Aarhus School of Marine and Technical Engineering

HVAC system modeling using EcosimPro: A digital twin for advanced control studies & virtual commissioning simulations

Lasse Ahrendsen

Industrial Control Systems Group, Controls Engineering Section
Beams Department

CERN-THESIS-2022-376
15/06/2022

Bachelor thesis

Submitted in Aarhus, Denmark on the 8th of June 2022 for examination and oral defense on the 15th of June 2022 for the degree of Bachelor in Technology Management and Marine Engineering

Authors name and study ID:	Lasse Ahrendsen - V18414
Project title:	HVAC system modeling using EcosimPro: A digital twin for advanced control studies & virtual commissioning simulations
Educational institution:	Aarhus School of Marine and Technical Engineering (AAMS)
Internship company:	European Organization for Nuclear Research (CERN)
Project supervisor:	D.Sc. Brad Schofield (CERN)
Supervising professor:	Associate professor Esben Sørensen (AAMS)
Thesis hand-in date:	Wednesday 8th of June 2022 at 12:00
No. of pages:	86 (xi + 75) 49.64 (119’141 signs of 2′400 per page including spaces)

ABSTRACT

At the European Organization for Nuclear Research (CERN), the upper management has increased its focus on environmental matters within the last couple of years. The Energy Management Panel has selected three main strategies to minimize environmental damage: "*increase efficiency, use less, and recover waste energy.*"

New projects have been allocated resources, and within the Industrial Control Systems group, a Ph.D. project is currently underway on energy-optimal control. It investigates the possibilities of deploying advanced controls, namely model-based predictive control (MPC) on cooling- and HVAC plants at CERN, to increase energy efficiency and reduce power consumption. As many of the cooling systems at CERN are critical, some even affecting beam availability if not operating, it is not straightforward to intervene in cooling plants that already are operating satisfactorily but not fully utilizing an energy-efficient operation.

In conjunction with the Ph.D. project, a pilot project was launched, ultimately aiming to deploy an MPC controller on a HVAC plant at CERN. A virtual commissioning environment was considered essential for the succession because the MPC can be tested and validated using a digital twin.

In this thesis, the development and validation of a high-fidelity plant model are carried out for use in the pilot project. Previous use cases where virtual commissioning has been performed were investigated, and available data from the as-built documentation and operational data were used to parametrize the model. A different set of data were used to cross-validate the model.

A two-step validation study was performed on the model to determine its applicability to be valid for a summer scenario. Finally, the model was handed over for use in the Ph.D. project to ultimately perform virtual commissioning proving good results which can be used to evaluate the performance of the MPC, validate it, and aid future decisions regarding the deployment on the pilot plot.

ACKNOWLEDGMENTS

As part of the 9^{th} and last semester for the degree of Bachelor in Technology Management and Marine Engineering, I, the author, had to attend an internship of at least ten weeks and write the Bachelor thesis that definitively concluded the studies. The internship was held at CERN within the Industrial Control Systems group and was of ~4.5 months duration. The thesis was also written in collaboration with the group about the associated project I was involved with on model development for process simulation.

In the spring of 2021, I attended a ZOOM meeting directly from CERN arranged by the Union MMF (Maskinmestrenes Forening). Different Danish staff members at CERN took part and elaborated on their line of work. That planted the idea in my mind that I wanted to finish my degree in precisely this place in my last internship.

Shortly after, I began working on writing the application to ensure that I got a project offer; much work later, the application was sent. After a few weeks, I got a project offer that I was initially skeptical about. The work area was far from the core of my abilities and frightened me. However, after a few days of consideration and talks with my family, a great friend of mine who would also happen to intern at CERN and Jesper Nielsen, a fellow Dane and marine engineer, all contributed to convincing me that this was a great opportunity. So I accepted the project offer, and I have never regretted it.

4.5 months went by, and it has been a monumental challenge professionally. Occasionally I had to stop and remind myself to enjoy the experience while it lasted. Nevertheless, I am grateful to have been offered this opportunity where I have gained valuable professional experience in CERN's unique place.

First, I would like to thank my project supervisor, Brad Schofield, for listening, being helpful, and understanding the difficulties I have met along the way. Faiq Ghawash, the Ph.D. student, for good cooperation when needed and for many interesting conversations, including them on model predictive control, which was entirely new for me. Esben Sørensen, my supervising professor, for guidance along the way and his visit on site in Prévessin, France.

Lastly, a thank you to my great colleagues in the Industrial Control Systems group who have made it easier for me, and especially thank you to my great office mates for having contributed to being part of the best office in the whole of building 864.

Aarhus, June 2022

Lasse Ahrendsen

REFERENCES

ABB, 2022a. *Model-Predictive-Control MPC.* [Online]
Available at: https://www07.abb.com/images/librariesprovider104/Industry-solutions/Oil-and-gas/past-and-future.jpg?sfvrsn=1
[Accessed 28 April 2022].

Belimo, 2019a. *Actuators for heating, ventilation and air conditioning.* [Online]
Available at: https://bit.ly/38Ibqpi
[Accessed 1 Maj 2022].

Booth, W., Viñuela, E. B., Bradu, B. & Sourisseau, S., 2017. Virtual Control Commissioning for a Large Critical Ventilation System: The CMS Cavern Use Case. *Proc. 16th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'17)*, October, pp. 92-98.

CERN, 2020a. *Environment report 2017-2018.* [Online]
Available at: https://hse.cern/environment-report-2017-2018/energy
[Accessed 23 February 2022].

CERN, 2021a. *Environment report 2019-2020.* [Online]
Available at: https://hse.cern/environment-report-2019-2020/energy
[Accessed 23 February 2022].

CERN, 2022a. *What we do.* [Online]
Available at: https://home.cern/about/what-we-do
[Accessed 6 February 2022].

CERN, 2022b. *Our history.* [Online]
Available at: https://home.cern/about/who-we-are/our-history
[Accessed 6 February 2022].

CERN, 2022c. *Birth web.* [Online]
Available at: https://home.cern/science/computing/birth-web
[Accessed 6 February 2022].

CERN, 2022d. *CERN's accelerator complex.* [Online]
Available at: https://home.cern/science/accelerators/accelerator-complex
[Accessed 16 March 2022].

CERN, 2022e. *LHC design report - Beam.* [Online]
Available at: https://lhc-machine-outreach.web.cern.ch/beam.htm
[Accessed 23 February 2022].

CERN, 2022f. *Pulling together superconducting electromagnets.* [Online]
Available at: https://home.cern/science/engineering/pulling-together-superconducting-electromagnets
[Accessed 6 February 2022].

CERN, 2022g. *Large hadron collider.* [Online]
Available at: https://home.cern/science/accelerators/large-hadron-collider
[Accessed 6 February 2022].

CERN, 2022h. *Vacuum empty interplanetary space.* [Online]
Available at: https://home.cern/about/engineering/vacuum-empty-interplanetary-space
[Accessed 6 February 2022].

CERN, 2022i. *Storage.* [Online]
Available at: https://home.cern/science/computing/storage
[Accessed 6 February 2022].

CERN, 2022j. *Data centre.* [Online]
Available at: https://home.cern/science/computing/data-centre
[Accessed 6 February 2022].

CMS Experiment, 2008a. *CERN.* [Online]
Available at: https://home.cern/resources/image/experiments/cms-images-gallery
[Accessed 6 February 2022].

Cooling & Ventilation Group, 2021a. *As-built PJ 312 SR1 HVAC.* [Online]
[Accessed 1 May 2022].

EA International, 2022a. *EcosimPro.* s.l.:s.n.

GIS Portal, 2022a. *GIS Portal CERN.* [Online]
Available at: https://gis.cern.ch/gisportal/IT_Equipment.htm
[Accessed 12 April 2022].

Larsen, S. B., 2018. *Projekter & Rapporter på Tekniske Uddannelser.* 1. ed. s.l.:Hans Reitzels Forlag.

NXCALS, 2022a. *CERN.* [Online]
Available at: https://nxcals-docs.web.cern.ch/current/
[Accessed 2022 March 2022].

Schofield, B., Viñuela, E. B., Booth, W. & Peljo, M. O., 2019. Waste Heat Recovery for the LHC Cooling Towers: Control System Validation Using Digital Twins. *International Conference on Accelerator and Large Experimental Physics Control Systems (17th)*, 05-11 October, pp. 520-524.

Schwenzer, M., Ay, M., Bergs, T. & Abel, D., 2021. Review on model predictive control: an engineering perspective. *The International Journal of Advanced Manufacturing Technology*, 1 November, Issue 117, pp. 1327-1349.

UNICOS, 2022a. *About UNICOS.* [Online]
Available at: https://unicos.web.cern.ch/
[Accessed 17 March 2022].

Wiedemann, H., 2015. Particle Accelerator Physics. In: S. Hassani, et al. eds. *Particle Accelerator Physics.* Fourth Edition ed. Chiang Mai(Stanford): Springer, pp. 11-12.

LIST OF FIGURES & TABLES

CONTENTS

1

INTRODUCTION

This chapter will briefly introduce CERN (with more background information in **2.1**) in the first section. Then, section **1.2** will derive and state the problem from the theses associated project, and section **1.3** will present the problem formulation with a research question and related objective. Furthermore, in section **1.4,** the delimitation of the thesis is argued.

1.1 CERN

CERN, short for Conseil Européen pour la Recherche Nucléaire, is a European scientific research organization, primarily within particle physics, which studies the smallest particles that make up matter to answer fundamental questions about the universe. The research also include broader areas but are not exclusive to high-energy and nuclear physics, studies on antimatter, and cosmic radiation (CERN, 2022a).

CERN was born in the time after World War 2 as European science was no longer world-class. A few visionary scientists had the idea to create a European atomic physics laboratory, and thus the seed was planted. Long story short, CERN was a reality. It was decided that Geneva, Switzerland, would be the ideal placement for CERN due to its placement in Europe, neutrality in the war, and the existing international organizations already based in Geneva.

Fast forward a few years to 1954, and the first shovel of earth was dug at the Meyrin site right at the Franco-Swiss border, and already in 1957 was the first accelerator, the Synchrocyclotron (SC), started up. The accelerator was operational for 33 years and closed in 1990, initially providing beams for particle physics research until 1964 before newer and more powerful accelerators came into operation. In 1964 the focus shifted to nuclear physics alone. In 1967 it started to supply beams for the Isotope mass Separator On-Line facility (ISOLDE), which researched areas ranging from pure nuclear astrophysics to medical physics. The SC provided beams until the closure in 1990, when the ISOLDE facility was transferred elsewhere at the accelerator complex.

Today, CERN is run by 23 member states instead of the 12 founding member states, which each have two official delegates in the CERN Council. One represents the respective nations' scientific interests, and the other represents the individual nations' governments' administration. Denmark has been a member state of CERN right since the very start, and before the experiments could take place, the theoretical work was carried out in Copenhagen, Denmark (CERN, 2022b).

CERN has, over time, become a world-famous institution known for many contributions both inside and outside of physics that have had an impact on the society we live in today. The best-known example to the layman is in 1989, when a British scientist working at CERN, Tim Berners-Lee, developed the World Wide Web on CERN soil to make it easier for scientists worldwide to share information during their collaborations (CERN, 2022c).

1.2 Problem statement

At CERN, energy- consumption, and efficiency are becoming increasingly important topics in a world where awareness of the environmental impact humans have on planet earth is rising. Within the latest years, the upper management at CERN has raised questions and started to enforce a not seen before focus on environmental matters. For the first time in 2020 and again in 2021, a public environment report was released to report transparently on performance records, objectives for improvement, and lastly, to reflect a proactive approach to environmental protection across the laboratory of CERN. In the report, the Energy Management Panel (EMP) presents three main strategies for energy: *"increase efficiency, use less, and recover waste energy"* (CERN, 2020a) (CERN, 2021a).

In making CERN's operation successful, many cooling systems are present, which are considered critical installations. The laboratory of CERN consumes a significant amount of electrical power, especially when operating and conducting experiments. One of the prominent consumers of the accelerator complex is the cryogenic- and associated systems. These systems enable the operating conditions of the superconducting magnets to be maintained, and consequently, a considerable amount of excess heat is produced, which must be removed and dissipated. Several different cooling plants perform this but removing substantial heat energy also requires much electrical energy. In the Industrial Control Systems group concerning the EMP's strategy of increasing efficiency, it is interesting to focus on the control systems of the cooling plants and the impact that these have on energy- efficiency and consumption.

Many of the cooling plants are deemed critical infrastructure, with some even affecting beam availability if not operating efficiently during LHC run. Consequently, it is essential to ensure that they can run efficiently and to specification at all times needed. Typically, they have constraints in terms of interventions, with some operating 24h/7d and others simply having a limited time to perform the actual intervention. It makes them sensitive to unforeseen problems, which could lead to delays or downtime.

The cooling plants, including their control system, can thus be considered critical. Moreover, the plants are already operating satisfactorily, though not necessarily energy-efficient. Since the plants already operate satisfactorily concerning availability and in their desired working regime, there has not been a strong incentive to modify the controls.

At CERN, energy optimization potentials exist in cooling plants, including their control systems. That fact, coupled with the increased focus of CERN on energy efficiency and environmental matters, has resulted in new projects arising and allocated resources that center around the strategies proposed by the EMP.

There is currently a Ph.D. project in the Industrial Control Systems group on energy optimal control using model-based predictive control (MPC) on cooling- and HVAC plants. The case is that more advanced control methods have not found use in CERN's variety of cooling plants, where mostly more traditional regulation loops with PID controllers are found. Typically, they have a single control objective as temperature regulation, but variable frequency-driven fans, among other actuators, can be over-actuated while reaching the control objective. This results in a waste of energy. Using advanced process control such as MPC can solve control problems more efficiently and reach a near-optimal operation with respect to energy usage. The Ph.D. project aims to investigate the possibilities of using MPC on cooling systems widely in the frame of CERN. The overall long-term objective is to establish knowledge and competencies within the group enabling the opportunity of operationalizing the deployment of MPC controllers into the production environment. Eventually, the Industrial Control Systems group may be able to provide services to relevant users at CERN interested in the possibilities more advanced controls offer.

In conjunction with the Ph.D. project, a pilot project has been launched to deploy an MPC on a carefully chosen plant. However, before deploying a control system on a critical cooling plant, it is convenient to evaluate the MPC beforehand in a safe and controlled environment. It can be done by virtual commissioning.

EcosimPro is a modeling and simulation tool of dynamic systems in which process plant models can be built and simulated as so-called digital twins. Using such a model as part of the virtual commissioning may validate the proposed MPC for the chosen pilot plant. A functioning plant model can be used to replace the actual plant for testing and validation purposes of the MPC. The potential to be realized is to identify problems or deficiencies and correct them, ultimately easing the existing plant's implementation process and run-in period.

Within the Industrial Control Systems group, a considerable effort has already been made to create models for process simulations on numerous applications at CERN. These have successfully been used in the past for different purposes. Since a flawless deployment of the MPC is expected, it is instrumental to use a process plant model developed in EcosimPro to perform virtual commissioning simulations of the MPC for successful validation.

1.3 Problem formulation

Research question

> ***How is it possible to create a high-fidelity digital twin of a cooling system for use in advanced controls?***

Objective

To take part in R&D activities of the Industrial Control Systems group within the simulation domain and develop a model of the chosen pilot plant. Ultimately, the plant model will be used to perform virtual commissioning for validation purposes of the MPC controller developed in the Ph.D. project.

It is necessary to establish a foundation of knowledge based on previous cases where virtual commissioning has been performed. Furthermore, historic control system data obtained from the actual plant is used to parametrize the model. Also a different set of historical data will be used to validate the plant model.

1.4 Delimitation

A great deal of work has been put into estimating model parameters against the operational data, and for this purpose, many different datasets throughout an entire year were applied. The external disturbance, the fresh air temperature changes, and the plant operation's nature are dictated by its control system. It is not unnormal to have two or even more sets of model parameters that are fitted to different seasons throughout a year. However, with the short timeframe of the project, it was decided to mainly focus on a single set of parameters fitted to the warmer months of the year. In such a scenario, the fresh air temperature is hot, and active cooling mode on the plant can be expected. There has still been effort into, e.g., the heating coil. Still, the validation studies will focus on a single validation dataset representing the described scenario because of the delimitation. Additionally, deficiencies and problems regarding the heat exchanger model present in the internally developed library of HVAC components in EcosimPro proved to cause some issues. Especially when using the component as a heating coil applying historic control signals to the valve caused simulation problems and further added to the above-argued delimitation.

2

Background

In the first section, this chapter will give a brief insight into what kind of research is being conducted at the laboratory of CERN, with subsection **2.1.1** explaining some of the technology behind the Large Hadron Collider (LHC). Then, in section **2.2**'s three subsections, more background information will be presented.

2.1 Research at CERN and the Laboratory

A wide variety of research is being conducted at CERN both on a theoretical and experimental level, where the main field is particle physics. As CERN is a laboratory, it is eminent to conduct experiments, assess current physics theories, and reach a greater understanding of the universe. With the particle accelerators present at CERN, research is being conducted for different purposes undertaken at the respective experiments. For example, the LHC and its four main experiments, ALICE, ATLAS, CMS, and LHCb, aim to collide particles at the highest possible energies and discover the smallest particles and their interactions in the universe. Almost opposite in the Antimatter Factory, the Antiproton Decelerator (AD) and its associated experiments aim to decelerate the particles to low energy levels trying to create antiprotons for studies on antimatter.

The LHC is the most powerful particle collider globally and has already contributed to extraordinary discoveries in its short lifespan, most notably the observation of the Higgs boson particle.

2.1.1 LHC at the accelerator complex

The biggest and the most famous particle accelerator at the CERN accelerator complex is the LHC (**Figure 1**). It is the most powerful particle accelerator ever built and is situated entirely underground in the old tunnel of the now replaced Large Electron-Positron collider (LEP). The LHC is a circular collider with a circumference of approx. 27 km with the collider both being under the French- and Swiss side of the border. Even though the LHC is the largest particle collider at CERN and globally, it is part of an accelerator complex that serves the LHC and other experimental programs for different applications of accelerators.

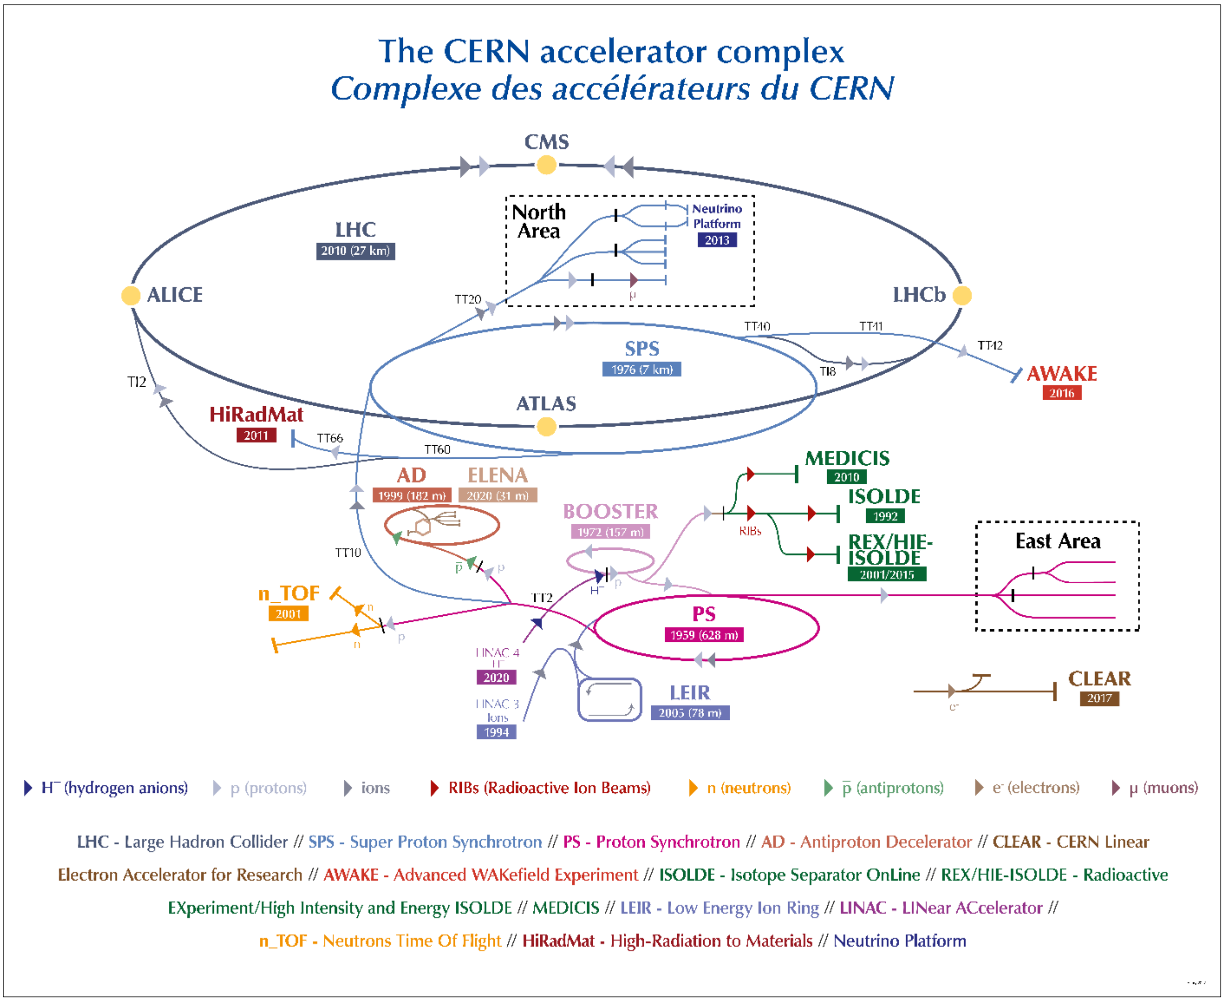

Figure 1: CERN's Accelerator Complex colorized and visualized how the different accelerators are interconnected (CERN, 2022d).

All accelerators and experiments are in some way linked together, and for the LHC to work, the particles must be increasingly accelerated to higher energy levels through several smaller accelerators starting from the LINAC 4 through the BOOSTER, PS, SPS, and finally, the LHC. The complete accelerator complex serves a rich and diverse array of experimental programs.

The source of proton particles used for the collisions originates in the form of hydrogen contained in a small bottle. However, only a tiny amount is conserved and needed for the purpose. At the beginning of the acceleration process, these protons are separated from hydrogen atoms and are increasingly accelerated to higher energy levels. The protons travel in bunches 30 cm long and contain $1.15e10^{11}$ protons at the start of nominal fill. The beam itself is made up of $2'808$ bunches which at top energy is travelling 7.5 m or 25 ns apart at 0.999999991 times the speed of light which is more than $299'000$ km/s.

In accelerator physics, the energy is indicated and measured in electron Volt $[eV]$ and defined as follows: "*1 eV is the kinetic energy a particle with one basic unit of electrical charge e $[C]$ would gain while being accelerated between two conduction plates at a potential difference of* 1 V" (Wiedemann, 2015). For the LHC, the energy equivalates to ~13 TeV (during LHC run 2), which is the maximum energy the two beams could release. To put it into perspective 1 eV is only $1.60217733e10^{-19}$ J, but the fact that there are so many particles accelerated to these higher energy levels amounts totally to 724 MJ of energy for the two beams (CERN, 2022e).

When injected into the LHC ring, the beams are separated into two beam pipes where they travel in the opposite direction and are ultimately steered to collide at different experiments around the ring, mainly the four experiments: ALICE, ATLAS, CMS (**Figure 2**) and LHCb where detectors are placed, and data can be obtained (CERN, 2022g).

