJECT IDEA NUMBER: 37

ML based controller for small cluster
- hybrid = data driven physical model
- TF permits flexible cost functions
- set-point for CPU and/or hot-aisle
- also thermal-aware load balancing

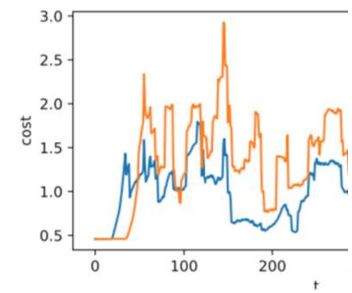

Ideas / next steps
- constrain ML/NN by physics e.g. fan laws, monotonicity,
- transfer learning by leveraging hierarchical priors
- generalizable vector embeddings predict DC metrics
- jointly model electrical and thermal parts for EDGE
- reinforcement learning using CFD to design control

Towards autonomous DCs

PROJECT IDEA NUMBER: 37