CEDAR

Cost Efficient heat pumps using DigitAl twins and Reinforcement learning

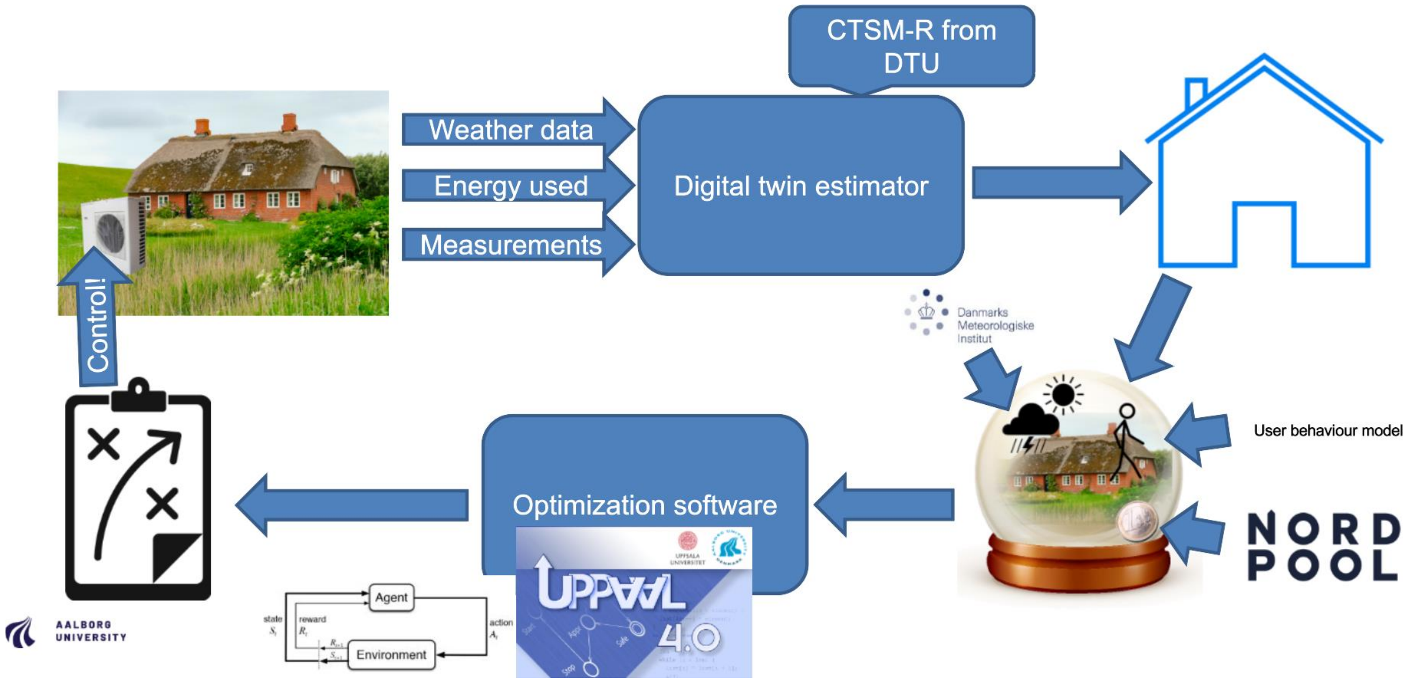

Figure 1: General flow of the approach studied in the CEDAR project.

Summary of project

The CEDAR project studies and develops next-generation technology for optimal control of heat-pump systems. In particular the project aims to construct a "install-and-forget" type of system for retrofitting in residential scale heat-pump systems.

The simplified flow of the envisioned solution is depicted in Figure 1. Given a single-family home, (1) observe the circumstances of the house (weather, energy consumption and internal temperature and humidity changes). Then (2) utilize the monitored data to construct a digital twin which is (3) decorated with auxiliary data-sources of the future context (weather forecast, future energy pricing, user behavior, …) to create a high-fidelity predictive digital twin. State-of-the-art stochastic optimization techniques (4) is then used to generate a strategy (5) for the future control of the heat pump. This process is then repeated over and over ad infinitum.

Internally the two core processes, namely the digital twin estimation and the stochastic optimization, relies on state-of-the-art techniques developed at the Technical University of Denmark (CTSM-R) and Aalborg University (Uppaal Stratego), respectively.

The novelty and strength of the approach lies in avoiding approximations and simplifications of the building dynamics to fit a specific Model Predictive Control (MPC) Framework. Instead the CTSM-R tool is utilized to estimate high-quality higher order thermodynamic models of a given building from the observed data – a model which also includes measures of disturbance from e.g. solar radiation. Importantly, such building models can be extracted with little or no knowledge of the physical building layout.

Conventional MPC frameworks require that such a model is abstracted or simplified into e.g. a linear model such that classical optimization techniques (Linear Programming) can be used. Instead, Uppaal Stratego provides a toolsuite for optimizing control of switchedcontrol non-linear, stochastic differential equation systems.

In particular, Uppaal Stratego utilizes a novel partition-refinement extension of classical reinforcement learning algorithms to provide near-optimal controller synthesis. The synthesis process itself is flexible with respect to the optimization criterion. This implies that the trade-off between cost and comfort can be adjusted freely according to any user specified function. In turn this opens for introducing pump-control flexibility such as temperature set-backs and target-bands which the optimization procedure can exploit towards an even higher savings by e.g. lowering the temperature at night.

A similar benefit of applying Uppaal Stratego is the ease at which peripherals such as accumulation tanks and photovoltaic power generation can be included as factors into the optimization problem; e.g. local power generation can be utilized when economically profitable, depending on the future energy demand of the house. As such, this project is in part envisioned as a stepping stone towards a holistic domestic energy optimizer for modern “prosuming” residential buildings.

In practical terms, the project is realized using off-the-shelve IoT sensor networks and edge-computing. We envision utilizing the inexpensive ZigBee family of devices for sensing and relying on low-powered mini-computers for the optimization and identification procedures. This facilitates a near offline application, ensuring robustness and stability.

Learnings and results

Preliminary studies using virtual house models demonstrate an up-to 30 %-40 % cost savings using the proposed method compared to a naive controller. Introducing user-specified flexibility (e.g. set-backs and target bands) an additional 11 %-point reduction has been demonstrated.

The CEDAR project aims to continue this work by (1) validating the laboratory results, (2) maturing the technological platform from research grade to consumer grade, and (3) further improve the technology of the digital twin estimation and the stochastic optimization procedure to ensure a robust and optimal control of heat pump units.

FACTS ABOUT THE PROJECT

IoT Category: Optimize heat pump operation
Goal: A self adapting and self optimizing cost efficient heat pump control. In particular facilitate that the heat pump control adapts to changes in the thermodynamics and changes in user behaviour.
Beneficiary: User, heat pump producers, Society
Data required: Weather forecast, day-ahead energy prices, sensors of the heat pump, temperature, and humidity sensors of the house.
Analysis method: Model generation via Continuous Time Stochastic Modelling and optimization via reinforcement learning on Euclidean Markov Decision Processes.
Modelling requirements: Fully data driven.
Quality-of-Service: Near real-time (minutes).
Time schedule: 2023-2024
Technology availability: TRL4 at project start.

Contact information

Peter Gjøl Jensen
pgj@cs.aau.dk