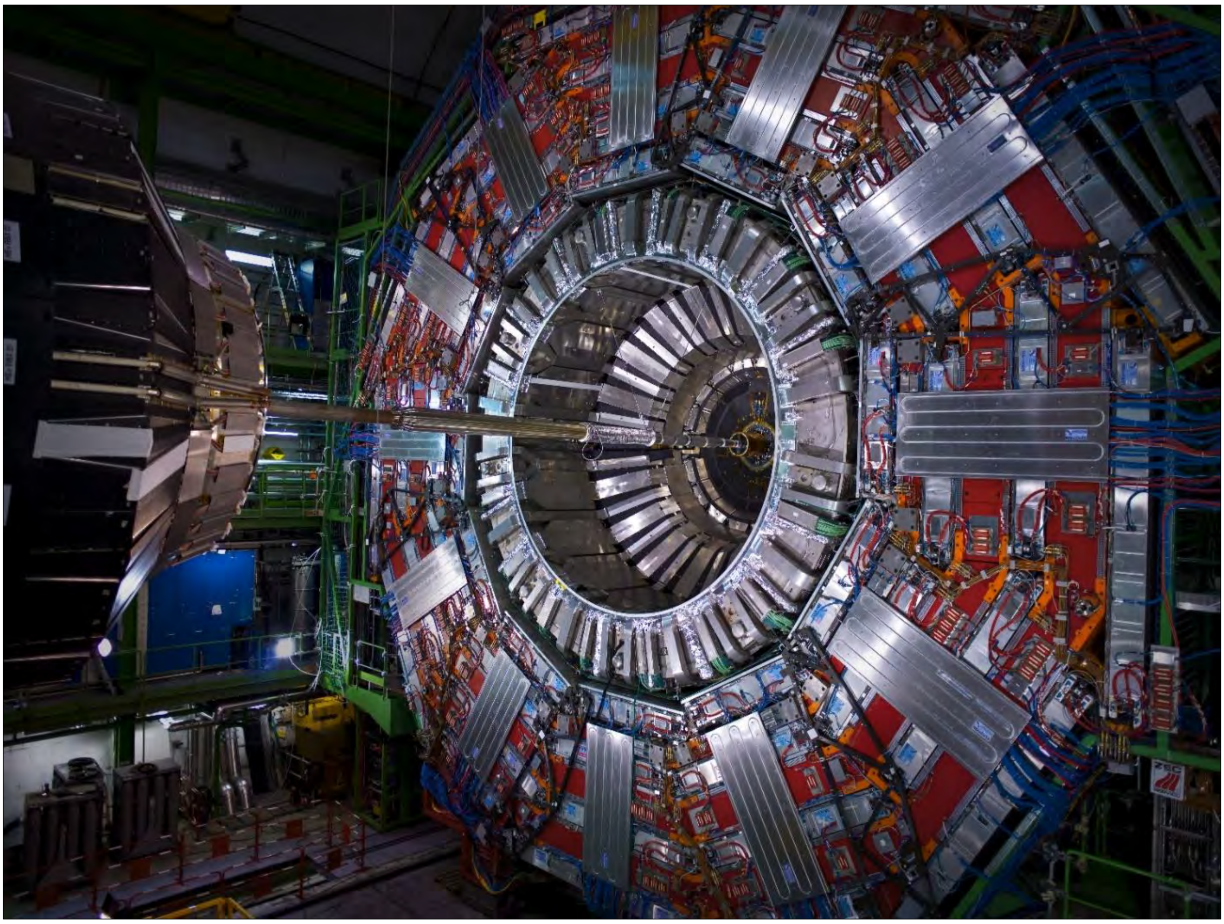

Figure 2: The detector of the CMS experiment (CMS Experiment, 2008a).

During operation, the collisions happen continuously until the level of protons reduces to a level where a new fill is needed to keep a high luminosity, which is the rate of collisions occurring. With high energies necessary for emulating conditions during the Big Bang, the luminosity is also particularly important as the collisions provide obtainable data.

Main contributing technologies

Reaching the energy levels mentioned in **2.1.1** does not come easily. With the LHC construction, the machine's size, the largest ever built by humanity, and the complexity of necessary auxiliary systems took a big step up from earlier accelerators. The LHC consists of a ring of different superconducting magnets (**Figure 3**) with dipole magnets to steer the trajectory of the beams around the ring and other magnets as the radiofrequency (RF) cavities that contribute to the acceleration of particles. A superconductive state for the magnets is needed. Therefore, there is an enormous distribution system of liquid Helium that enables the possibility of cooling down the coil of a unique electric cable made of Niobium–titanium within the magnets to 1.9 K. With the superconductive state achieved, it is possible to conduct electric current without resistance. $11'000$ A of direct current is applied to reach strong magnetic fields of 8.3 T. These magnets are essential for the succession of the LHC's purpose but are producing excessive amounts of heat which must be removed (CERN, 2022f) (CERN, 2022g).

Figure 3: Technicians measuring on the connecting busbar between two magnets in the LHC tunnel (CERN, 2022g).

CERN also houses one of the most significant cryogenic cooling systems globally, with Helium as the refrigerant. It also houses an extensive vacuum system. It is necessary to create an ultrahigh vacuum of up to 10^{-14} atmospheres in the beam pipes for the particles to travel at the speed of light. Other vacuum systems are also used in the LHC system as an insulator for the cryogenic system, which experiences enormous differences in temperature in the order of a ΔT of 300 K facilitating a significant and undesirable heat transfer (CERN, 2022h).

Exceptionally enormous amounts of data are stored in CERN's data center from the experiments. For example, the LHC alone produces ~90 petabytes $90e10^6$ gigabytes per year, which is a vast amount of data that needs to be sorted, structured, computed, typically by external computation power, and analyzed by the physicist to test out their theories (CERN, 2022i). The external computation power is available at the CERN Data Centre with $470'000$ processor cores and $11'000$ servers running 24h/7d (CERN, 2022j). The data centre also produce a vast amount of heat which must be removed by cooling plants.

2.2 Simulation & advanced control

Two of the areas that the Industrial Control Systems group and provides services of are within the domains of simulation and advanced controls. As these are of interest in this thesis, an introduction will be made to both. In the first subsection, previous use cases at CERN where virtual commissioning was performed using EcosimPro are investigated to establish a solid foundation for the project by drawing on experts' experience who have done projects of similar character in the past.

In subsection **2.2.2**, the simulation activities with EcosimPro will be explained more in detail as it is the primary tool used to develop- and simulate the model. In addition, it includes the working principles needed to conduct the project's work.

In subsection **2.2.3**, a brief introduction to model-based predictive control is given as it is not expected that the thesis reader has any prior knowledge of it.

2.2.1 Virtual commissioning: previous cases

Virtual commissioning has been performed at CERN for many years. In this context, in the Industrial Control Systems group. Previous cases exist, and two papers will be used to establish a foundation of knowledge for the project on model development for process simulation. The area is of no prior experience to the author.

The *"Virtual Control Commissioning for a Large Critical Ventilation System: The CMS Cavern Use Case"* presented in (Booth, et al., 2017) and *"Waste Heat Recovery for the LHC Cooling Towers: Control System Validation Using Digital Twins"* presented in (Schofield, et al., 2019) will be investigated.

The CMS cavern case

The case was a big project undertaking two years of work to provide a simulation model of the CMS cavern HVAC plant that could be used in the virtual commissioning of the upgraded control system, which was to be implemented during the LHC long shutdown 2 (2019-2020). It is noted that the CMS plant is running 24h/7d and is critical during LHC run, after running, and during interventions, where interventions have a limited timeframe.

Since the CMS ventilation system is critical, the virtual commissioning environment was deemed preferable. The aim was to ensure a less problematic and efficient upgrade of the control system performing virtual commissioning prior to the actual deployment of the control upgrade on the real plant.

The project's starting point was to continue the development of the CERN HVAC library in EcosimPro, which was initiated earlier, as the state of the library was not sufficient for the ultimate purpose of the project. Then, with both new components developed and existing ones improved, they moved to the next stage, which was to develop three different plant models as test models to prove the worth of the library and the tool in its entirety by testing and validation.

A significant difference between the test models and the CMS plant was the lack of operational data for the latter. The SCADA systems and PLCs were old before the upgrade and did not log data into the accelerator logging system as newer plants employing WinCC OA as SCADA do. On the other hand, the CMS plant development depended on reliable input data to aid in the modeling work (**Figure 4**) and was modeled from first principles.

The case with the test models is that besides the available input data, much control system data was available, which can be used in identifying the plant and estimating model parameters. This is especially useful for the modeling work and subsequent validation of the model, which is why the modeling work of the CMS plant was prone to be vastly more complex.

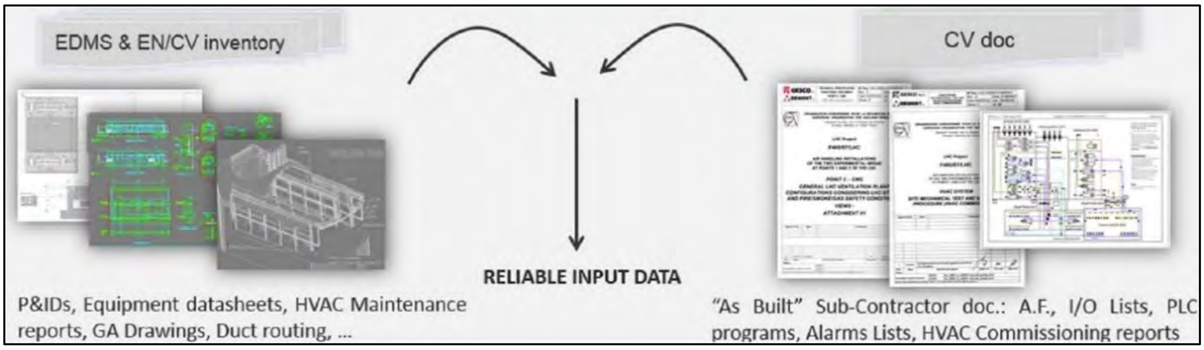

Figure 4: Retrieval of all available data of the CMS cavern HVAC plant (Booth, et al., 2017).

With the data available from the test plants (excluding the operational data), some model parameters cannot be set, and the operational data can be used. However, a considerable amount of trial and error estimation was required since other solutions are limited. Examples as room heat load, damper authorities, and valve parameters had to be estimated by tuning these model components against the operational data.

For the test models, their existing control system configuration was implemented. The models could do closed-loop simulations with the fresh air temperature- and humidity disturbances as outside inputs for a given period. These were implemented so that it was possible to also do simulations with a real PLC connected. At the time, the future control upgrade was under development. A preliminary control algorithm based on an initial functional analysis was implemented and tested in the PLC lab (**Figure 5**). The development was concluded with the setup used in the actual virtual commissioning of the control system upgrade.

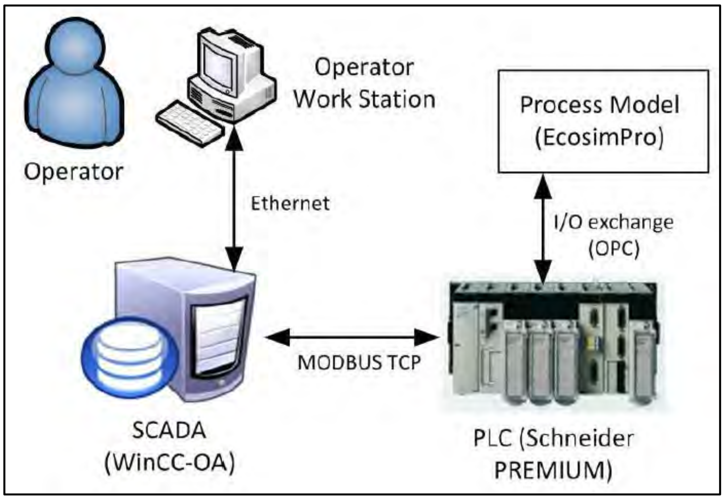

Figure 5: PLC lab test setup for virtual commissioning (Booth, et al., 2017).

The WHR case

In the waste heat recovery (WHR) case by (Schofield, et al., 2019), virtual commissioning was performed in slightly different circumstances and for a different purpose. Cooling towers are used at several CERN locations to remove massive quantities of excess heat. Different clients transfer heat through a secondary cooling circuit to the primary cooling circuit, the cooling tower open-loop water circuit. In the combined return line from clients to the cooling towers, a WHR plant is to be operated in the future, which was the reason for concern in the study.

As the cryogenic clients are critical to the operation of the LHC, concerns were raised about whether the impact of a sudden shutdown of the WHR plant would result in a too high supply water temperature. In addition, it was questioned whether the current control system configuration would be able to cope with such a scenario within an acceptable timeframe and effectively counteract a shutdown of the critical cryogenic systems.

The cooling tower model already existent in EcosimPro proved insufficient in this case, with the towers old and unique to CERN with no data available. Instead, a significant amount of operational data was available. Therefore, the development of a new cooling tower model was initiated. A significant amount of published work is available on evaporative cooling towers, and extensive work was put into finding a suitable approach to model the cooling towers mathematically. A semi-empirical approach was chosen using the available operational data, and the project's central area was centered around that purpose.

When the model was validated and concluded, a production-ready version of the current control system configuration loaded into a PLC was used on the developed model as the virtual commissioning setup. Various fault scenarios, including the one mentioned above, were simulated (**Figure 6**). It was found that a sudden loss of the WHR in nominal conditions results in a sharp rise in the common line return water temperature. However, a very slight increase in supply water temperature is seen, way under the requirements. The model and virtual commissioning setup were successfully used to conclude that the existing control system configuration could cope with the different fault scenarios that were simulated.

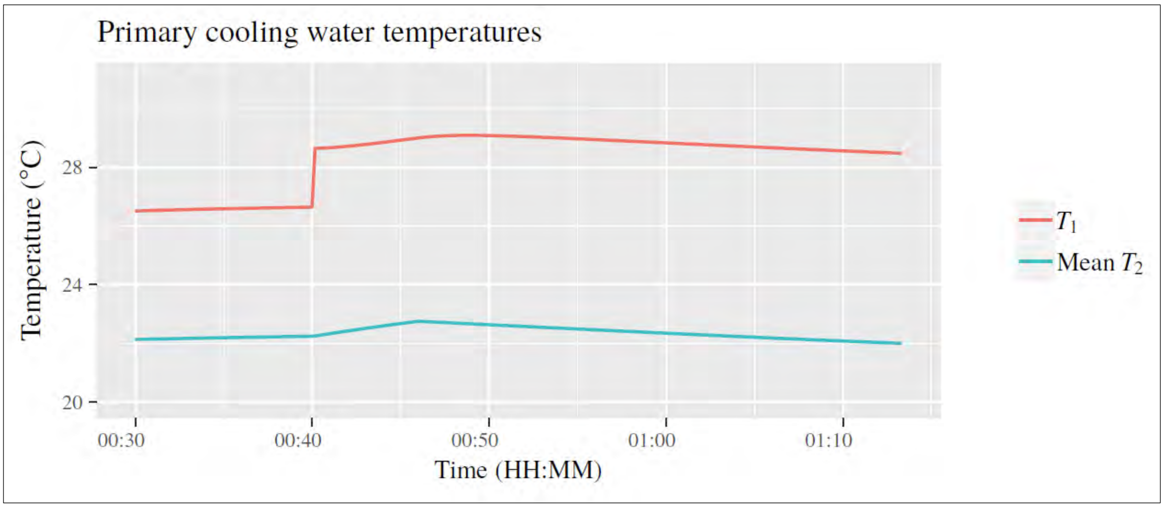

Figure 6: Effect on the return- and supply water temperature during a WHR shutdown in nominal conditions (Schofield, et al., 2019).

Summary

For both cases investigated, different approaches to the modeling work in EcosimPro were used, extensively using the operational data available for the plants in different ways (in the CMS case, for the test models). Different objectives were also present for each case, where both used a virtual commissioning setup to fulfill them. Not precisely alike scenarios, both cases centered around the control systems, with the CMS case to eventually ease the implementation of an upgraded control system and the WHR case to evaluate its existing control system during specific fault scenarios. For both cases, a virtual commissioning setup was deemed a good option as it can provide excellent results in a safe environment, not interrupting the existing system.

Especially in the CMS case, extracting all available data was emphasized. It was used in the modeling work, which required extensive knowledge and depth into the specific process plant, and consequently, valuable insight into the systems was obtained. Different malfunctions and component deficiencies were discovered in the CMS case, which negatively affected the regulation. In the WHR case, the current control system implemented was concluded to efficiently reject a sudden heat load increase caused by the shutdown of the WHR plant. It was discovered to operate very energy-inefficiently hence in both cases, the insight into the processes provided invaluable insight. It resulted in the discovery of operational- and energy-efficiency optimization potentials.

2.2.2 Modelling & simulation with EcosimPro

EcosimPro is an object-oriented mathematical tool used to model an extensive range of dynamic systems and run simulations as experiments. The created models can be simplistic without integrated control systems or fully equipped, highly complex representations of a real system, the so-called digital twins.

EcosimPro has been extensively used in several domains at CERN as the Cryogenics group and for cooling and ventilation systems. Cryolib is a cryogenics-specific library developed in collaboration with CERN, and the development of the CERN HVAC library mentioned in 2.2.1 took place in the Industrial Control Systems group in the past.

Some areas that EcosimPro has been used for:

- Virtual commissioning of control systems
- Training of operators
- Optimization
    - Deriving the most appropriate control strategies to apply
    - Control optimization with tests of control strategies in simulations
    - Controller optimal tuning

EcosimPro is intuitive for the engineer, modeling systems and reversing the process where typically 80 % is programming and 20 % is modeling. Ideally it should be possible to build the model by connecting components picked out of the relevant library, put in parameter data, and run a simulation.

Then the programme internally handles all the variables and equations and solves the mathematical model of the system.

Components as valves, heat exchangers, fans among others are picked out of the relevant library and drag-and-dropped into the schematic connecting them by lines to in- and out ports. The modeling is intuitive as one can use components already modeled and relatively quickly create a model than can run experiments. However, a full model consisting of many components will have different individual model parameters that must be set or estimated, e.g., using operational data.

Creators of EcosimPro, EA have defined four different levels of users. These are explained in the following **Figure 7**:

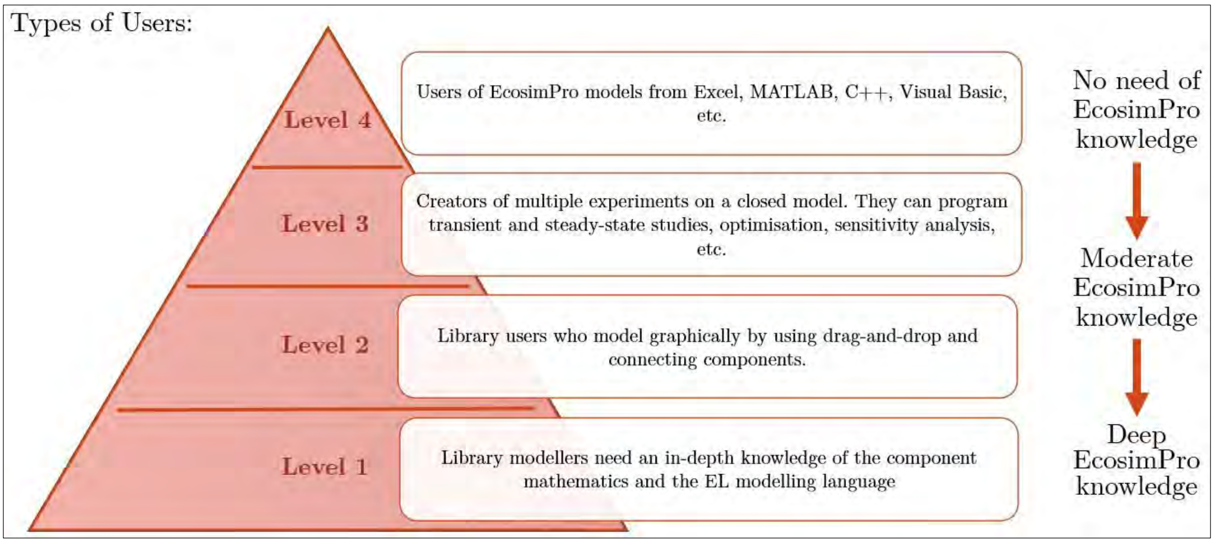

Figure 7: Different user levels in EcosimPro (EA International, 2022a).

Use of EcosimPro in the context of the thesis

In this thesis, the user level will be at level 2, modeling a plant from scratch, using existing components models in the internally developed CERN HVAC library. User-level 2 requires a moderate to deep knowledge of EcosimPro and an equal understanding of the theory behind the components used and process plant technology, specifically HVAC plants. To perform the modeling work successfully, assuming a sound library of components, the understanding of HVAC plants is used to obtain data for setting parameters and do simulations to analyze behavior and estimate other parameters where data cannot be obtained. There were frequently more minor investigations into the components (level 1) during the modeling work, where they are mathematically modeled and written using the EL language. It was used to better understand the components and their behavior during the modeling work and provide answers when needed. No other consideration to make changes at this level was made. It requires the highest possible knowledge of the theory and mathematics behind the components and the EL language itself.

2.2.3 Advanced control with model-based predictive control

As the readers of this thesis are assumed to understand PID control, this subsection aims to briefly introduce the concept of model-based predictive control, which is considered a more advanced method of control than a PID. It will enable the reader to understand how an MPC controller differs from a PID controller at a superficial level.

The MPC concept

Optimal control has spawned the concept of model predictive control over time, but it is not new. In the 1980s, MPCs were deployed on advanced process plants in the chemical- and oil refinery industries. Today it is still considered an advanced process control method, with research still being conducted in the area.

The concept differs substantially from traditional feedback controllers such as PID controllers. First, one must rely on a good system model to achieve reasonable control which can be achieved using a mathematical model of the system derived from operational data from the plant.

The MPC is an optimization strategy based on a model of the system; why the most critical asset of an MPC is the dynamic model of the process plant, which is used to forecast system behavior and optimize the forecast to produce the best decision - the control move at the current time. Therefore, models are at the very center of the MPC concept.

The MPC will run forecasts of the model forward in time for different actuation strategies and optimize over the control input in a predicted time horizon and essentially determine the immediate next control action based on that optimization. When the control action is applied, there will be a reinitialization of the optimization where the time window is moved a time Δt to find the next immediate control action. The control input u over the predicted horizon will change at each time step as the predicted output, and measured output deviates over that exact period. While u is optimized over the period at several time steps that will bring the desired value y closer to its objective, only the next immediate control action at the first time step is applied before reoptimizing once again (Schwenzer, et al., 2021). It is visualized in the following **Figure 8**.

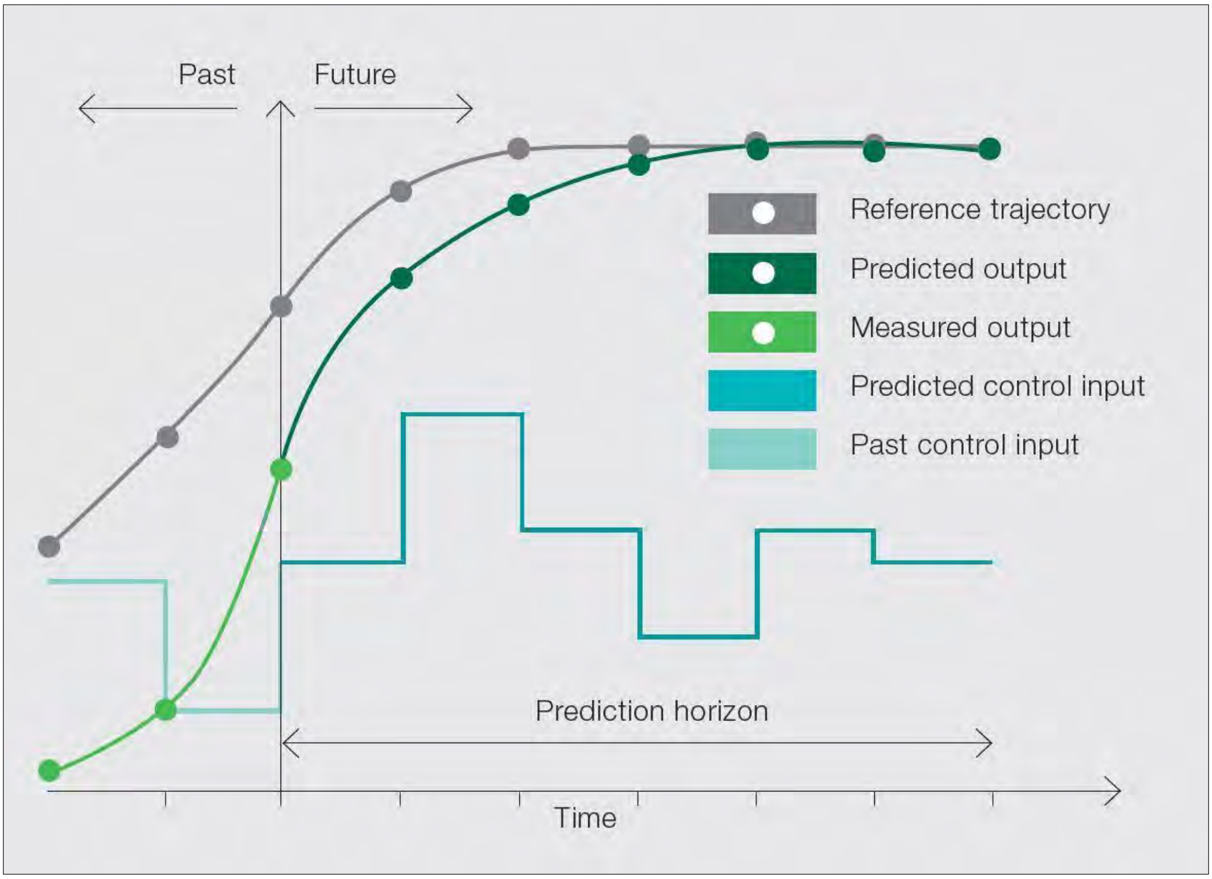

Figure 8: Visualization of how the MPC works with the predicted horizon (ABB, 2022a)

A powerful possibility of using MPC is using constraints, and they can be applied to both inputs and outputs. In combination with the so-called cost function is very interesting in terms of energy-optimal control, as is the subject of the Ph.D. project in the Industrial Control Systems group.

The use of constraints could be when a minimum- and maximum temperature constraint is set for the room air temperature with an ideal setpoint defined. An upper- and lower limit on the supply air temperature is set as constraints. The cost function can be defined as a specific energy consumption calculation that has to be minimized at all times while satisfying the mentioned constraints. MPC can predict the future behavior of the plant outputs and is thus possible to react accordingly to satisfy the constraints.

It results in considering the uncontrollable external input to the system as outdoor temperature, variations in heat load, a temperature in the room between the constraints while minimizing the specific energy consumption. Simultaneously it aims to reach the ideal setpoint if it is the most energy-efficient.

To further improve control- and energy performance, it is possible to consider, e.g., an outside air temperature forecast that will enable the controller to pre-heat or cool the room in advance, potentially achieving further savings on energy.

3

METHODOLOGY & DATA

This chapter will introduce the methodical approach for conducting the work of the thesis and what data has been used.

3.1 Workflow

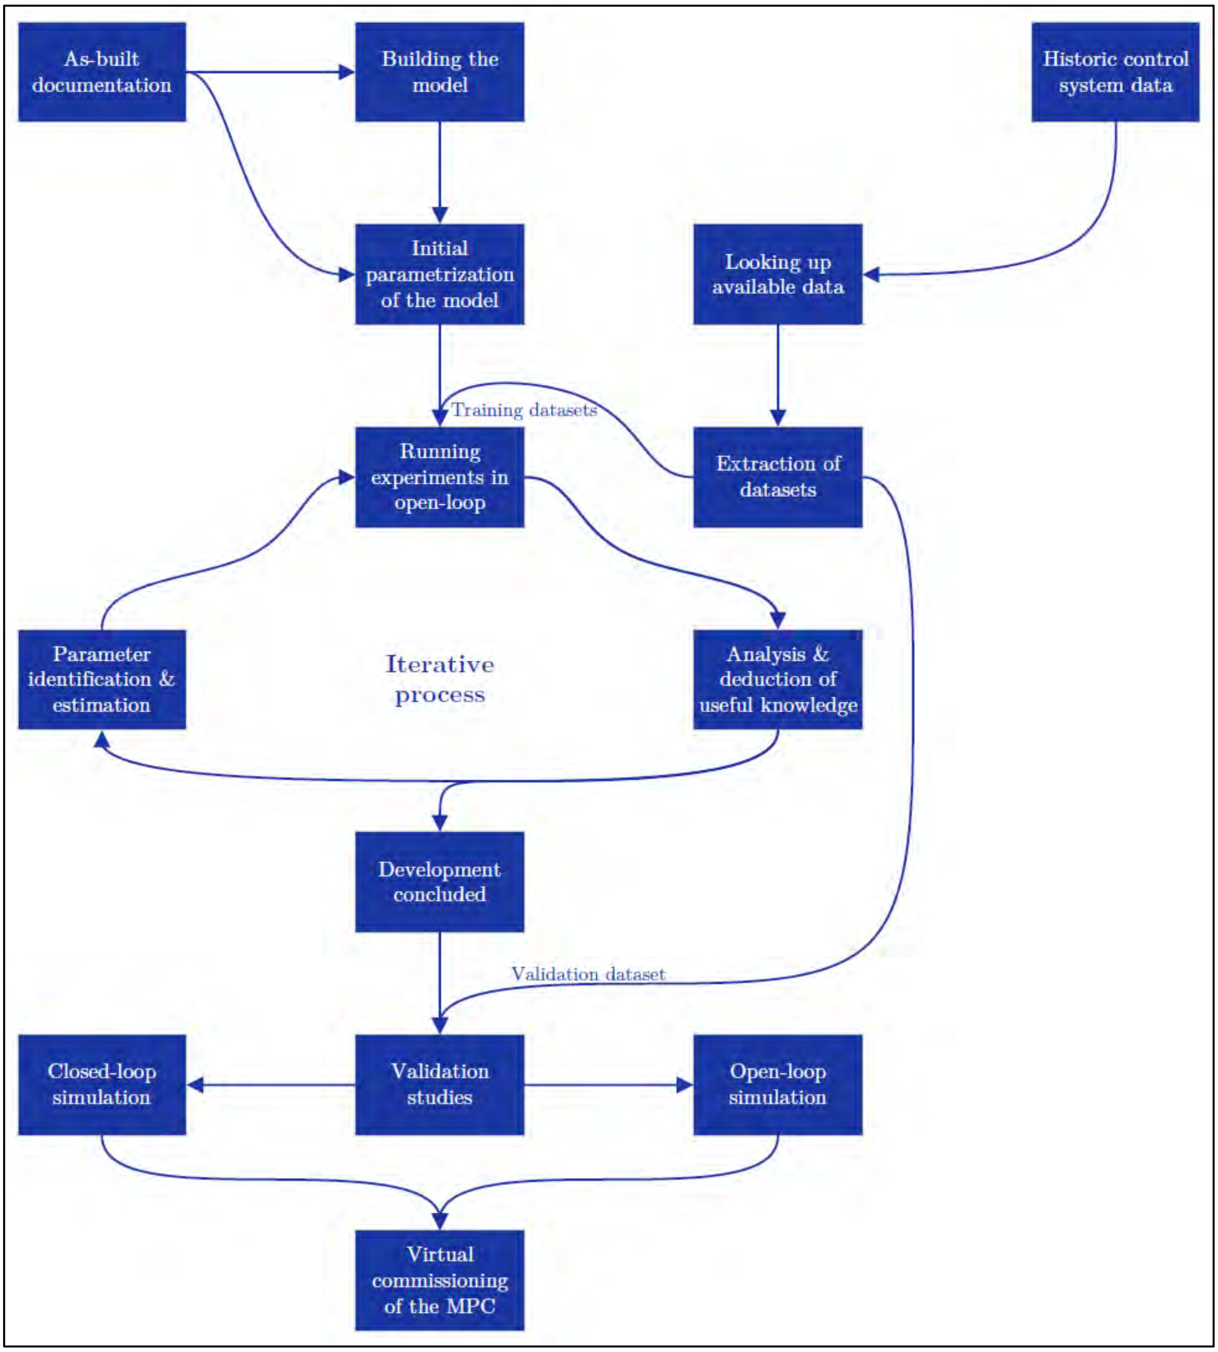

Figure 9: The methodical approach is visualized as a workflow scheme. Note the iterative process, which was a main part of the theses work (Own archive, 2022).

The above figure compactly represents the overall methodical approach and how the data used has been applied. It is presented as a workflow scheme covering the work carried out from the beginning to the end. The aim is to help the reader get a better general understanding of how the work is carried out. It is worth mentioning that the central area of work was within the iterative process seen in the figure.

The following sections of this chapter will go more into detail about the data and methodical approach.

3.2 Data used & methodical approach

The data which has been used can be divided into the following categories, which then will be explained in the following subsections:

- **3.2.1** As-built documentation
- **3.2.2** Historic control system data
- **3.2.3** Running experiments in open-loop
    - With the findings being applied for estimating model parameters

3.2.1 As-built documentation

Based on the investigation findings in **2.2.1**, it was learned that a big emphasis should be laid on mining all available data of the plant and examining it thoroughly, as all data potentially could contribute to valuable information. These data were made available by the Cooling and Ventilation group and consisted of a complete folder containing the as-built documentation.

The documentation consisted of the functional analysis for the plant, which is essential from a process- and control point of view. Furthermore, selection of components, e.g., valves, dampers, filters, sensors, commissioning and test reports, and maintenance guides, among much more.

The data acquisition was the first step, but the material was continuously revisited during the modeling work. That ensured a solid and consistent foundation from which the work took place. As a result, both quantitative and qualitative information can be found. In addition, a lot was learned about the plant from the initial commissioning and operation and maintenance of the plant, including changes/discoveries made.

Some information extracted from this data is definite and was not further considered after implementation in the model parameters. For example, identifying the fan in the air handling unit, the characteristic curve was derived and plotted in a table format in the fan model. Other data, for example, datasheet information on the cooling- and the heating coil, was implemented in their respective models and generally left out of consideration afterward. These are examples of quantitative data that were easily transferable to the model. However, qualitative data in the form of textual descriptions of the plant or operation/control, for example, from the functional analysis, is not transferable in the same way as the data mentioned above but plays an equally important role in the correct understanding and development of the plant model.

The investigation proved that the modeling work requires extensive work to get an in-depth understanding of the respective plant, which is needed. Limitations but also possibilities were encountered during the work, and the knowledge derived from this data provided a foundation of which solutions to encountered problems could be solved.

3.2.2 Historic control system data

During the investigation in **2.2.1**, it was emphasized that this data plays an instrumental role.

Assuming a solid library of components and suitable datasets that cover the whole operating range of the plant and its actuators, a good model can be created using this data. Furthermore, the operational data of the plant can also be used to validate the developed model.

Therefore, the historic control system data is vital for modeling work. First and foremost, both sensor measurements (e.g., temperature measurements) and output/feedback signals to/from actuators (e.g., valves) from the control system were used in the identification of the actual plant and after the initial parametrization of the plant model using the data described in **3.2.1**.

Extraction of datasets

Regarding the LHC construction a home-made CERN framework, UNICOS was developed as a need for the LHC cryogenics control system development. It deals with two layers, the supervision- and control layer. It de facto became the norm at CERN to use UNICOS when developing industrial control systems (UNICOS, 2022a).

The supervision layer of UNICOS uses the new logging system NXCALS based on Hadoop Big data technologies which save data from the systems and store it in CERN's long-term data storage. The Remote Operations Gateway (ROG) made it possible to access all plants employing WinCC OA as SCADA. It is possible to find the relevant sensors and actuators and identify their naming tags. CERN provides an intuitive web-based graphical user interface (GUI) called Timber, where users with the identified naming tag, for example, on temperature measurement, can access the data for a defined period and plot graphs. It was used extensively to identify specific datasets to be used in the parametrization of the model. With more than one year of data, needs were identified to try and parametrize components in the model. For that purpose, it was ideal to look up and search for specific operational scenarios in the data which could be extracted and then used. These can be referred to as training datasets.

Not all operational scenarios can necessarily be expected in the plant operational data, e.g., the heating- and cooling coil was discovered to only have recorded data of them operating roughly in the range between 0-5 % opening. Therefore, many different datasets were used throughout the year. However, as it is difficult to have one set of model parameters to work consistently over a full year, it is not unnormal to have different sets of parameters, e.g., summer and winter. Therefore, as argued in **1.4**, the focus was narrowed down to find a single set of parameters that would fit a few consecutive months consistent with the warmer months and not a full year.

It is possible to query the data directly from Timber, but other measures were used because it was easier to perform repetitive extractions and is computationally faster. In addition, CERN has an exclusive Jupyter-based notebook application called SWAN available for its users. The notebook can be used with several programming languages, most notably the Python language used broadly. SWAN aims to enable its users to query data from the system and provide a platform to perform interactive data analysis in the cloud with SPARK cluster computing power available (NXCALS, 2022a).

For the project, the SWAN notebook was mainly used to extract datasets and modify them to fit its purpose. Because of how data are logged into the system, it was needed to resample the data. Data are not continuously saved at a specific time delta but instead saved at a maximum time delta observed to be more than 24 hours if no changes are recorded. With changes recorded, plots have been saved with less than 1-second intervals, depending on the rate of change. Therefore, the data had to be synchronized with the beforementioned Python script using the SWAN notebook. The datasets, which consist of columns and rows, were saved in .csv files and converted to .txt files, where these could be loaded directly into EcosimPro.

3.2.3 Running experiments in open-loop

The building of the model is presented in **4.2**, and with the model being able to compile, generate the partition (the mathematical model), and verify that simulations were able to run successfully, experiments could take place. The simulations were run as open-loop, meaning the approach was to use the extracted data explained in the former subsection. For example, consider 24 hours with operational data available for the period. The control signals sent to the actuators of the actual plant are also used to control the models' actuators. The fresh air temperature for the same period are also used as input to the model.

Control signals to the actuators are used with few unimportant exceptions as well as wherever possible, input measurements from the actual plant are applied to the model. A realistic and practical comparison platform is present when the model is running, and this data is applied. It is possible to compare simulated values to measured values, whereas the observed differences can be considered the model parameters. The assumption is that a good model of each component is the case with relevant parameters available to adjust.

As the model has limitations, there are some differences that the model could not consider. However, the main differences were the model's parameters, which could be estimated against the operational data.

As mentioned in **3.2.1**, some parameters were relatively easily extracted and applied. For example, even though some minor problems and discrepancies were considered, it was straightforward to parametrize the fan model and apply the correct scaling to reproduce an airflow consistent with actual plant measurements. Other cases were a bit more problematic, but it was essential to use the operational data actively to estimate model parameters.

Certain aspects are essential to consider to make this approach practical. First, a good understanding of the processes and thermodynamics that make up, in this case, an HVAC plant is imminent. To mimic the plant's behavior, some dynamics needed to be investigated to deduce a reasonable explanation and try to implement the findings into the model.

For example, the supply air temperature, which is an important one in the context of the modeling work, required to thoroughly examine a period where some deviations in dynamic behavior were present, which needed to be sorted out. It was complex as several variables affect one variable (supply air temperature), but manipulating some of the variables will also affect each other internally. A careful approach was taken, and findings from experimenting with the model parameters were learned to understand the model itself better.

It enabled to point out possibilities and limitations in the model from the plant modeling perspective. However, this approach required an excellent and efficient monitor for plotting the simulation results. The monitor is presented in **4.3.1**.

3.3 Validation studies

With the development of the model concluded, the purpose of these studies was to look at the model state and examine how it performs regarding the actual plant to evaluate the applicability of the developed model.

The validation study was performed on the model in two steps:

- Open-loop simulation
- Closed-loop simulation

The open-loop simulation directly compares the model against the actual plant measurements. Statistical measures will be used to quantify the precision of the output (temperature variables), as this kind of simulation will give the most unambiguous indication of the model's ability to reproduce actual plant behavior. The dynamic behavior will not be evaluated quantifiably but will be considered when evaluating the simulation results and discussed. However, good results on the precision will also indirectly indicate that the dynamic performance of the model is good.

The closed-loop simulation is the next step in the validation process. It differs from the open-loop simulation and can only, to a lesser degree, be compared directly to measurements. However, a strong point of model validation is if the current control system configuration can run the model. Closed-loop will, like open-loop simulation, also differ. Although the model is simulated in a closed-loop, good results can necessarily be expected regarding the actual regulation. If that is the case, the differences will likely be in the control signal inputs, which is the most relevant measure to evaluate.

3.3.1 Statistical measures

The following two statistical measures were applied to the four temperature variables to quantify the precision of the model.

Root Mean Squared Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{\sum_{i=0}^{n}[y_i - y_i]^2}{n}}\ [°C] \tag{1}$$

Mean Absolute Percentage Error (MAPE):

$$\text{MAPE} = \sum_{i=0}^{n} \frac{|y_i - y_i|}{y_i} \times 100\ [\%] \tag{2}$$

Where $y_1, y_2, \dots, y_n$ are historic control system values, $y_1, y_2, \dots, y_n$ are predicted values by the model and n is the sample size. The unit of RMSE is the same unit as y, e.g., $[°C]$ and MAPE are in %.

Both measures are commonly used in the literature when comparing a model to measurements and are relatively easy to apply and understand.

The RMSE can be interpreted directly in the unit of the variable itself, in this case $°C$ and is negatively-oriented. An RMSE value of 0 $°C$ will represent a perfect model fit. As the error between the two variables is squared before taking the root, significant prediction errors will have a relatively big negative effect on the RMSE value, which can be an advantage if large prediction errors are unwanted. The value will also always be positive because the error is squared.

The MAPE is slightly different from RMSE and measures how much the simulated variable, on average, is different from the measurements. Likewise, the RMSE is negatively oriented. As the value is represented in percent, a better relative measure of how large the model's error is can be interpreted. Both measures were chosen because they are thought to give a reasonable and realistic picture of the model's prediction abilities. A single value is calculated for each measure on the respective variables compared.

4

Model development

This chapter focuses on the actual work that has been executed in the modeling of the SR1 plant. As learned in 2.2.1, in-depth plant knowledge is needed to consider the complexity. Therefore a tremendous amount of effort has been put into learning about the SR1 HVAC plant.

In the first section, the SR1- facility and plant are introduced. Then, section 4.2 explains the approach to building the SR1 plant in EcosimPro. Finally, in section 4.3, a selection of relevant specific details in context to more concrete areas of the modeling work are highlighted and explained.

4.1 Introduction to the pilot plant of SR1

The continuation of the Ph.D. project has turned its focus toward HVAC plants - more specifically, the air handling unit on them.

Previously work on the cooling towers present at CERN has been investigated. This exact plant was initially chosen for the Ph.D. project alone based on its suitability for the MPC development work. The AHUs of the SR1 plant had an exemplary configuration, including a variable frequency-driven fan and rich instrumentation.

The aim is ultimately to deploy an MPC on a pilot plant. However, there is a dependency on good cooperation with the Cooling and Ventilation group, including a green light from the management before deployment.

The SR1 plant is less critical than other plants, further suiting the plant as the pilot for the project.

4.1.1 SR1 facility

SR1 is a surface facility placed in building 2175 at LHC point 1, where the ATLAS experiment is also placed. Similar buildings are placed at the other points (8 in total) named 2275, SR2, 2375, and SR3. The building is containing IT-equipment, namely electronic racks and some power converters, which essentially is the machinery that ejects heat, which has to be removed. The equipment is for test purposes as the “real” machinery is found underground.

The underground facility is of great criticality in achieving the objective of the experiments, essentially getting physics data from them. Therefore, the basis of the existence of the

SR1 facility is solely for test purposes. At the same time, it also means that the HVAC system of SR1 is less critical.

The building consists of the main hall where the inside has placed a selection of rooms, including a so-called clean room. That room has an HVAC plant, the SR1 'Salle Blanche' plant (not related to the SR1 plant), which maintains a positive pressure to ensure no outside contamination as sensitive detector parts are assembled there. This selection of rooms placed in the main hall is not at the full height of the main hall itself. Thus, there are several meters above these rooms, and it means at the same time that the two outer ends of the main hall are connected over the top of them. Therefore, it can be seen in **Figure 10** that the size, shape, and heat ejection equipment of each room, respectively to the left and right, is not similar.

The SR1 plant consists of two AHUs placed in two separate outside rooms at each end of the main hall, close to the equipment it effectively needs to cool down (or heat). Supply air from each AHU is blown into a distribution plenum through a false floor. This false floor is an approx. 2-meter-tall basement consisting of various installations and pipings, and both AHUs are also connected through this. In addition, fixed position dampers are placed along a corridor in the false floor. These are closed towards the center and opened in the ends helping to distribute air.

In **Figure 10** is seen:

1 AHUs (UARX10001 in R-401 and UARX10002 in R-403)
2 Electronic racks
3 Power converters
4 The clean room of SR1 (marked by red lines)

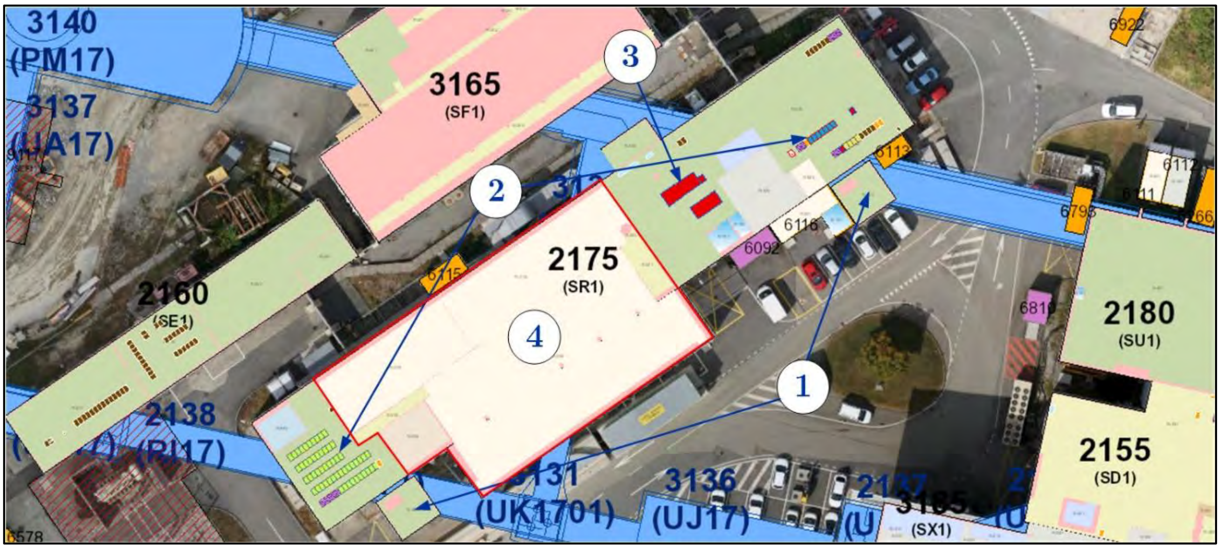

Figure 10: Air photo of the SR1 facility of B.2175 with the IT-Equipment layer at room level available (GIS Portal, 2022a).

In **Figure 11**, a schematic is seen of the complete plant of SR1 taken from the Process and Instrumentation Diagram (P&ID). Two AHUs feed supply air into the distribution plenum, the false floor, from which grilles are strategically placed on the floor at the bottom of the electronic racks from which airflow will stream through and provide cooling (**Figure 12**). In addition, some racks have cooling internally provided by a chilled water circuit that is not part of the SR1 plant.

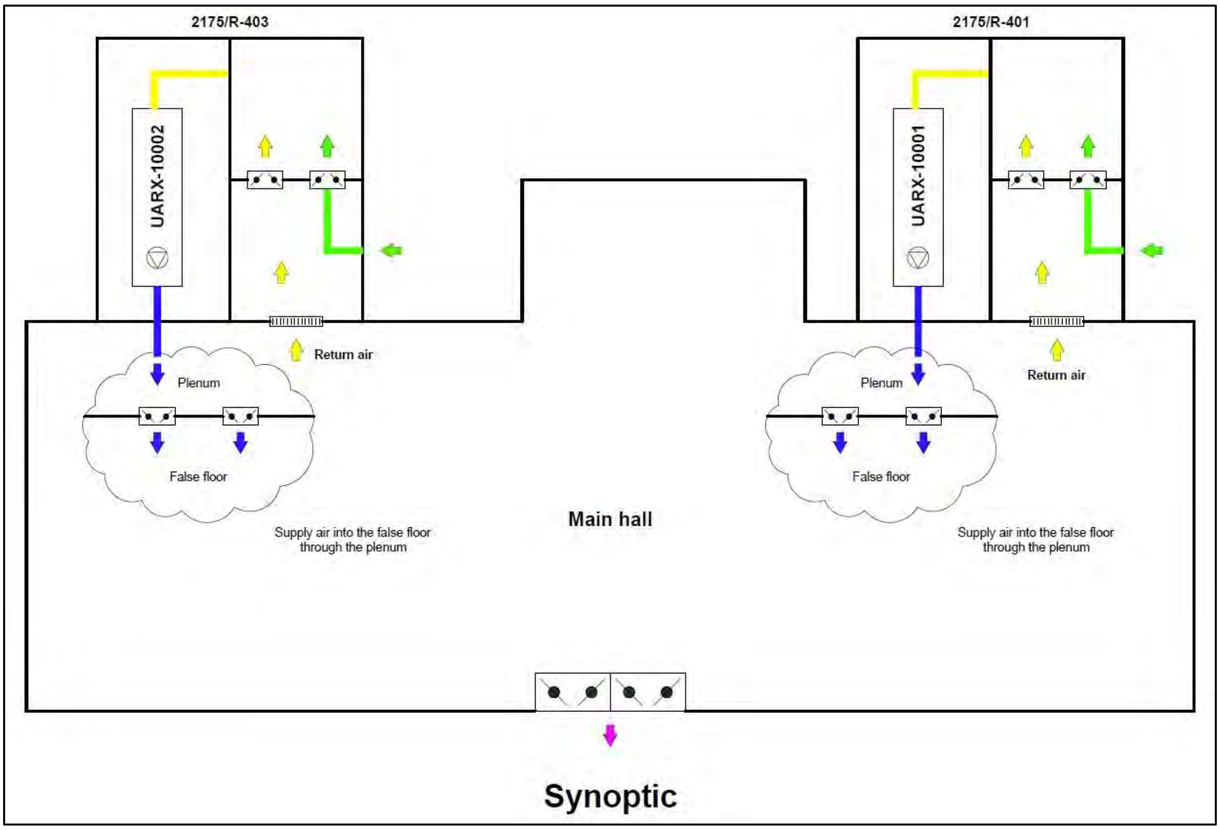

Figure 11: A snippet of the P&ID showing a schematic of the complete SR1 HVAC plant (Cooling & Ventilation Group, 2021a).

Figure 12: A rack where the grille can be seen at the bottom. Through the grille, it is possible to see the false floor and some of the installations also placed there (Own archive, 2022)

4.1.2 The air handling units of the SR1 plant

The complete HVAC plant system consists of two identical subsystems. These subsystems are an air handling unit, including a filter, a heating- and cooling coil supplied with chilled/hot water, a fan, and various instrumentation (Figure 13 & Figure 14). The chilled water is produced on-site by a sizeable chilled water plant that supplies different clients, including the SR1 plant, most notably the cryogenic systems, which have large heat rejection requirements. The hot water is produced by a shared hot-water production facility at the Meyrin site, distributing hot water to clients in different places, including the SR1 plant.

Each AHU can operate in the following modes:

- Free cooling mode
- Active cooling mode
- Active heating mode

In free cooling mode, only the actuation of dampers and possibly the fan affects the measured value, the main hall air temperature. In either active- cooling or heating mode, a control signal is requested and applied to the control valve and allowing a flow of either chilled or hot water, providing cooling or heating.

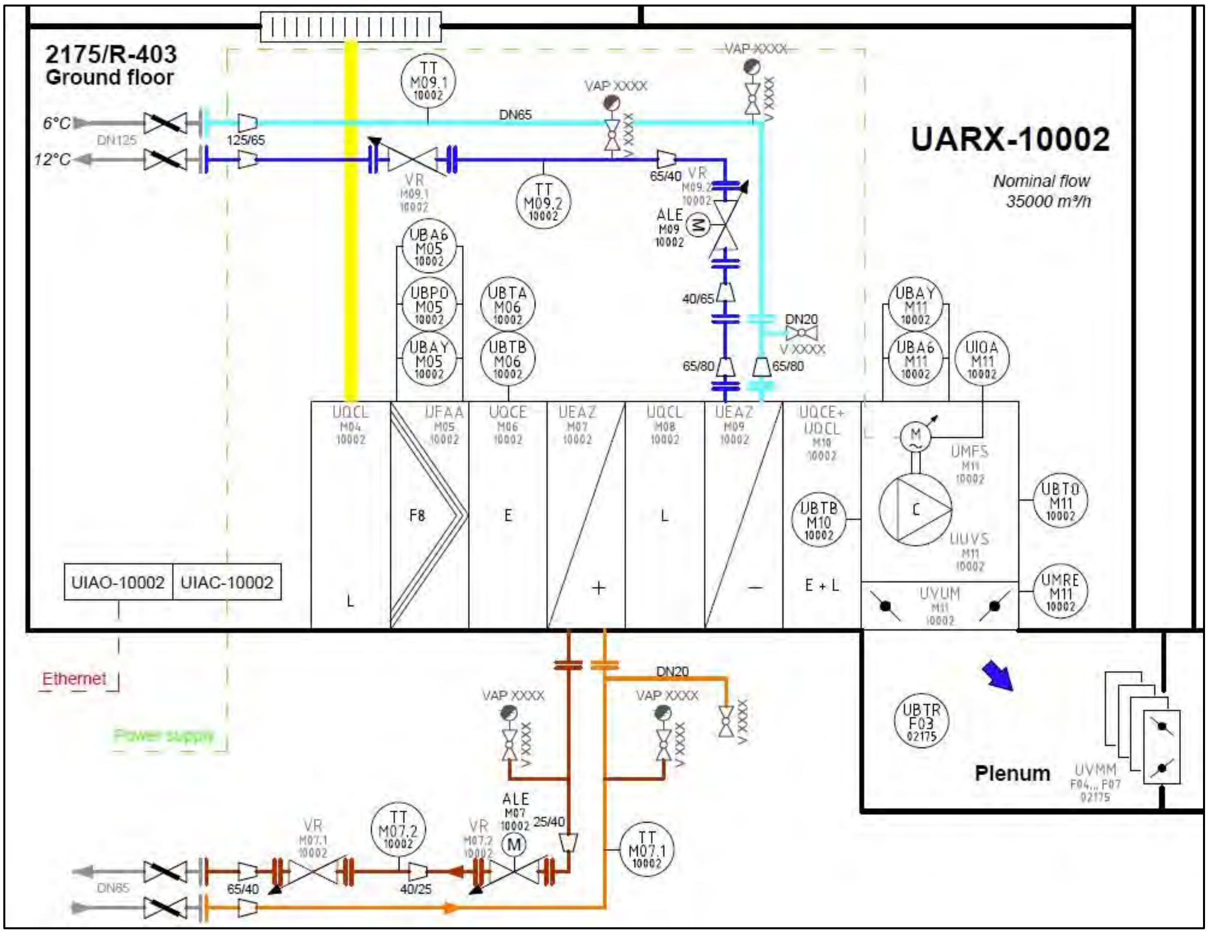

Figure 13: A snippet of the P&ID where the AHU is seen (Cooling & Ventilation Group, 2021a).

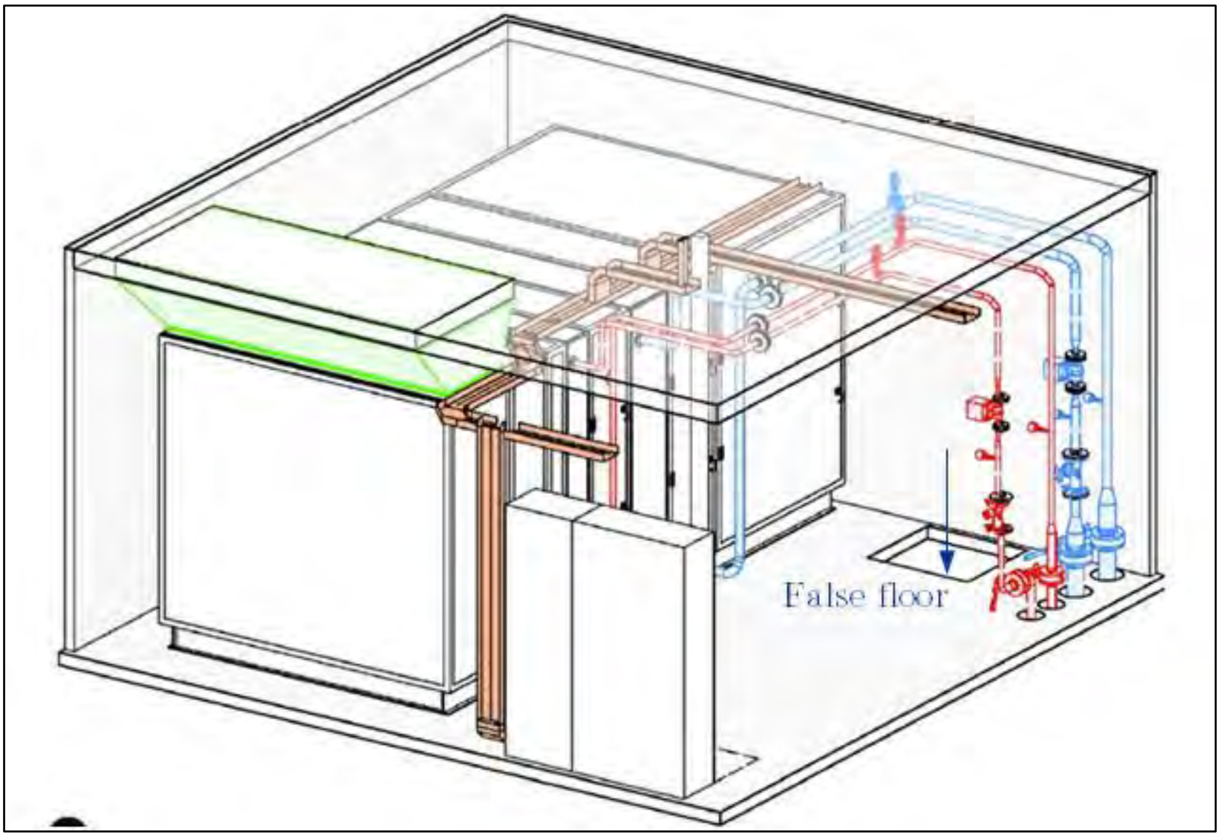

Figure 14: 3D visualization of the air handling unit, where the operator's entrance to the false floor can be seen (Cooling & Ventilation Group, 2021a).

Each sub-system is over-sized as a single AHU operating at maximum capacity can take the entire heat load if the other is out of operation. However, the nominal operating conditions are well under the maximum capability of each AHU. As explained in the former subsection, this is doable due to the distribution plenum common to both sub-systems.

Both AHUs have a PLC that communicates through a Profibus connection. In addition, they share a few operating parameters that ensure the continuation of the process operation if one of the sub-systems stops operating.

The air drawn into the AHU is a mix of fresh- and recycled air from the main hall, respectively, through an outside air intake and an extraction grille (**Figure 15**). A fresh- and recycled air damper controls airflow into a small room, which can be considered a mixing plenum from where the air mix is drawn into the AHU (**Figure 16** & **Figure 17**). At the inlet of fresh air, there is ducting to the fresh air damper. However, just before it is, a fixed position damper installed and adjusted to allow a minimum of $1'000\ m^3/h$ of fresh air into the process regardless of the two other dampers. This portion of air is mixed with the recycled air before entering the recycled air damper.

As the system can be a scenario where primarily fresh air is drawn into the process, and at the same time no active extracting of air is present, two exhaust air dampers (**Figure 18**) prevent a too big positive pressure inside the building. The two dampers are each connected to one of the two AHUs. They are simply operating at a fully open/closed position determined by the transition point, which is 20 % opening for the fresh air damper respectively of each of the AHUs.

Figure 15: One of the recycled air extraction grilles places in the main hall (Own archive, 2022)

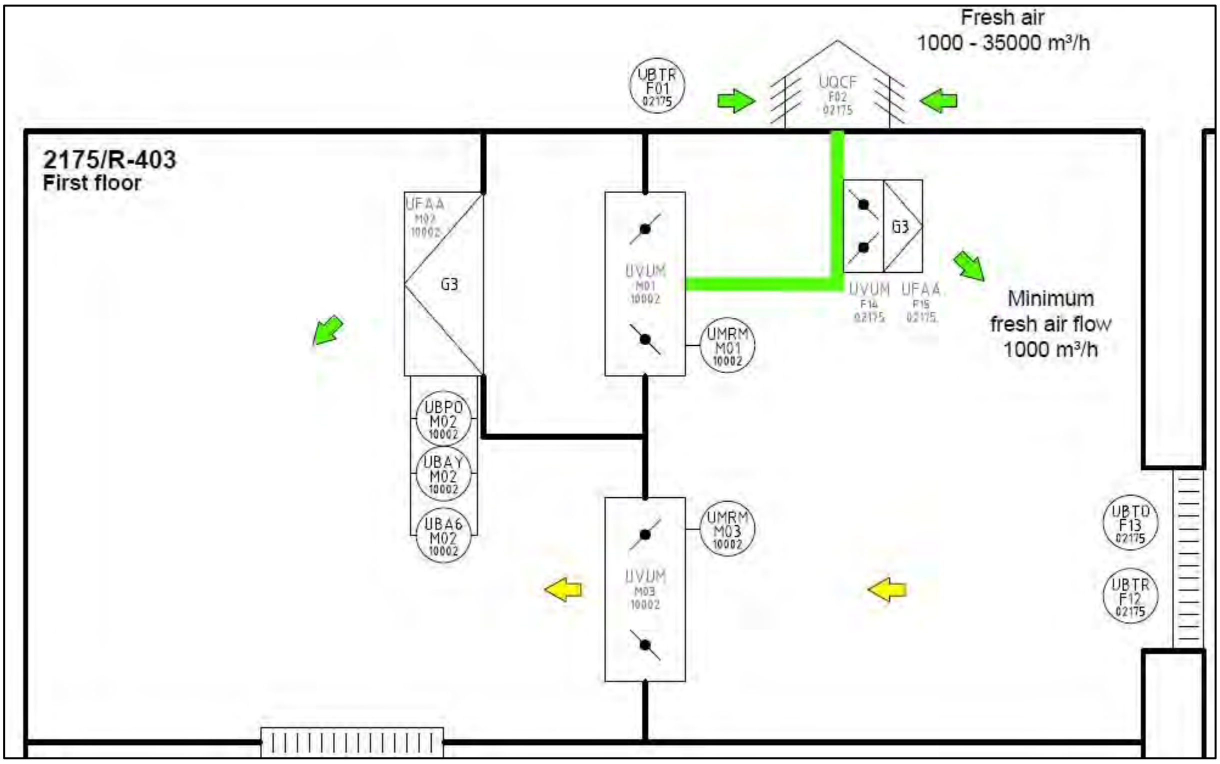

Figure 16: A snippet of the P&ID showing the entrance of fresh- and recycled air through a damper for each air-stream and into the mixing plenum (Cooling & Ventilation Group, 2021a).

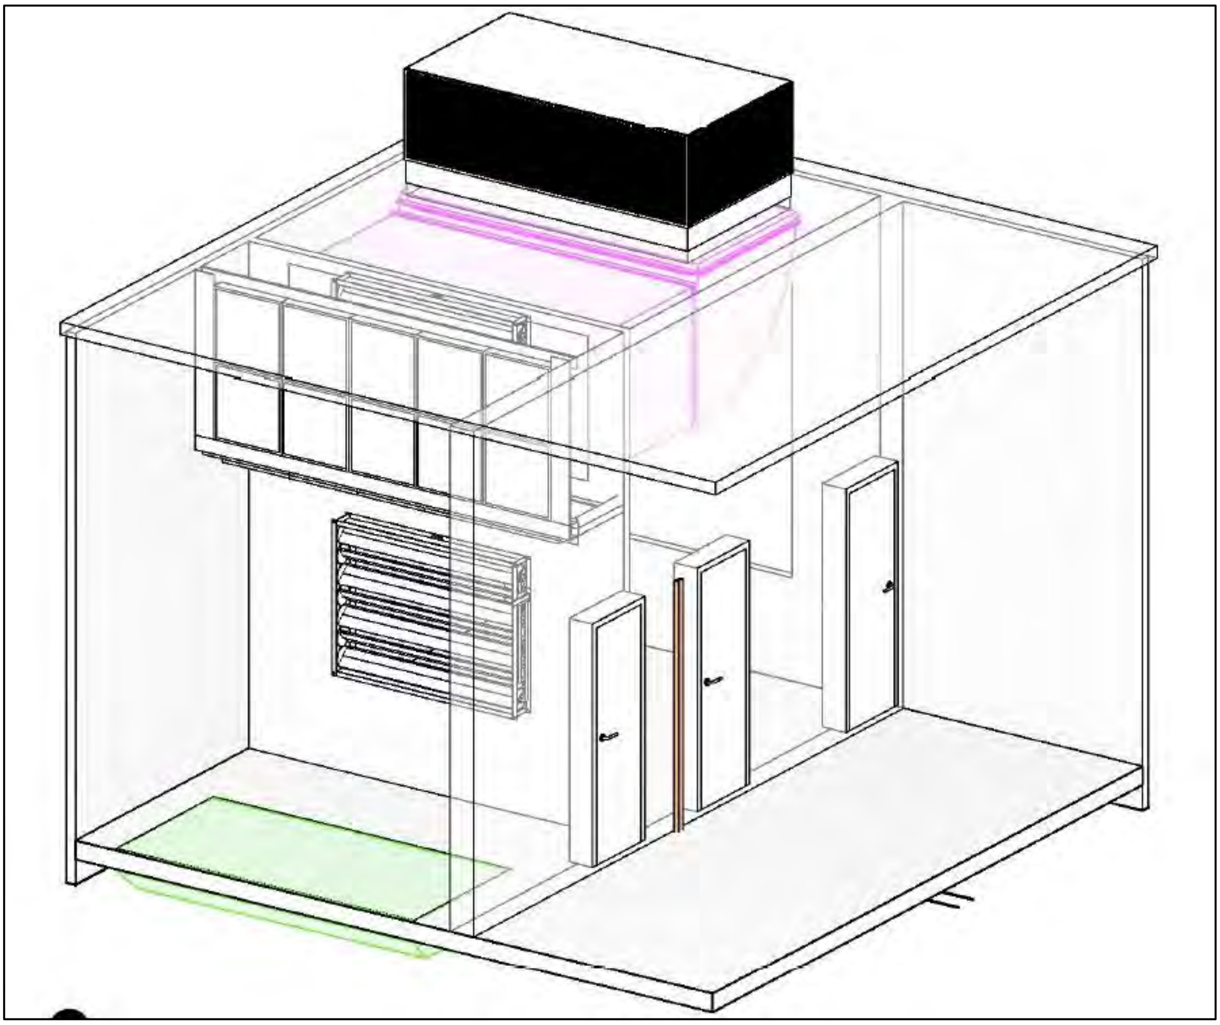

Figure 17: 3D model showing the R-401 first floor. The entrance of fresh- and recycled air through the dampers into the mixing plenum can be seen (Cooling & Ventilation Group, 2021a).

Figure 18: Both exhaust air dampers can be seen and are placed centermost and at the top of the main hall. The clean room can be seen directly underneath the dampers (Own archive, 2022).

Like the rooms, their sizes and equipment content are not the same; a slight variation in the operation of each AHU is observed.

Things like the heat load unequally admitted into the hall, the distribution of air, and the placements of various sensors will influence.

4.2 Building the model

The schematic in EcosimPro is a visual representation of the system and the worksheet where the modeling work takes place. With the different libraries available, including the HVAC library of CERN, components can be drag-and-dropped in the schematic. Lines connect components to in- and out ports. Each component consists of a mathematical model, and by coupling components, a complete mathematical model of the whole system is created (by generating a partition). There are different ways of modeling a system and driving the model. For example, the model can be provided with mass/volume airflow sources, temperature, pressure, relative humidity, and a sink in a completely open-loop simulation. It depends on the actual case and what is relevant. As explained in **3.2.3**, the method used is an open-loop model to run simulations. The model essentially is locked in with as much historic control system data to feed into the model as possible. The model can then produce simulated values that can be directly compared to operational data.

The chosen approach to produce airflows in EcosimPro is based on pressure drops in the system network. As there is a relation between Δp and airflow in a ventilation system and different components causing a pressure drop, an approach where Δp's was calculated provided a reference pressure drop at nominal airflow conditions was chosen as the most suitable. With information available, it was used to establish a reasonable starting point. The fan model component is the most essential to determine the airflows in the system. However, the pressure drops around the system network caused by ducting, other components, etc., are also essential to mimic the existing system and its dynamics.

The system is relatively large, and for other beneficial reasons, the complete SR1 plant was built in three layers with a schematic each. These three layers are compiled into a component and the middle and bottom layers into a symbol. The symbols can be placed in other schematics and interconnect the components. For example, the bottom layer consists of the two AHUs from the air inlet to the air supply (**Figure 19**). An in- and out air source port was placed in the schematic to port the air source from the middle layer (**Figure 20**). The top layer only consists of the different sources from the historic control system data, fixed source values, and other components such as scaling (**Figure 21**).

The advantage of modeling in several layers is that only the uppermost layer needs to generate a partition. The simulations can be run from this layer only, noting that all layers are interconnected. Each component needs to be compiled before generating the complete model's partition (the mathematical model). Given that a model is completed, the bottom- and middle layers can be left, and copies of the top layer can be provided with new source data files for running multiple experiments with different datasets.

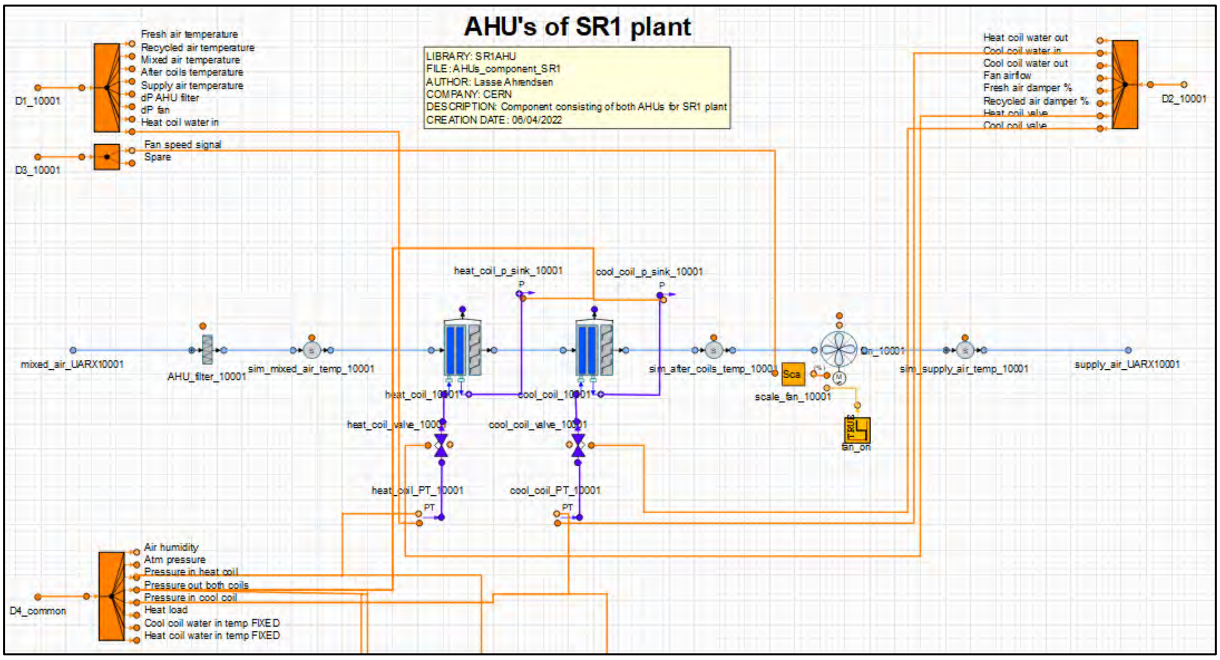

Figure 19: A snippet of the schematic of both AHUs modeled. Only one AHU is seen in the figure, and the source files are depicted (orange). On the left- and righthand side, the humid_air port can be seen from where the air source is ported to and from the model's middle layer (**Figure 20**) component (Own archive, 2022).

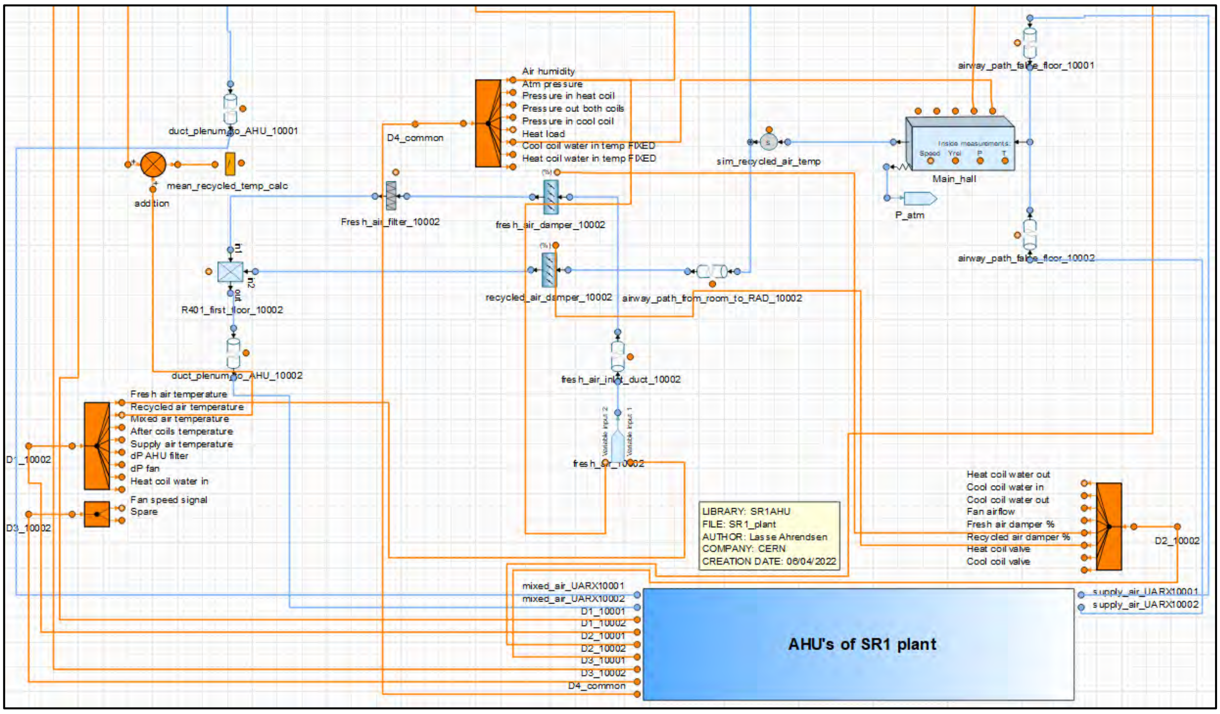

Figure 20: A snippet of the middle layer schematic of the plant model. Excluding the AHUs, as shown in **Figure 19** above. The sources are depicted (orange) (Own archive, 2022).

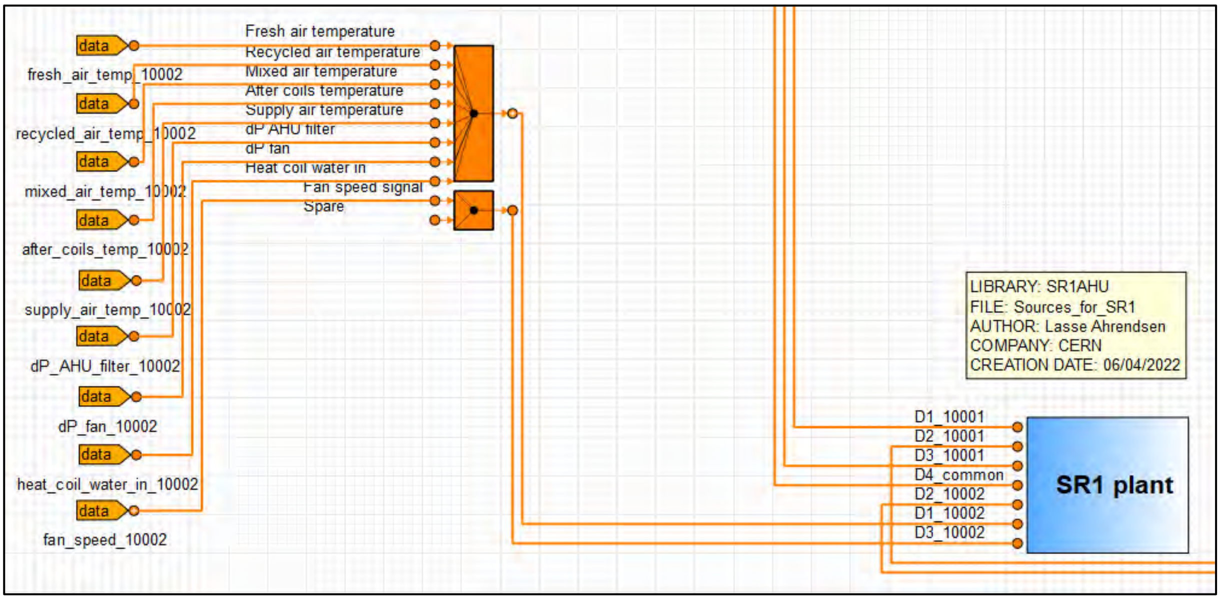

Figure 21: A snippet of the upper layer schematic where all source files are placed and ported into the component of the SR1 plant, which is the middle layer, as seen in **Figure 20** (Own archive, 2022).

4.3 Parameter identification & estimation

As emphasized, all data available in the as-built documentation, discussions with co-workers, good cooperation with EN-CV's process experts, and a site visit were instrumental in finding suitable parameters that will make the model behave like the actual plant.

Various selections of concrete cases in the overall modeling work process are presented in the following subsections. A great deal of work was put into analyzing the data, including the simulation variables. In addition, it was extensively used to estimate and correct parameters in specific scenarios drawing on the in-depth knowledge needed to model the SR1 plant.

4.3.1 Setting up the simulation monitor

As mentioned in **3.2.3**, it is crucial to have a good monitor for plotting simulation results and investigating and learning about the plant behavior, both the actual plant and the model. It enables to deduce knowledge that can be applied in the modeling work to estimate parameters. All relevant sensor information from the plant is extracted in the different datasets.

Some of them are used to drive the model by connecting them to their input control ports (e.g., control signal to cooling coil valve), and some are merely used to compare real measured values to simulated values. For example, in **Figure 22** one of the tabs contains plots (seven tabs in total).

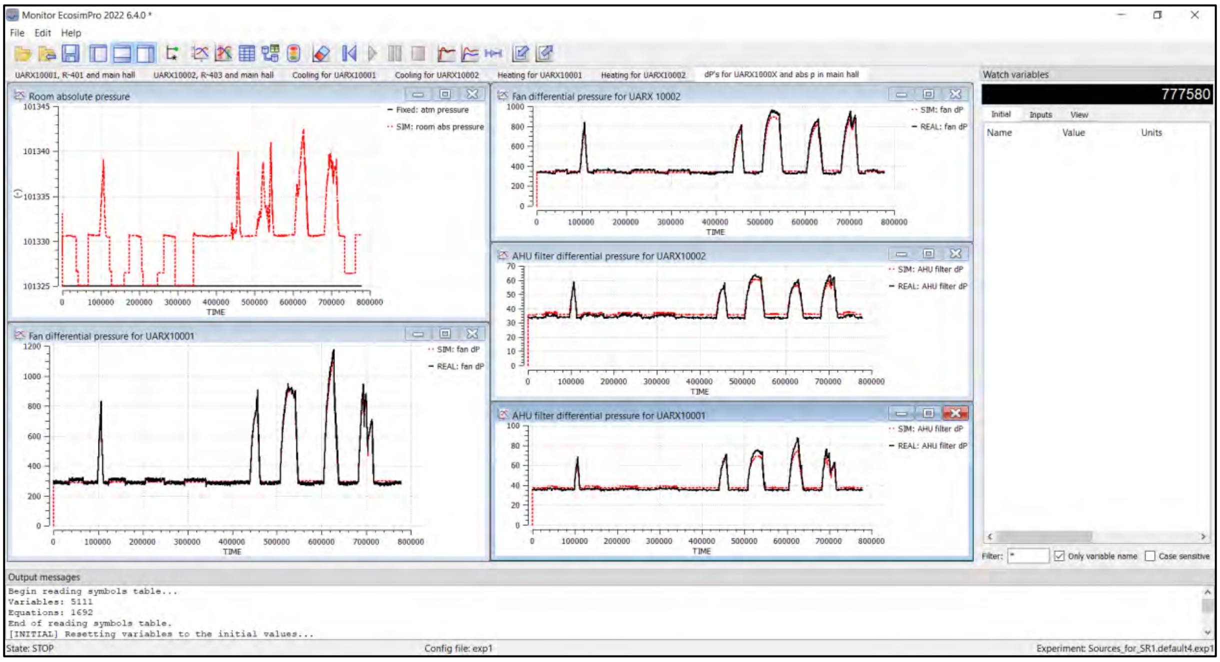

Figure 22: Screenshot of the tab: dP's for UARX1000X and abs p in the main hall. An explanation of the containment of plots is found in the text following the figure (Own archive, 2022).

A significant amount of effort was put into building a good, manageable, and efficient monitor that would help and ease the investigation process, which was needed to examine many different and specific operational scenarios.

The different tabs are categorized as follows:

▪ **UARX10001, R-401, and main hall** ▪ **UARX10002, R-403, and main hall** o Plots containing: - Air temperature: fresh air, recycled air, mixed air, after coils, supply air - Airflow: fresh air, recycled air, mixed air/fan air - Control signal: fresh- and recycled air damper, fan speed
▪ **Cooling for UARX10001** ▪ **Cooling for UARX10002** o Plots containing: - Cooling effect: cooling coil - Control signal: cooling coil valve - Water flow: cooling water - Air temperature: before coils/mixed air, after coils - Water temperature: cooling coil in, cooling coil out
▪ **Heating for UARX10001** ▪ **Heating for UARX10002** o Plots containing: - Heating effect: heating coil - Control signal: heating coil valve - Water flow: hot water - Air temperature: before coils/mixed air, after coils - Water temperature: heating coil in, heating coil out
▪ **dP's for UARX1000X and abs p in main hall** o Plots containing: - Absolute pressure: room (main hall) - Differential pressure: fan, AHU filter

Table 1: Overview of tabs and plots within the tabs in the simulation monitor which was an essential tool in the modelling work of parametrization (Own archive, 2022).

There are both simulated- and measured values present for many of the variables. It is not the case for all variables in the whole system, but even if no comparison can be made directly, it is still helpful to see what output the simulation provides as it can still be helpful in the work. For example, there is only one flow measurement in the system, but it could be interesting to know the airflow distribution between the fresh- and recycled air, but it is not directly possible to know. A good understanding of the process will still provide helpful insight into what happens inside the model, given the possibility of getting a substantial number of simulated variables out of the model to be displayed in the plots.

The tabs were categorized into suitable tabs containing plots relevant to each other. The tab: **UARX10001, R-401, and main hall** contained plots where the recycled air temperature, both simulated and measured, could be compared and the same for the supply air temperature. Output signals to the actuators, namely the dampers and the fan, are displayed as the system's airflow. This tab was worthwhile to investigate the airflow in the system and how it changed with different damper positions and fan speed. Also, in the parametrization of the damper authorities were, this tab used. It was rather complicated as it was impossible to parameterize one of them at a time, given their coupled behavior. Several plots with variables affecting each other were needed to account for all at once.

4.3.2 Fan, airflows & pressure drops in the system network

A natural starting point was to parametrize the fan model. The fan handles the airflow in the system, but the whole system network has to be taken into account as it has an influence.

Information about the network and the fan itself, to a certain extent, could be found in the available data. For example, for the fan component in EcosimPro, the characteristic fan curve (**Figure 23**) can be plotted among other parameters and initial conditions (**Figure 24**).

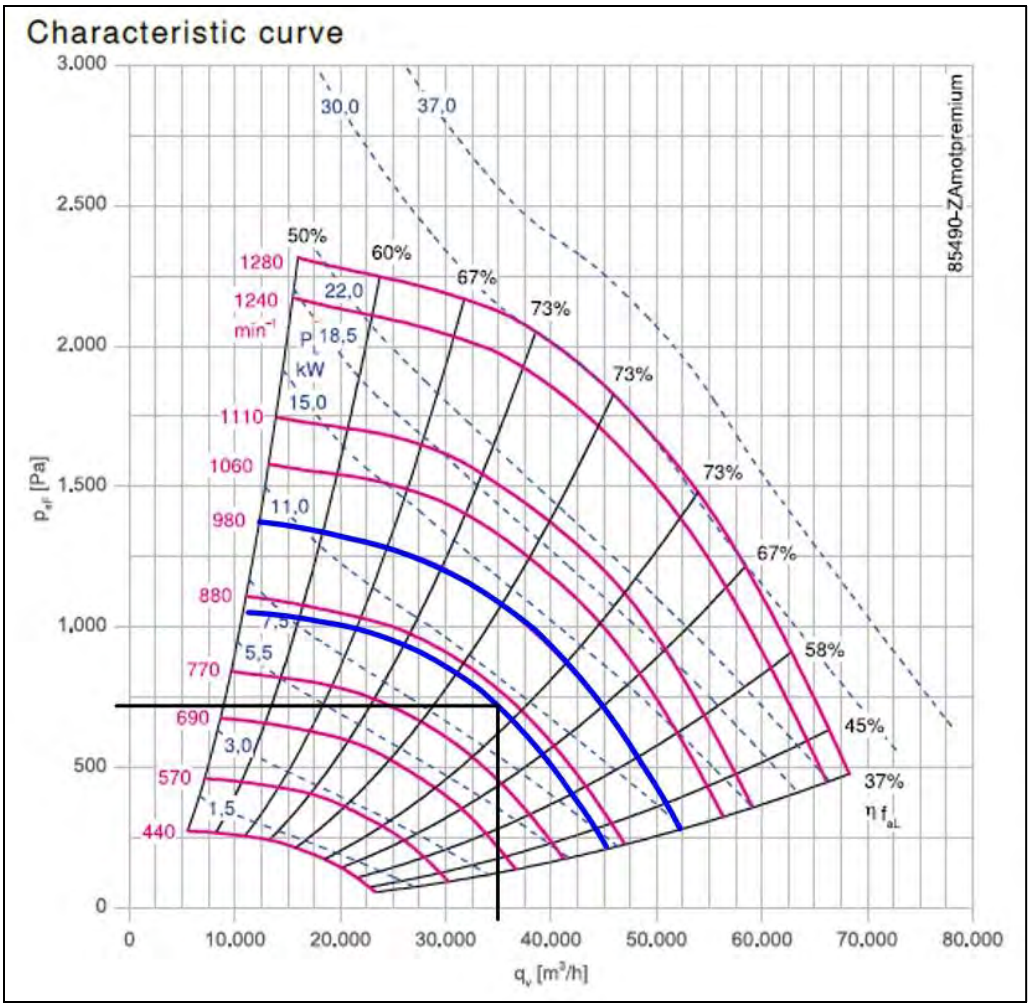

Figure 23: Fan characteristic curve found from the datasheet. The upper blue line is the one that is plotted in the fan component and is for a maximum speed scenario of the fan. The lower blue line and the point where the intersections meet are from actual measurements performed during the commissioning phase of the plant in nominal conditions (35′000 m^3/h) which were used as a reference (Cooling & Ventilation Group, 2021a).

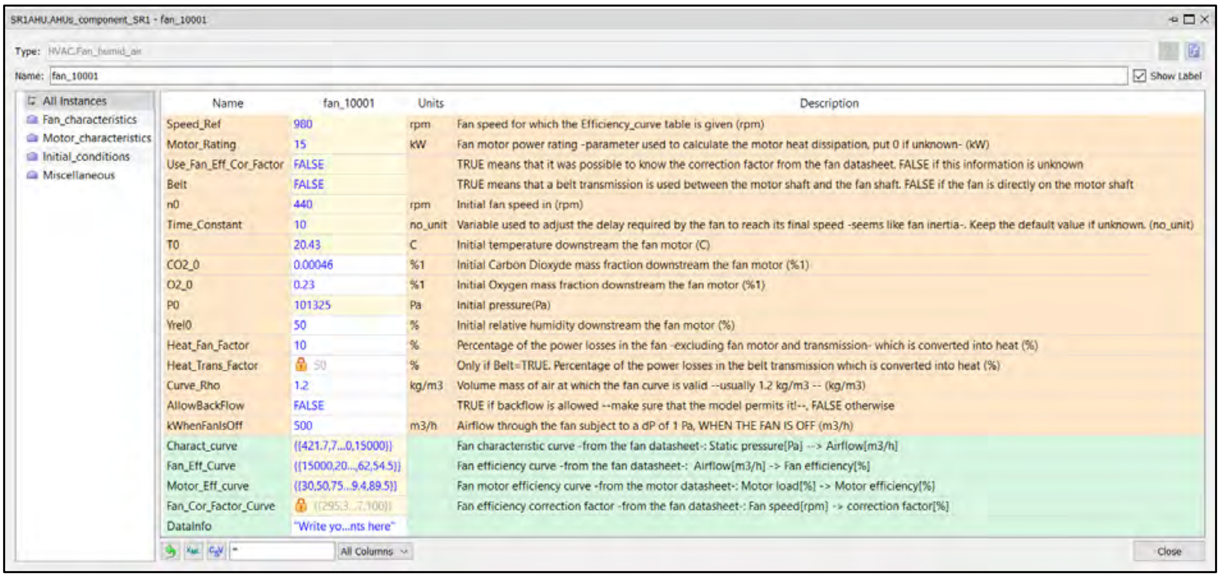
SR1AHU.AHUs_component_SR1 - fan_10001

Type: HVAC.Fan_humid_air

Name: fan_10001

Show Label

Name	fan_10001	Units	Description
Speed_Ref	980	rpm	Fan speed for which the Efficiency_curve table is given (rpm)
Motor_Rating	15	kW	Fan motor power rating -parameter used to calculate the motor heat dissipation, put 0 if unknown- (kW)
Use_Fan_Eff_Cor_Factor	FALSE		TRUE means that it was possible to know the correction factor from the fan datasheet. FALSE if this information is unknown
Belt	FALSE		TRUE means that a belt transmission is used between the motor shaft and the fan shaft. FALSE if the fan is directly on the motor shaft
n0	440	rpm	Initial fan speed in (rpm)
Time_Constant	10	no_unit	Variable used to adjust the delay required by the fan to reach its final speed -seems like fan inertia-. Keep the default value if unknown. (no_unit)
T0	20.43	C	Initial temperature downstream the fan motor (C)
CO2_0	0.00046	%1	Initial Carbon Dioxyde mass fraction downstream the fan motor (%1)
O2_0	0.23	%1	Initial Oxygen mass fraction downstream the fan motor (%1)
P0	101325	Pa	Initial pressure(Pa)
Yrel0	50	%	Initial relative humidity downstream the fan motor (%)
Heat_Fan_Factor	10	%	Percentage of the power losses in the fan -excluding fan motor and transmission- which is converted into heat (%)
Heat_Trans_Factor	50	%	Only if Belt=TRUE. Percentage of the power losses in the belt transmission which is converted into heat (%)
Curve_Rho	1.2	kg/m3	Volume mass of air at which the fan curve is valid --usually 1.2 kg/m3 -- (kg/m3)
AllowBackFlow	FALSE		TRUE if backflow is allowed --make sure that the model permits it!--, FALSE otherwise
kWhenFanIsOff	500	m3/h	Airflow through the fan subject to a dP of 1 Pa, WHEN THE FAN IS OFF (m3/h)
Charact_curve	({421.7,7...0,15000})		Fan characteristic curve -from the fan datasheet-: Static pressure[Pa] --> Airflow[m3/h]
Fan_Eff_Curve	({15000,20...,62,54.5})		Fan efficiency curve -from the fan datasheet-: Airflow[m3/h] -> Fan efficiency[%]
Motor_Eff_curve	({30,50,75...9.4,89.5})		Fan motor efficiency curve -from the motor datasheet-: Motor load[%] -> Motor efficiency[%]
Fan_Cor_Factor_Curve	({295.3...7,100})		Fan efficiency correction factor -from the fan datasheet-: Fan speed[rpm] -> correction factor[%]
DataInfo	"Write yo...nts here"		

Figure 24: Attributes of the fan component with parameters that can be changed. Note the row `Charact_curve` where the fan curve is to be plotted (Own archive, 2022).

The fan component needed to have two input sources, one for the fan speed signal and one on/off to set the running status of the fan. These are provided in the bottom layer, ported from the top layer, as seen in **Figure 19**. The fan speed signal from the historic control system data needs to be applied with scaling as the minimum fan speed (1 %) is set directly at the variable frequency drive supplying the electric motor. The minimum speed is set to 15 Hz. As the fan is always on when the system is active, the value was merely set to `True` permanently.

To achieve a correct airflow in the model after the characteristic curve is implemented was to use the simulation monitor, as explained in **4.3.1**, to do trial and error estimation by adjusting the `Ref_dP` for the ducting's in the system.

These values were already set as information was given in the available documentation for the other components. Only slight adjustments were needed, e.g., a pressure drop measurement was available for the filter in the AHU for which a better fit could be obtained.

As seen in **Figure 25** by the four marked circles, some dynamics were not at first captured by the model, and an investigation into the scene was necessary. The fan speed signal was steady-state at 1 % (min. speed). By comparing plots, looking at the P&ID, and observing the simulations (looking at actual measurements), what happens around the time these four events occur. It became apparent that the dampers switched positions rather rapidly. The dampers have coupled behavior which means that the sum of both damper openings is always equal to 100 %.

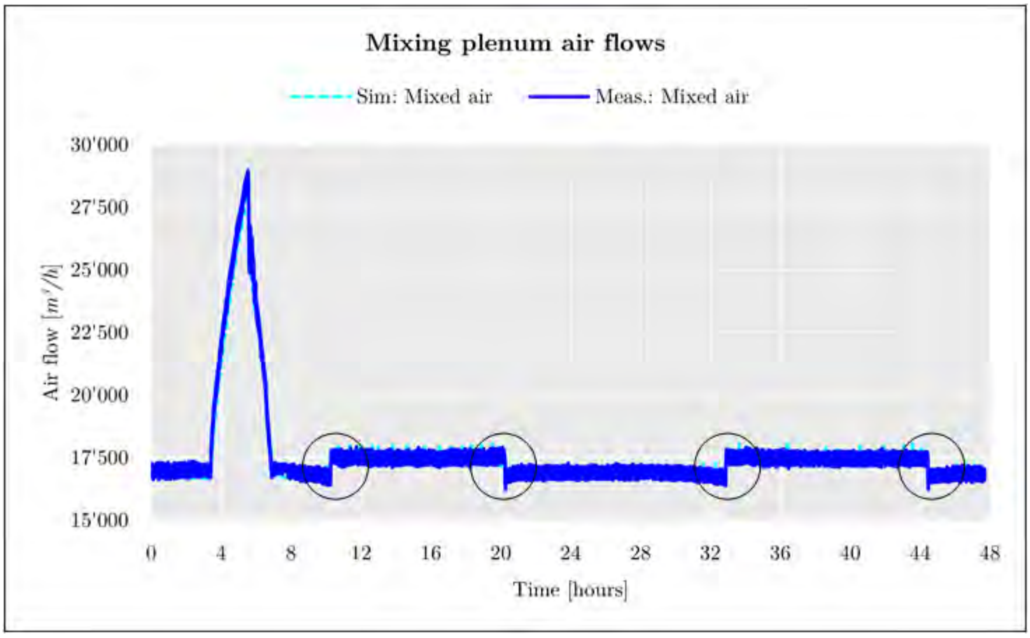

Figure 25: A plot of mixed/fan airflow (for UARX10001) is used to set the fan model parameters with a training dataset so that the fan can reproduce the measured values, also taking into account the dynamics caused by dampers changing position marked by the four circles (Own archive, 2022)

After examining the scenario further, it was deduced that the airway path between points B and A has more resistance to airflow than the path between C and A (**Figure 26**). It can be explained reasonably because the fresh airway path consists of actual ducting, a damper, and a filter. The recycled air path does not have any actual ducting's, only the extraction grille and the damper itself. Note that both the fresh- and recycled air dampers are identical, why this further add to the reasonability of the deduction.

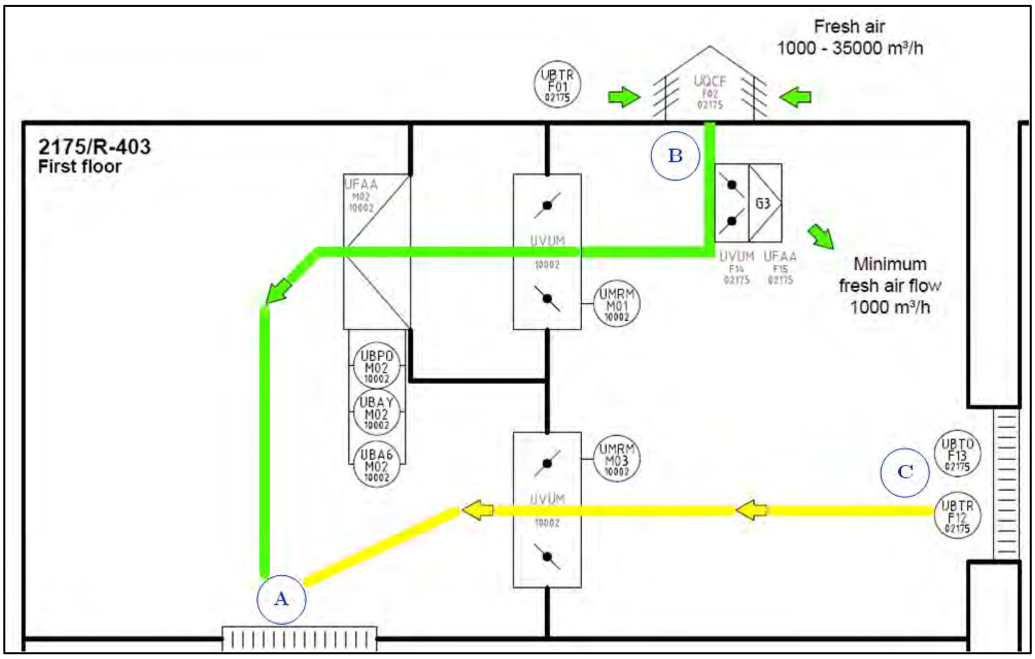

Figure 26: Visualization of the airflow paths (Cooling & Ventilation Group, 2021a).

With the behavior investigated and deduced with a reasonable explanation, it was implemented in the model. In **Figure 27** two ducting's respectively named, `fresh_air_inlet_duct_10001` and `airway_path_from_room_to_RAD_10001` can be seen. The two ducts are the ones that are interesting in this scenario. Other ducting's in the system are carrying the mixed airflow; why the position of the dampers will not affect the difference in airflow. The knowledge learned about the system, and its behavior in the concrete operational scenario was straightforward to implement in the model.

The `fresh_air_inlet_duct_10001` needed to be adjusted to a higher `ref_dP` value and unequal to `airway_path_from_room_to_RAD_10001`. It results in a bigger pressure drop when the airflow is coming through that airpath compared to the other. Consequently, affecting the total airflow in the system as the bigger resistance will lower the airflow in the scenario (or slightly bigger in the opposite scenario).

As seen in **Figure 25** it was successfully implemented in the model.

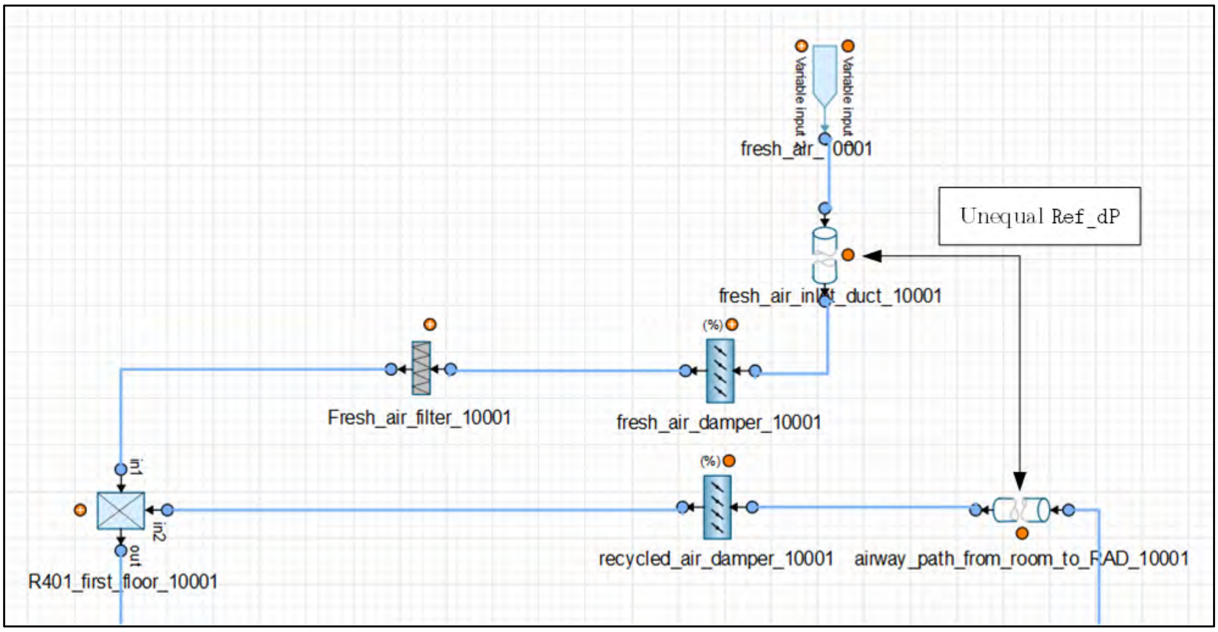

Figure 27: `fresh_air_inlet_duct_10001` and `airway_path_from_room_to_RAD_10001` duct components where they were set with unequal dP's to mimic the actual plant behavior (Own archive, 2022).

A few different datasets were used to parametrize the fan, and a good fit was observed. No data, however, is stored with a fan speed above 75 % for both AHUs, so it was not possible to compare the model in this operational scenario. However, experiments at 100 % fan speed indicated that the values appeared realistic to what could be expected in such a scenario. As earlier explained, with the AHUs as identical subsystems, some variance in operation had been observed in the data. It can be considered due to variations of the heat load distribution locally in the room.

Furthermore, resistance in airway paths is different, the orientation of the building concerning the effect the sunlight has locally in the room and at last, simply down to sensors being slightly inaccurate, which was observed numerous times. With that in mind, it was still considered the minor differences between the two AHUs, and parameters were successfully fitted for the other AHU (UARX10002) as it generally had a slightly higher airflow. However, minor discrepancies are seen when the dampers are operating in approx.

70/30 % to 30/70 % and is considered due to a not entirely consistent set of damper authority parameters.

4.3.3 Air dampers

The dampers already mentioned in the former subsection also significantly influence the airflow. However, the main thing is how it affects the airflow and distribution between the fresh- and recycled air when the dampers are changing position, which is essential.

The damper authorities significantly influence the percentage of airflow allowed as a function of the percentage of opening.

Every damper will have an inherent flow characteristic that has been measured in a controlled environment with constant differential pressure. A generic example can be seen in **Figure 28**.

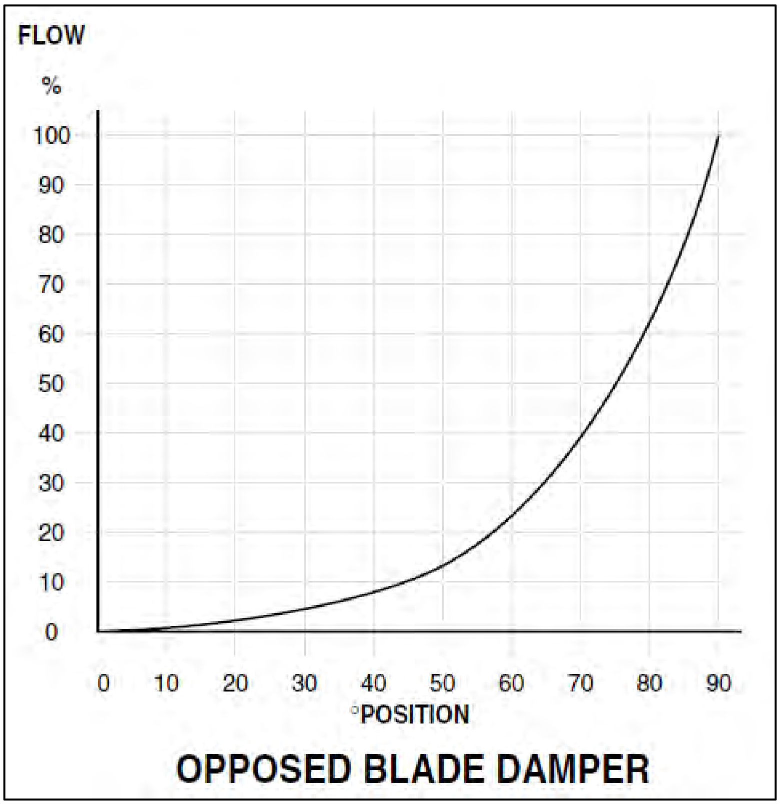

Figure 28: Generic example of how an inherent flow characteristic look for an opposed blade damper, as is the case with the dampers present on the SR1 plant (Belimo, 2019a).

When installed in a system network, only in rare cases will the flow characteristic follow the inherent curve. Therefore, the damper authorities measure how the flow profile for the damper is in the actual installed system. However, measurements can only decide them on the actual plant, so it is not usually a known factor. Therefore, experiments are needed to find consistent damper authorities in the model, as is the case with the SR1 plant.

The damper authorities are calculated using the following equation:

$$\text{Damper authority} = \frac{\text{Open Damper } \Delta p}{\text{Total System } \Delta p} \times 100 \ [\%] \tag{3}$$

The Total System Δp is defined between points A and B in a system where it is the damper and only the damper that controls the airflow.

The damper is not the only part in a system (between points A and B) that causes pressure drops and affects airflow. Using **Figure 26** as an example, it was found that the airway path between B and A had more resistance than the airway path between C and A. Given that the dampers are identical, this information can again deduce that the fresh air damper (point B to A) should have a higher damper authority (higher = closer to 0 %) than the recycled air damper. The extent and what the damper authority values are cannot be known, only that the fresh air damper authority should be higher. In **Figure 29** can be seen a generic example of how different damper authority values affect the flow profile.

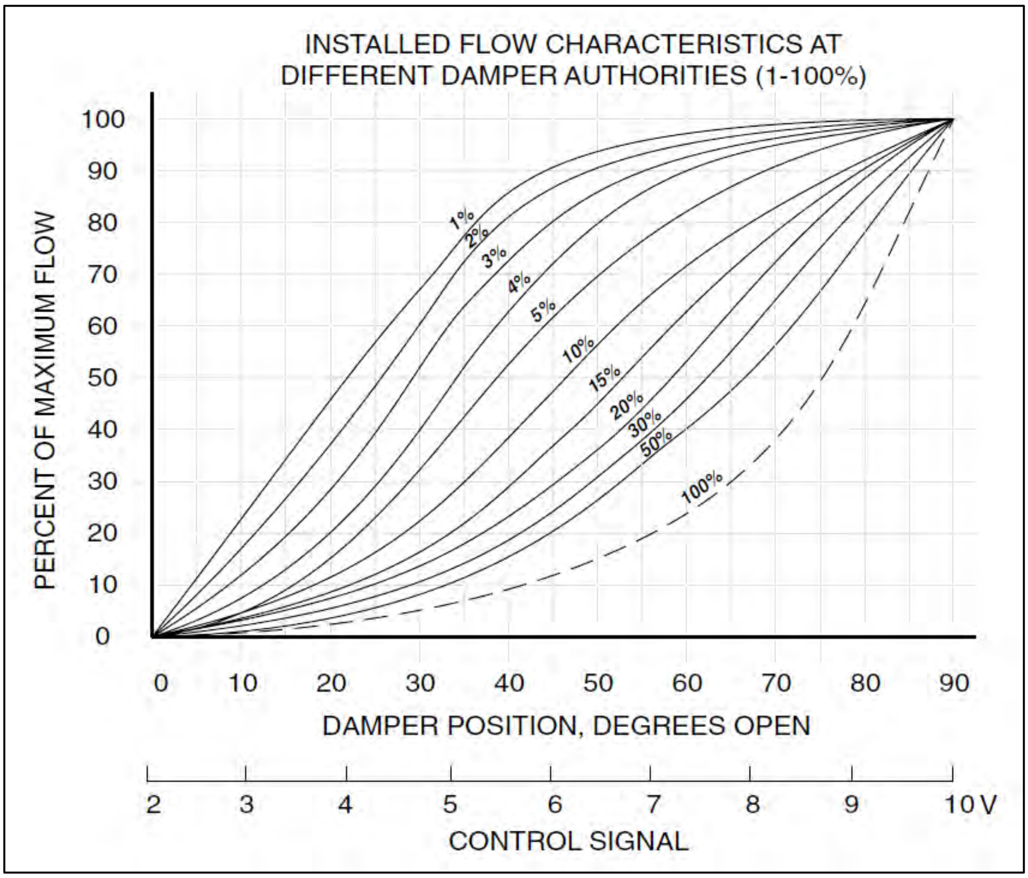

Figure 29: Generic example of an installed flow characteristic for an opposed blade damper (Belimo, 2019a).

About the parametrization of the dampers on the SR1 plant, the fresh air- and recycled air damper have coupled behavior. Knowing that they have different flow characteristics makes it very unintuitive compared to the fan model parametrization. The coupled behavior of the dampers means they are locked together so that when one is opening, the other is closing.

With different flow characteristics and no other knowledge than the fresh air damper authority should be higher than the recycled air damper authority, proved this to be an underwhelming starting point. Both dampers allow airflow into a mixing plenum with different flow characteristics; consequently, a resultant mixed airflow and resultant temperature will be present, but this was then the measure used.

In **Figure 30**, an example of a resultant total flow can be seen as a function of the mixing ratio between the fresh- and recycled air damper. This example is a case with a near-optimal matching of the dampers with an almost linear mixing ratio fit as the dampers are operating and a nearly constant resultant airflow.

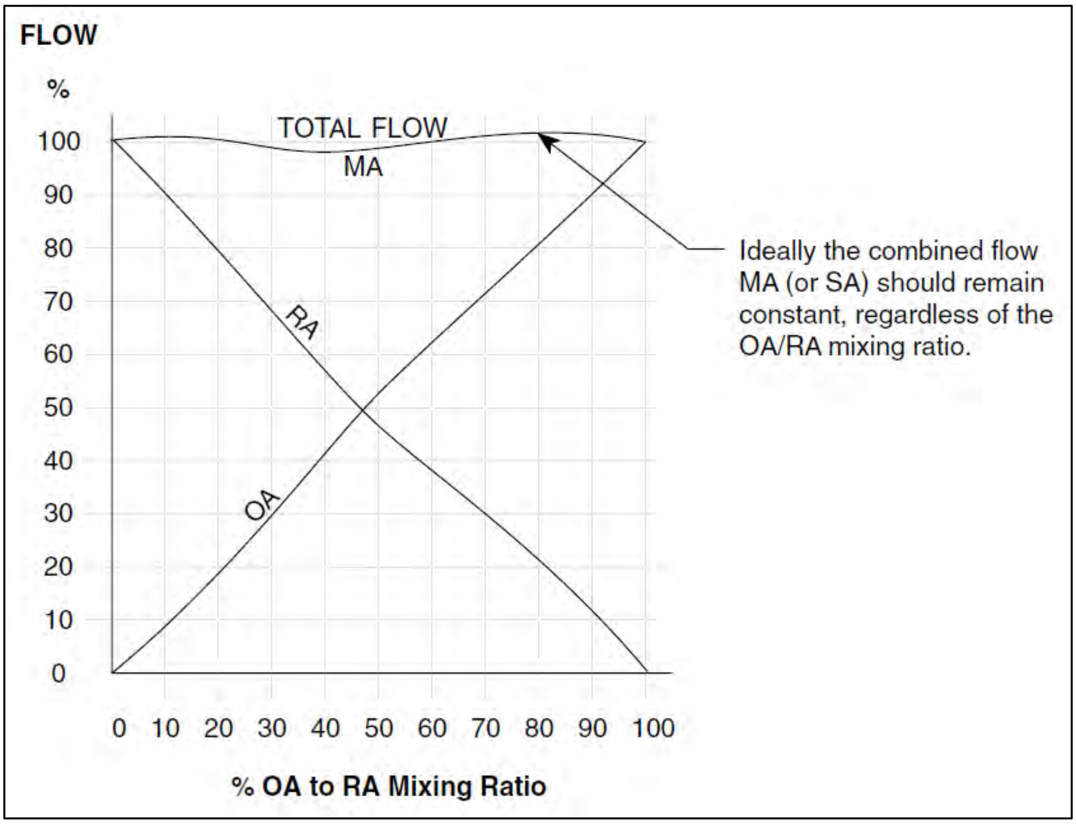

Figure 30: Generic example of remarkably good matched dampers that have nearly a nearly constant airflow and almost identical and linear characteristic airflow percentage as a function of the air mixing ratio (Belimo, 2019a).

As already seen in **Figure 25** on page **38**, the airflow changes when the dampers swiftly change their positions opposite. A difference of roughly 800 m^3/h is measured with a min. fan speed corresponds to nearly a 5 % difference in the mixed airflow. With this knowledge, it is the case that the resultant airflow characteristic is not staying constant, and knowing that the mixing ratio has non-linear characteristics.

When changing them, several different parameters influence each other, making it vastly tricky to identify a good set of parameters for the damper authorities. A careful approach to the investigation was used, and the primary variable to use as a measure was the mixed air temperature. Already modeled were the airflows in the system during steady-state and transient simulations, dampers fully opened/closed, and the fans operating up to 75 %. These results were consistent and thus considered reliable. Suppose the simulated recycled air temperature was relatively consistent with the measured value. Different datasets could be applied to try and experiment with and identify parameters that could mimic the dynamic behavior and precision of the measured value. In a scenario where the operational condition of the plant is in free cooling mode, it is only the changing in damper openings that adjusts the temperature of the mixed air and also the supply air and thus makes it possible to have an excellent platform to do the experimenting. The different datasets applied ideally need to include the entire operating range of the dampers. This difference in behavior is reflected in the data in different seasons, e.g., in the summer, the fresh air damper is primarily open, and the opposite scenario in the winter. In the spring/fall, behavior is seen where the dampers more often and consistently operate with a ratio closer to 50/50 %, which also revealed that the airflow fit deviated slightly more than presented in **Figure 25**.

As mentioned, many different factors influence each other, so it was not possible to the same extent as with the fan model to find a consistently good set of parameters. However, compared to measured values, the airflow discrepancies are considered mainly down to the damper authorities.

An indication of this is possible in **Figure 31**, where the discrepancy is believed to be caused by the damper authorities as the airflow else is consistent. Therefore, the discrepancy is, even though the awareness considered minor to no meaning.

The difference is roughly 3 % in the scenario, but slightly more significant discrepancies have also been observed.

After much analysis, parameters that could produce relatively good results were found and applied, although not perfectly consistent in all operating ranges.

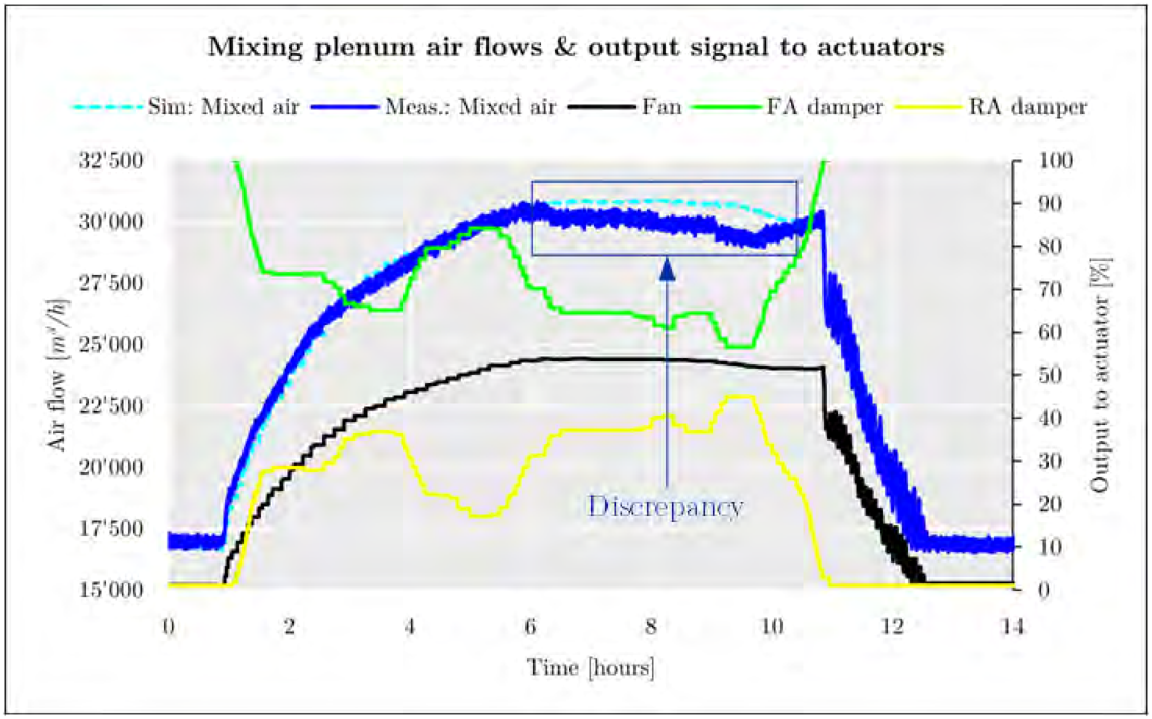

Figure 31: Plot with airflows (y_1 axis) and the three actuators output signal (y_2 axis) that influence the airflow. A minor discrepancy is displayed which is thought to be the damper authority settings not being completely consistent (UARX10001) (Own archive, 2022).

Minimum fresh air flow damper

As presented in **4.1.2**, there is a fixed position fresh air damper present before the fresh air damper, which allows a minimum of $1'000\ m^3/h$ at all times. However, it was impossible to model directly in EcosimPro, as the modeled components would not allow two input sources of fresh air. It can be down to several reasons, e.g., the appearances of algebraic loops, which is not ideal (as they are difficult to solve mathematically). Therefore, other measures had to be thought of to try and model this fresh air flow input to the process as it is not unimportant.

In **Figure 32**, the approach taken to find a practical solution is seen. The solution is not ideal but results in a better outcome than not taking any measures to fix the encountered model limitation.

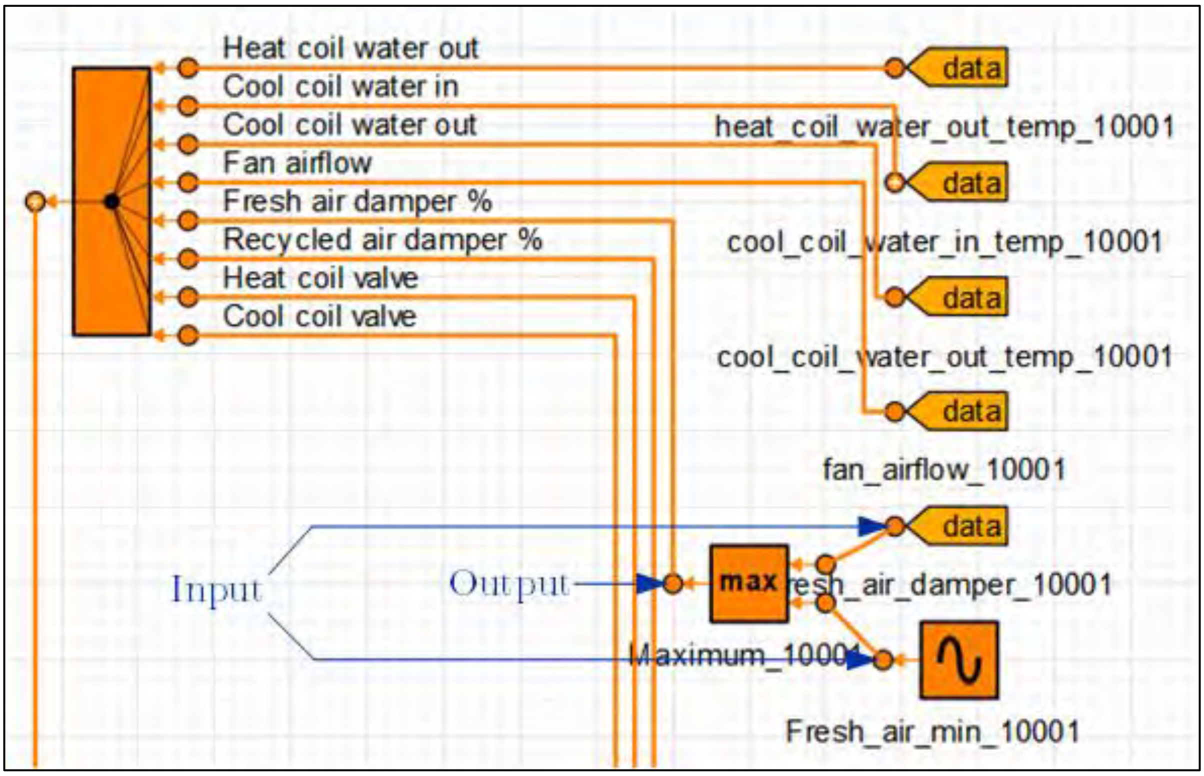

Figure 32: Max component which outputs the biggest value of the two inputs (Own archive, 2022).

A max component available in EcosimPro was used. The component has two input ports and one outport port. The fresh air damper signal is fed into the component in one of the input ports. Then the solution was to manipulate the fresh air damper position in the lower region by setting a fixed value in the other input port, e.g., 15 %. The component will always output the biggest value of the two inputs. If the signal reduces to 0 %, there is no airflow; the fixed value of the other input will then be the bigger value, which will dictate the output and allow a specific airflow aimed at approx. $1'000\ m^3/h$, which can be seen achieved in **Figure 33.**

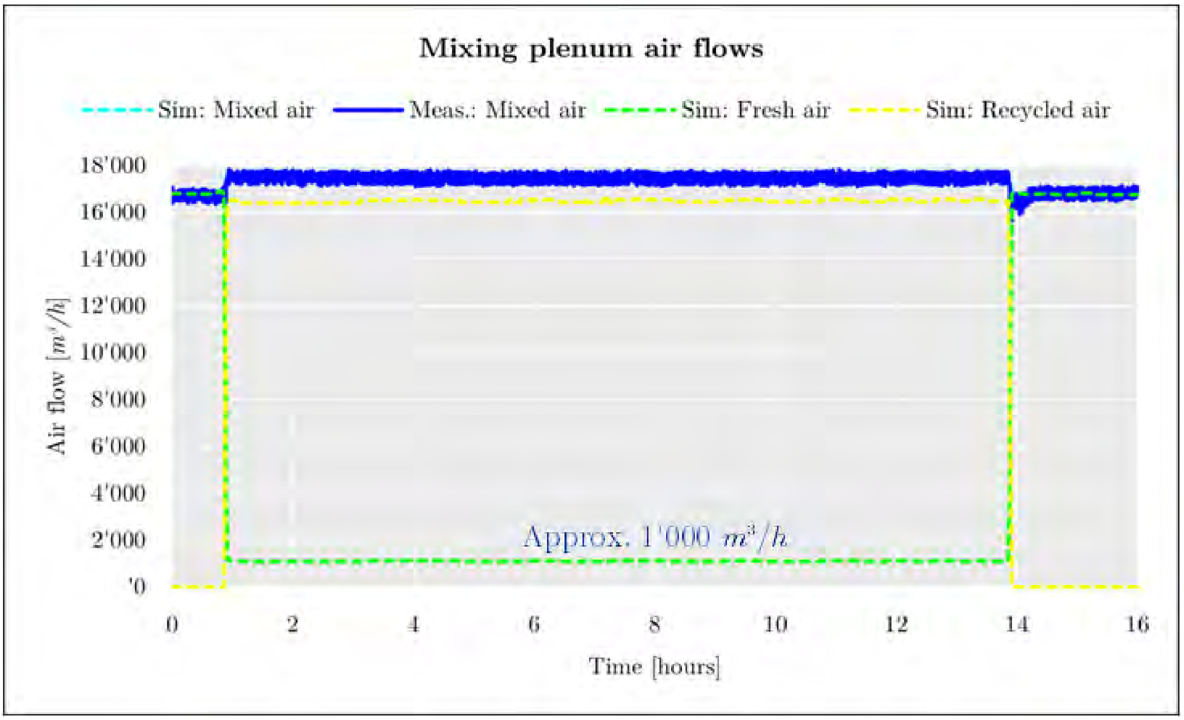

Figure 33: Depiction of the effect the manipulation of the fresh air damper has in order always to reach a minimum airflow of $1'000\ m^3/h$ when the fresh air damper is at 10 % opening (which is a minimum requirement also) (UARX10001) (Own archive, 2022).

Exhaust air dampers

As presented in subsection **4.1.2**, two exhaust air dampers are also present in the main hall. The exhaust air dampers could not be modeled for similar reasons as explained with the minimum fresh air flow damper with two input sources of fresh air. The function of the exhaust air dampers is merely to prevent a too big positive pressure inside the room, as it could have adverse effects on the building. When primarily fresh air is blown into the room, a positive pressure will build up with no active extraction of air. When primarily recycled air is blown into the room, a stable room pressure almost equal to the surrounding air absolute pressure is present. As this is not an essential part of the model affecting the main objectives, it still aimed to reproduce realistic behavior though with no pressure measurements. It was done by simulating a considerable leakage in the room to replace the dampers that could not be implemented. The dynamics that the dampers would be responsible for could not be taken into account, but the values outputted by the simulation are considered realistic. With the model built as described, there is a relation between airflow and the position of the fresh air damper with a more extensive positive pressure inside the room when the fan is ramping up and at the same time if the fresh air damper is fully open. On the actual plant, it would be expected that the most significant positive air pressure inside the room would be in a scenario where the fresh air damper is close to but still <20 % open (20 % dictates when the exhaust air damper opens) and possibly if the fan is ramping up.

The evolution of the absolute room pressure can be seen depicted in **Figure 34**.

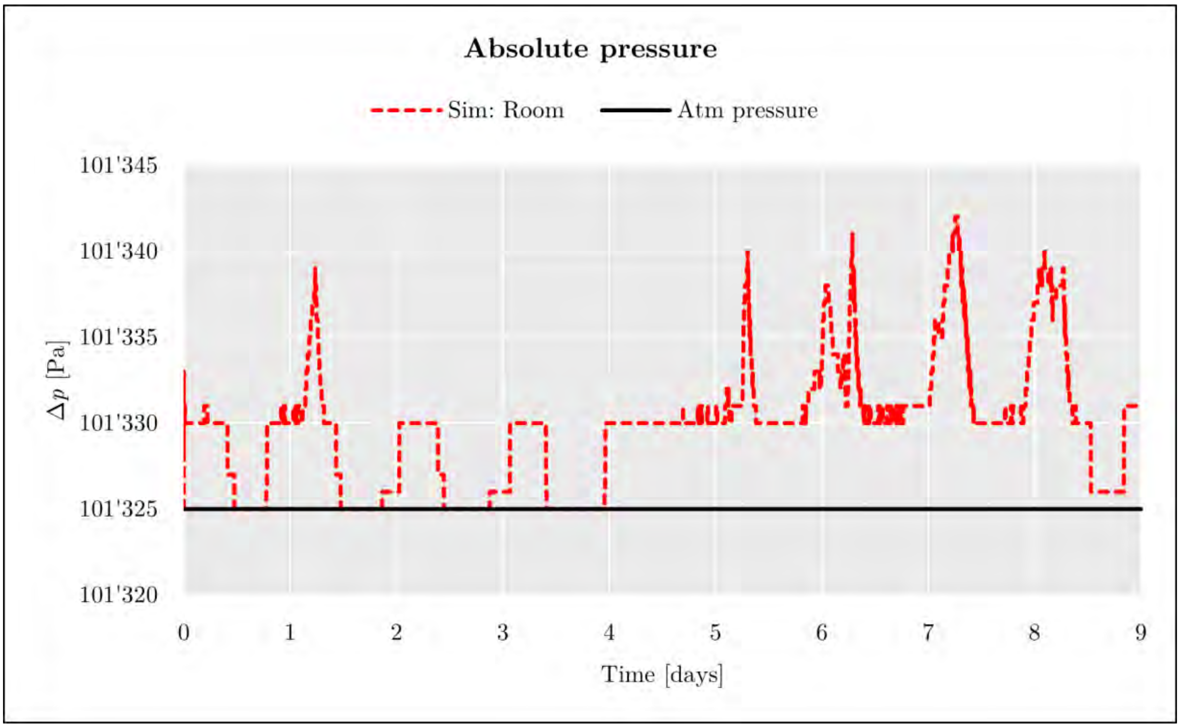

Figure 34: Evolution of room absolute pressure over 9 days. Note that both the fan and dampers of the AHUs are operating, and not only the dampers (Own archive, 2022).

4.3.4 Heat exchanger & control valve

Each AHU also includes a heating- and cooling coil to exchange heat by either adding or removing it. It is done by controlling the flow of water with a control valve. Free cooling mode is a priority seen from an energy-optimal perspective (not considering fan speed), but the control system will request either heating or cooling. Typically, there will mainly be operation in active heating- and cooling mode in the winter and summer.

Heat exchanger model

The heating- and cooling coil models were mainly parametrized given the data from the manufacturer's datasheet and not considered afterward.

Early on, experimenting with the heat exchanger model and the closely related valve model (controlling water flow through the coil) and following an investigation, the deficiencies were found in the model. The CERN HVAC library made for EcosimPro has previously successfully been used, including the heat exchanger model, but even though deficiencies were found.

Several things were found to affect the model applicability:

- No consideration of a fluid volume in the coil
- No consideration of in- and out pressures in the coil water network
- Simulation problems occur at times

The investigation revealed that the model would work well assuming a continuous control signal above 0 % is applied to the valve. It will allow a flow of water and consequently an exchange of heat. It has been the case in earlier models that have been used. With the method used in this thesis for estimating parameters using historic control system data, the problem emerged as the control signal to the heating- and cooling coil valves operated between 0-5 % opening and had a characteristic behavior that was almost like on/off signals. It provided many oscillations, which could be expected given that the dynamics should be relatively fast, and the same behavior is seen in the operational data. The problem in this context is that the valve is also operating at 0 % opening in between openings. However, because of the non-existent consideration of a water volume, the model will immediately stop transferring heat if there is no flow through the coil. On the actual plant, a continuous, though more minor (and decreasing) heat transfer is the case that this model will not consider.

Another discovered problem was the in- and out pressures provided as input sources to the component. They need to be provided for it to work, but even though these are not taken into consideration and will therefore not affect the water flow rate.

The oscillations mentioned proved to be a fundamental problem that needed to be solved somehow. The feedback signal used from the historic control system data and its on/off behavior between 0-5 % openings meant that not enough heat was added/removed in the periods when the valve was operating.

Control valve

Closely related to the heat exchanger model is the valve controlling the water flow in the coil. With no parameters affecting the water flow, it is merely the control valve CV-value that will dictate it. Moreover, the flow characteristics (like the air dampers) are also interesting. However, a limitation in the model is present as it is only possible to select a linear characteristic valve or an isopercentual (logarithmic) characteristic valve. The valve on the plant is an equal percentage valve with an inherent flow characteristic. This characteristic will change when installed in a system, and with no means of manipulating the authority, it could prove challenging to reproduce the real valve behavior. For the exact case, there is mainly operating in the range between 0-5 % opening, which effectively means that the parametrization can only be based on this small portion of data which is not enough to get a complete meaningful result. Often, an equal percentage valve is chosen to achieve a relatively linear flow characteristic, and therefore a linear valve was chosen in the model. The way of parametrizing the CV-value of the valve was then to do trial and error estimations to achieve the nominal flow rate of water read from the test- and commissioning report of the plant with a 100 % open valve.

Measures to solve the discovered problems

Other measures had to be implemented to solve the deficiencies of the heat exchanger model and limitations in the valve model. As explained about how the valve models were parametrized, two problems need to be solved:

- Wrong flow characteristics
- Oscillations and insufficient heat added/removed given the on/off behavior and 0-5 % openings of the valve

With the approach chosen to parametrize the valve model, a discrepancy was accepted given the limitations of the valve model. Experiments deduced clear indications that the valves present on the plant had relatively high authority. It means that small valve openings will provide a relatively high water flow rate. Consequently, meaning that when applying the feedback signal as control input to the valve in the model, an insufficient amount of heat was added/removed because a lesser flow of water is the case in the model applying the same control signal.

For the problem with no consideration of a water volume, a solution was implemented with good results, but still, some discrepancies and oscillations needed to be accepted.

As mentioned, the control signal applied to the valve in the model was the feedback signal from the historic control system data. The solution was to use a component available in EcosimPro to manipulate the input signal to the valve using a lookup table with the x value being the feedback signal coming from the historic control system data and the y value the manipulated signal sent to the models' valve. The manipulation was done whenever the valve was operating between 0-5 % opening; a continuous signal was applied instead of 4 % (**Figure 35**). This value was again found by trial and error estimation as no other measures were available. The mean value of the signal for a defined period where it was operating was studied. The value was relatively low, but given the knowledge about the discrepancy between the flow rate with small openings, an over-compensation had to be made in the model to get a similar amplitude of heating/cooling. As the feedback signal was not precisely consistent, the y value was sometimes lower than the threshold set. Thus the appearance of some oscillations was still present but did not affect the overall effect the solution had.

The solution was implemented successfully, eliminating most of the problems with oscillations, correcting the amplitude of heating/cooling, and accepting the discrepancy in the valve model.

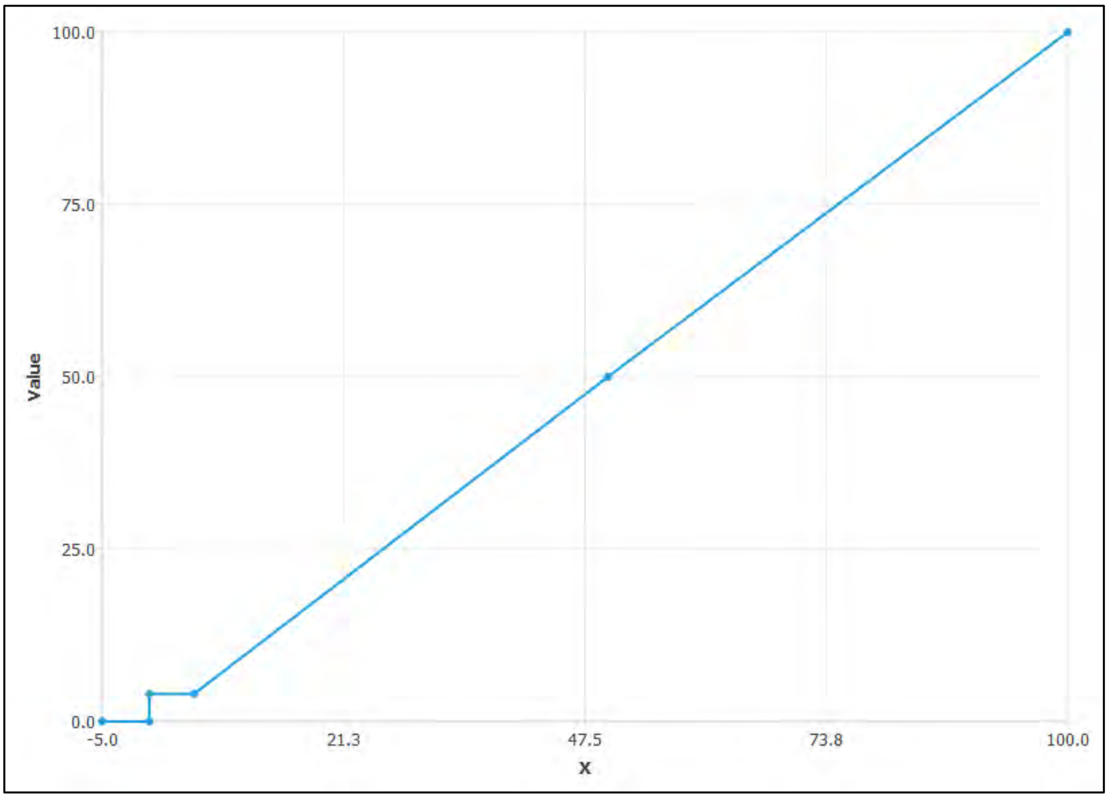

Figure 35: Lookup table component used to manipulate the control valve signal, y (Own archive, 2022).

4.3.5 Summary of the model development work

In section 4.3, a selected number of concrete areas of the model development was highlighted and explained as model parameters were identified and estimated. The final model version consists of $5'111$ variables and $1'692$ equations which is inevitable aiming to model a high-fidelity digital twin of the plant.

The focus was on the methodology and the way of thinking and approaching the work. Measures were taken to analyze the plant and its operational data to learn about it, discover possibilities and limitations in the model, and apply the best possible solutions.

The fan model was relatively straightforward to parametrize and showed excellent results with only minor discrepancies considered without meaning. However, some dynamics were discovered and analyzed to account for it in the model, which was also implemented.

The air dampers were not equally easy, and a careful approach was needed to find consistent damper authorities. Much time was spent on analyzing the experiments run and changing the parameters. However, because the dampers have coupled behavior, it was not possible to consider only one at a time. Adding to that is the influence of these variables on other variables, which made it even more difficult. A set of parameters were found that produced relatively consistent results. Discrepancies were consistent in some areas where the dampers are operating but considered acceptable.

The heat exchanger models and their control valves consumed a significant amount of time. Deficiencies were discovered in the heat exchanger model, and limitations in the valve model were found. The solution found and applied was not straightforward and had some discrepancies that needed to be accepted, but overall, the solution was valuable and good.

The model development naturally had to conclude at some point. The problems mentioned consumed much time, and thus the model development was prolonged. The model's state was not producing good enough results until late in the project period. The development work can always be continued as the iterative process allows the model to get better. However, the results produced in the end were considered good enough and near their full potential.

5

SIMULATION RESULTS AND DISCUSSION

With the model development concluded, this chapter includes results from the final simulations performed. Three different simulations will be presented. Common for these is the same validation dataset used.

The data is five days between the 19th to 24th of August 2021, which can be considered a month in the warmer half of the year where active cooling can be expected. First, a two-step validation study is performed on the developed model with an open-loop- and a closed-loop simulation presented in sections **5.1** and **5.2** respectively. After that, the model applicability is discussed in section **5.3**.

Section **5.4** will present the simulation, which ultimately was the objective - a virtual commissioning simulation with the developed MPC controlling the plant model.

5.1 Open-loop simulation for model validation

The open-loop simulation is performed similarly to the multiple experiments in the iterative process of parameterizing the model before the development concluded. The validation dataset extracted was applied, and the simulation results is visualized and presented between **Table 2** through **Table 7** on pages **51**-**56**.

The exact simulation type has a platform for direct comparison between measured and simulated values. As seen in **Table 2**, the measured recycled air temperature is a mean value of two sensor readings and therefore not directly comparable, but almost. The simulated value also coincides with the main hall air temperature in the model as it cannot consider different zones within a volume. However, the recycled air temperature is still the most meaningful to use for the parametrization. The statistical measure RMSE and MAPE, respectively, from equation **(1)** and **(2)** in **3.3**, is applied to the measured and simulated variables on the four different temperature measurements in the system. These are chosen because they all are included in the current regulation loop and therefore considered most important. The results are presented in **Table 8** on page **57**.

5.1.1 Plots

See **Table 2** through **Table 7** on pages **51**-**56**.

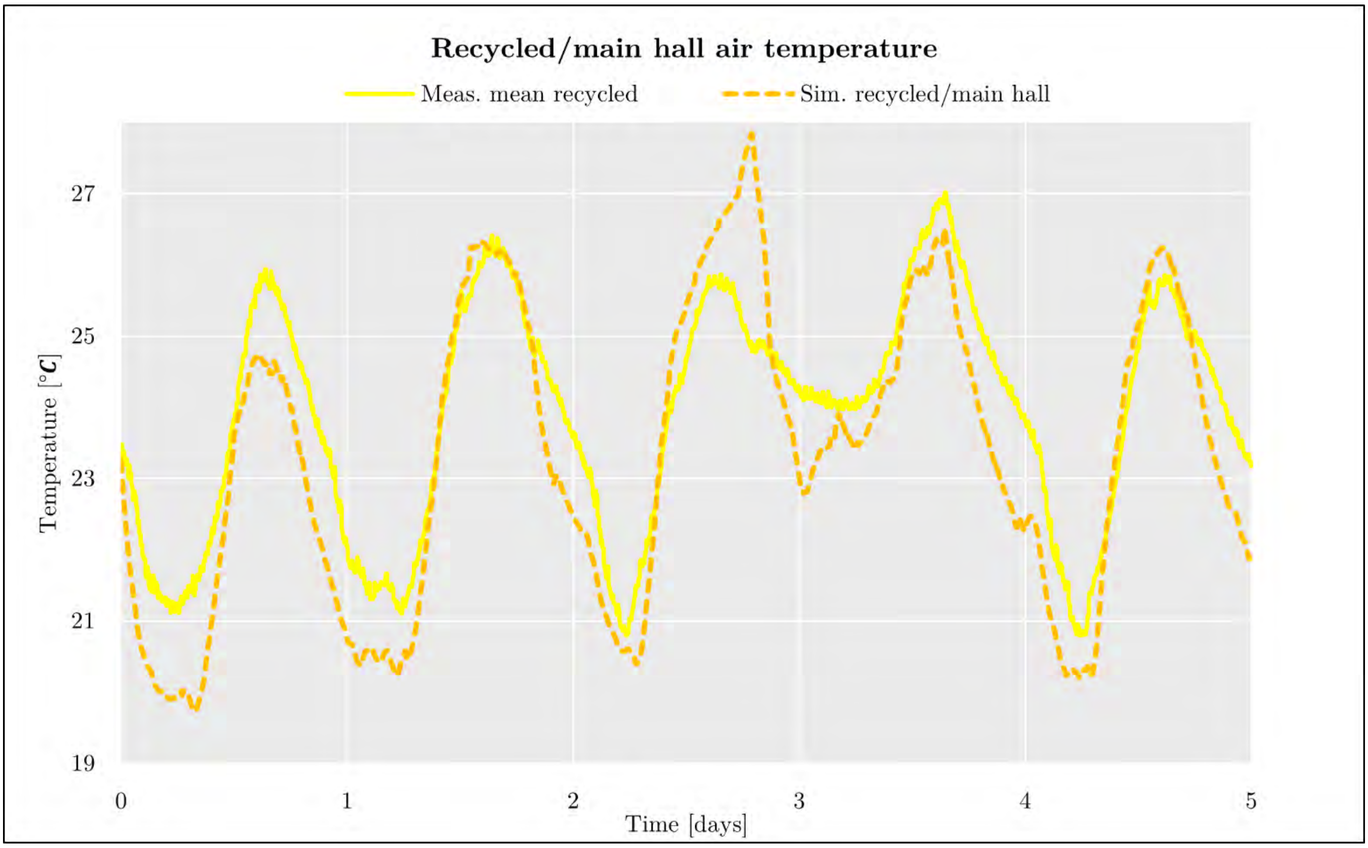

Table 2: Plot of the recycled/main hall air temperature with both the measured and simulated value (Own archive, 2022).

(a)

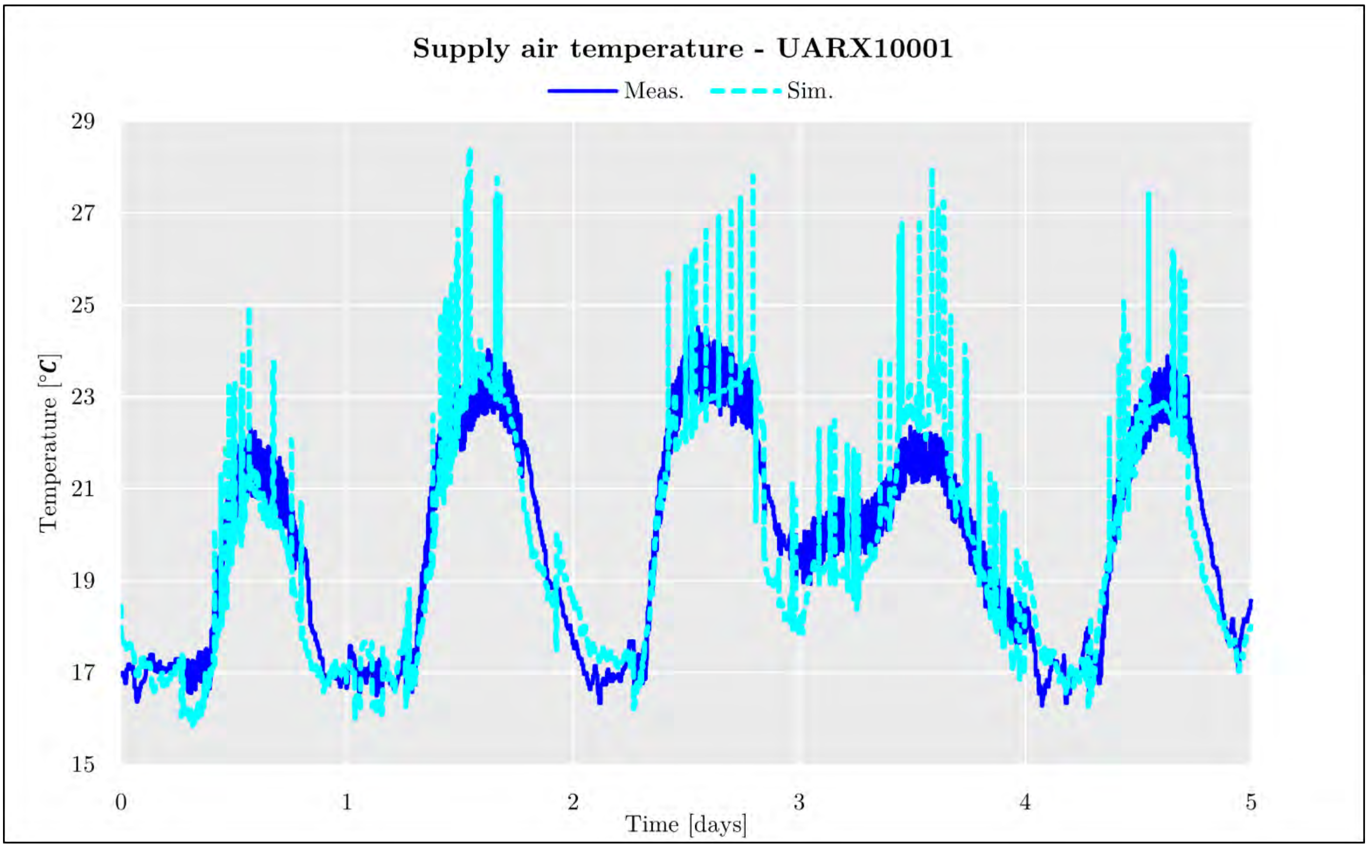

(b)

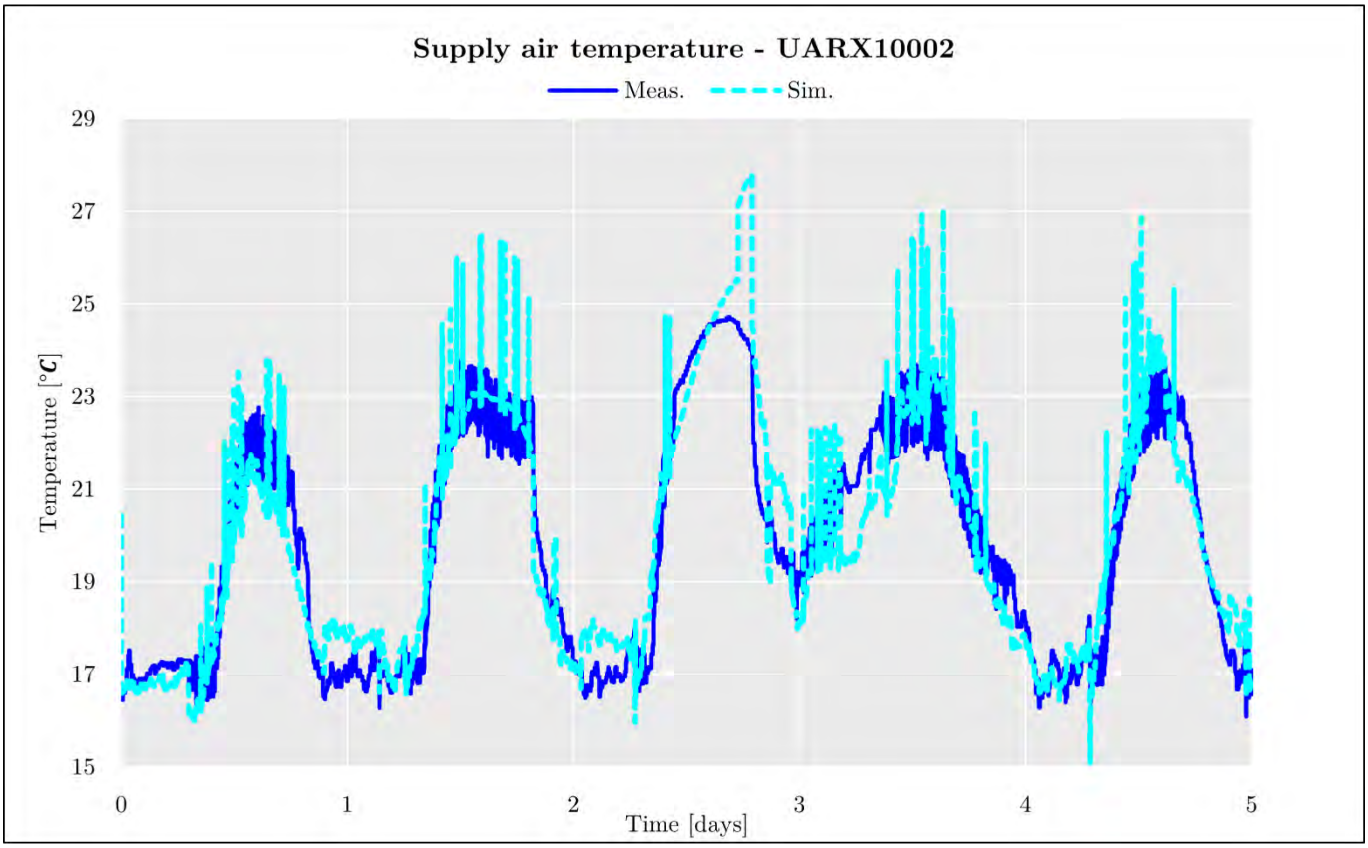

Table 3: Plot of the supply air temperature for both AHU's in **(a)**, **(b)** (Own archive, 2022).

(a)

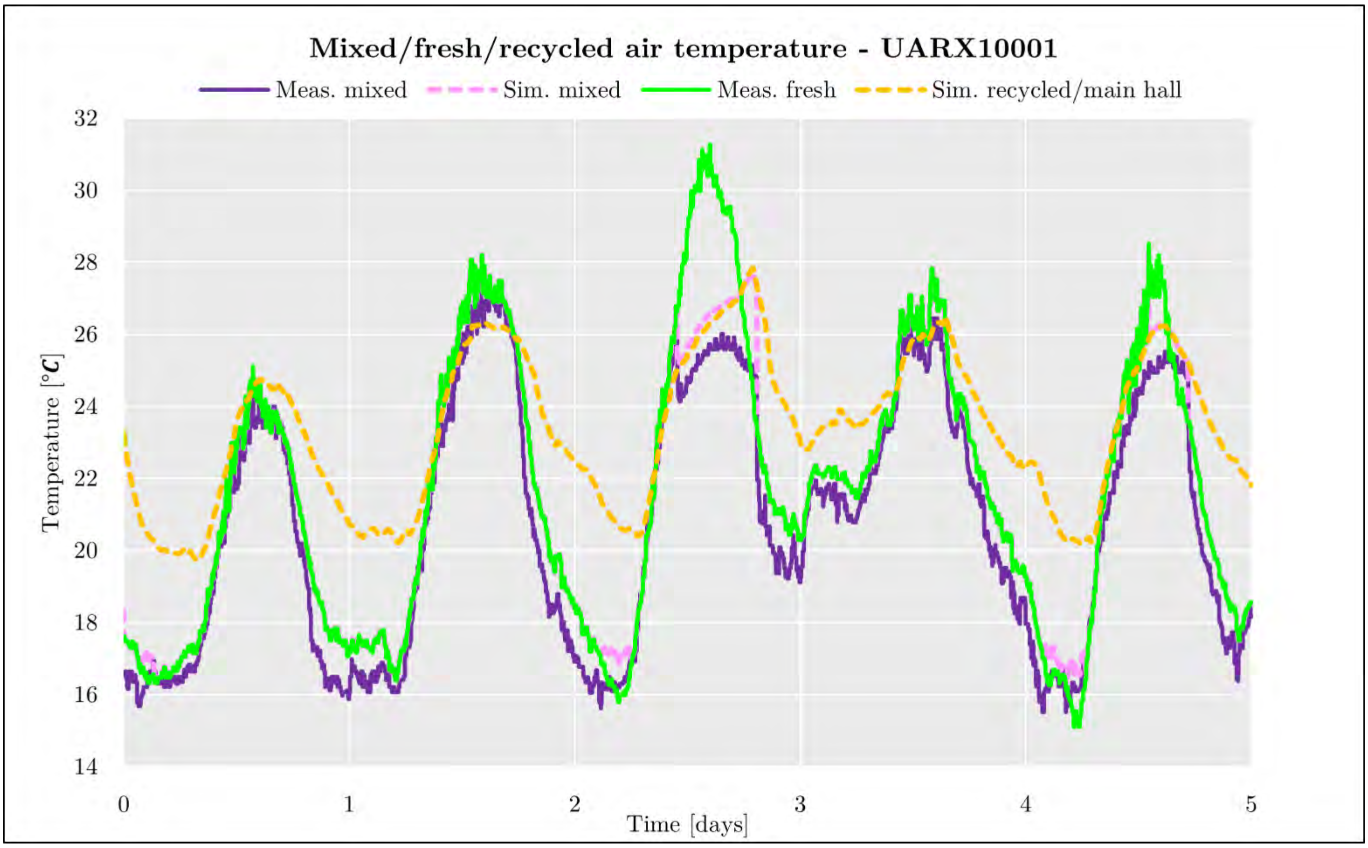

(b)

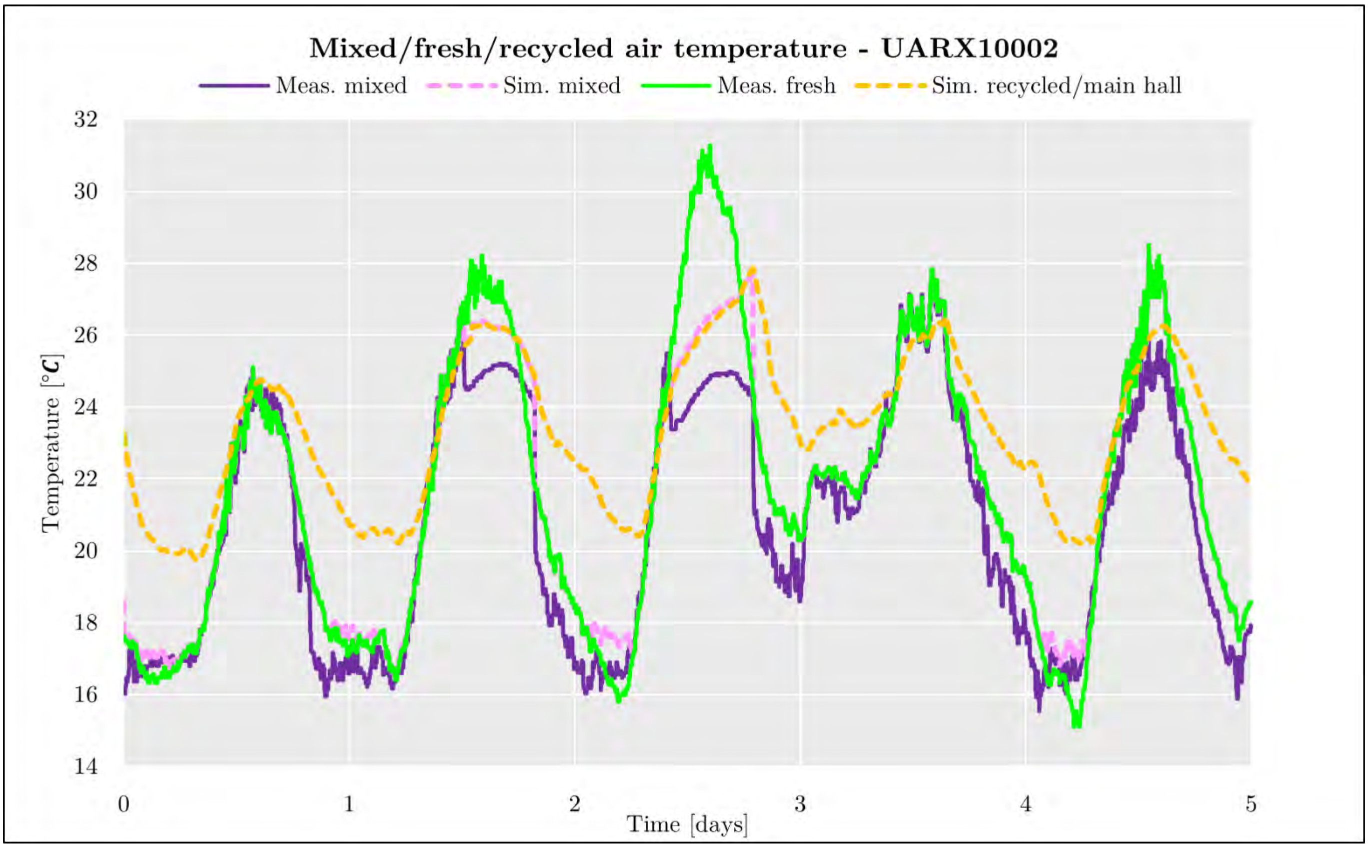

Table 4: Plot of mixed, fresh, and recycled air temperature for both AHU's in **(a)**, **(b)** (Own archive, 2022).

(a)

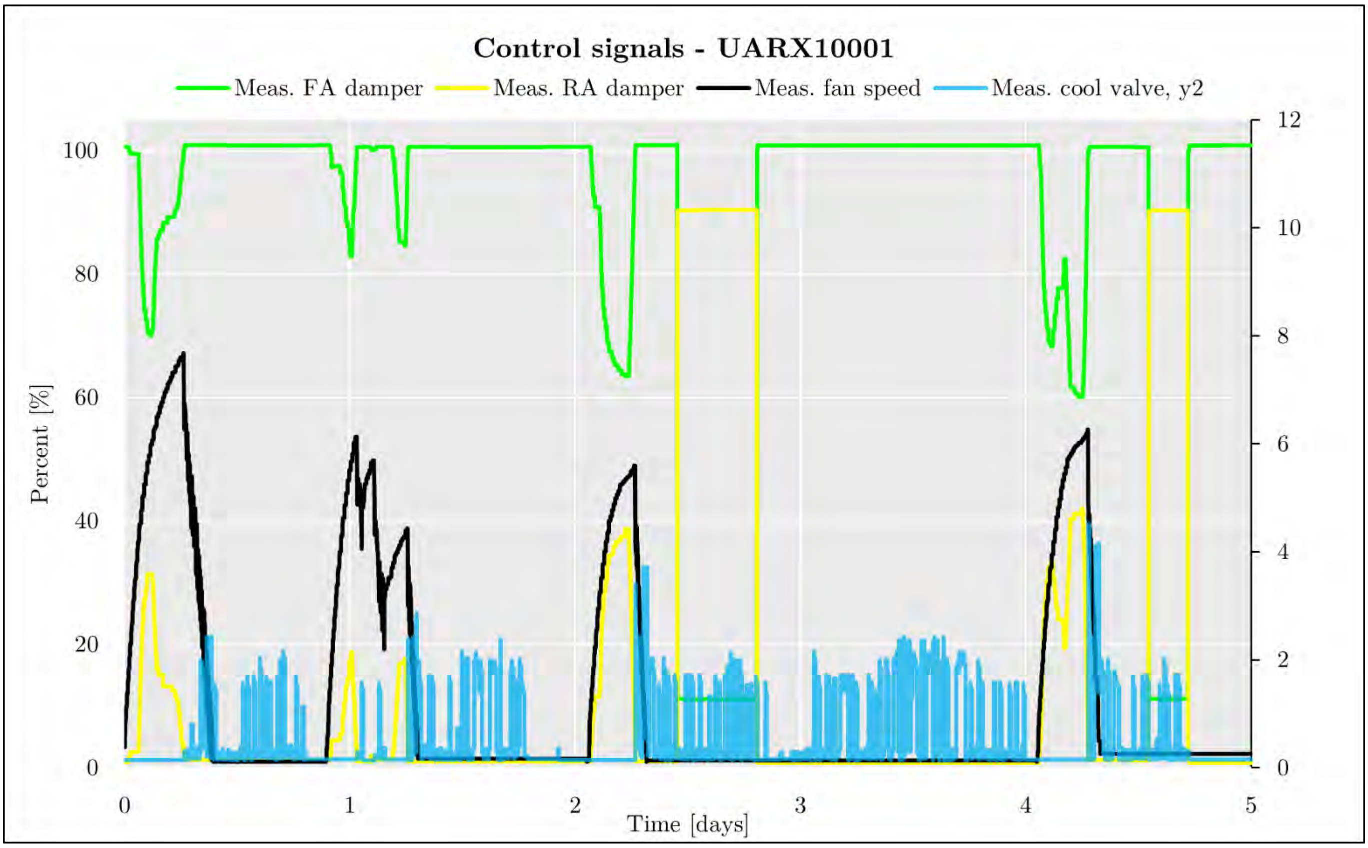

(b)

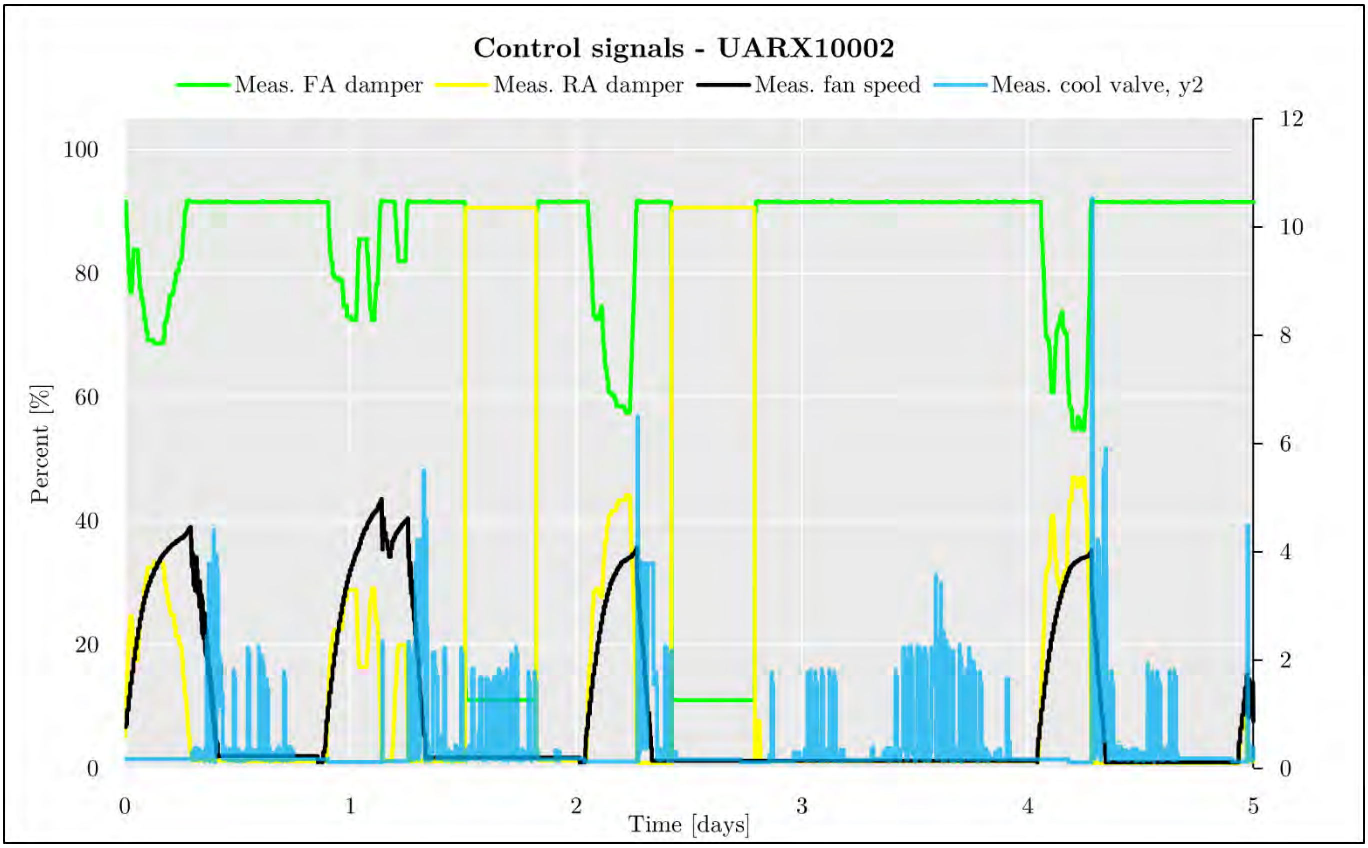

Table 5: Plot of FA- and RA damper, fan speed, and cool valve control signals in **(a)**, **(b)** for both AHU's (Own archive, 2022).

(a)

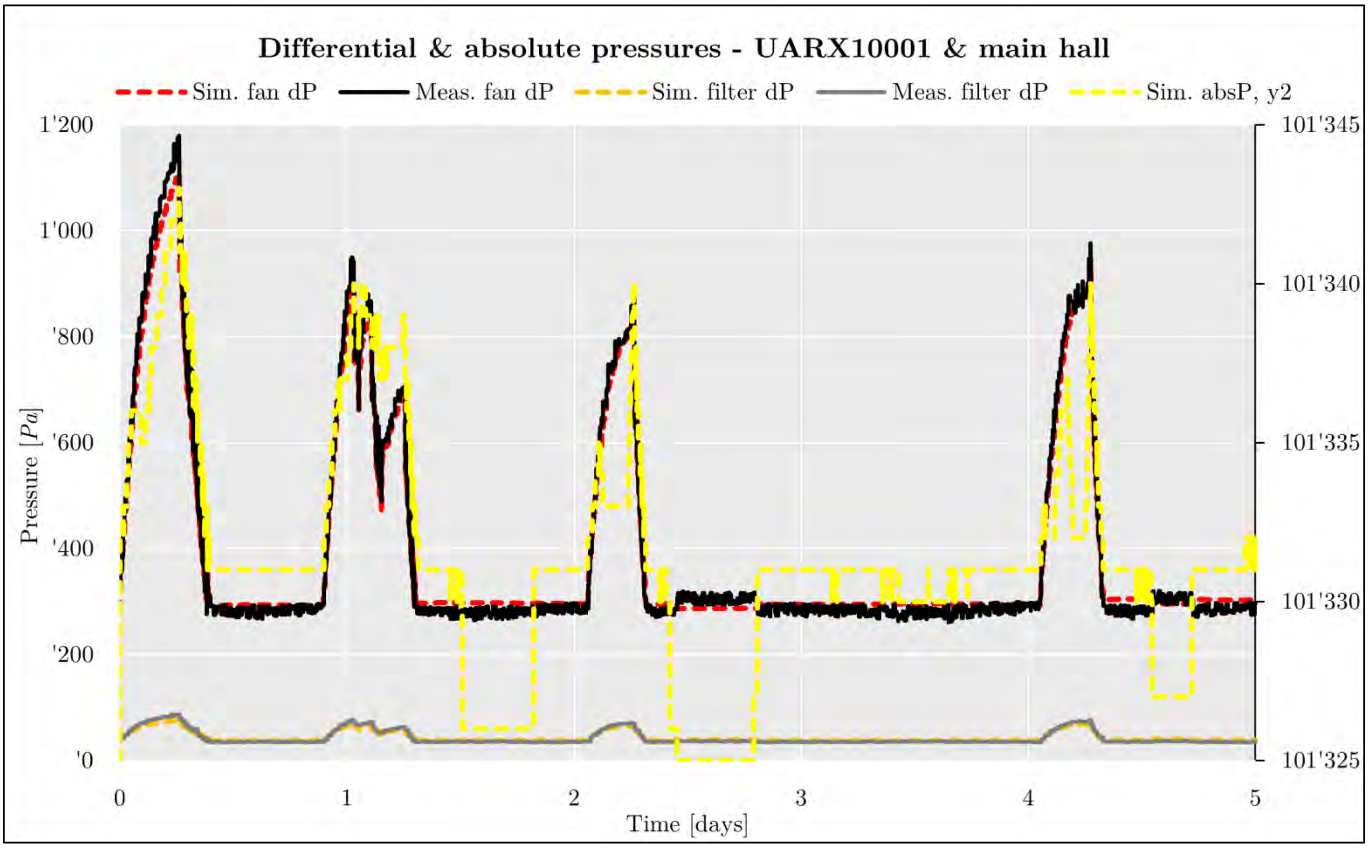

(b)

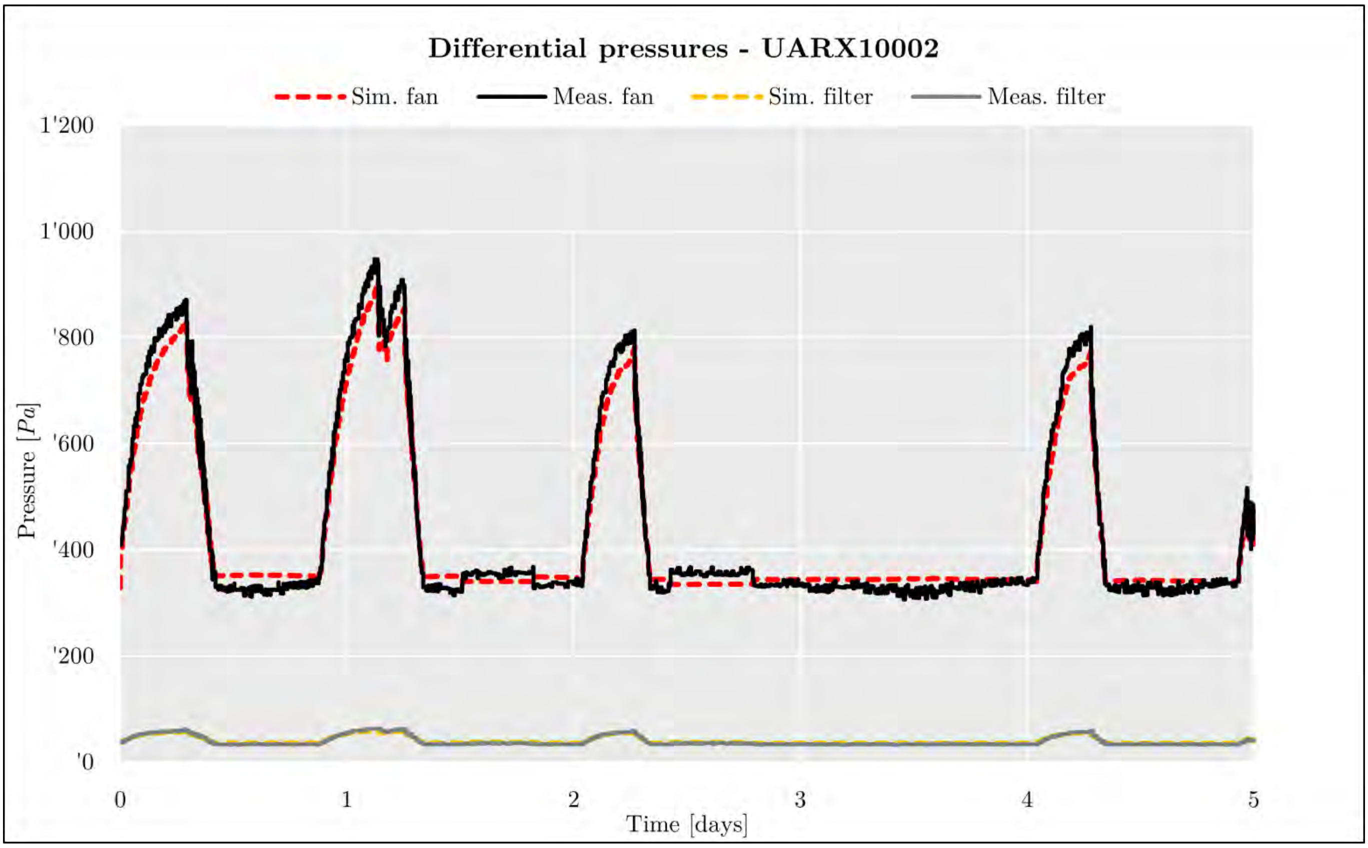

Table 6: Plot of dP's across the fan and filter for both AHU's in **(a)**, **(b)**. In **(a)** the main hall absP is seen (Own archive, 2022).

(a)

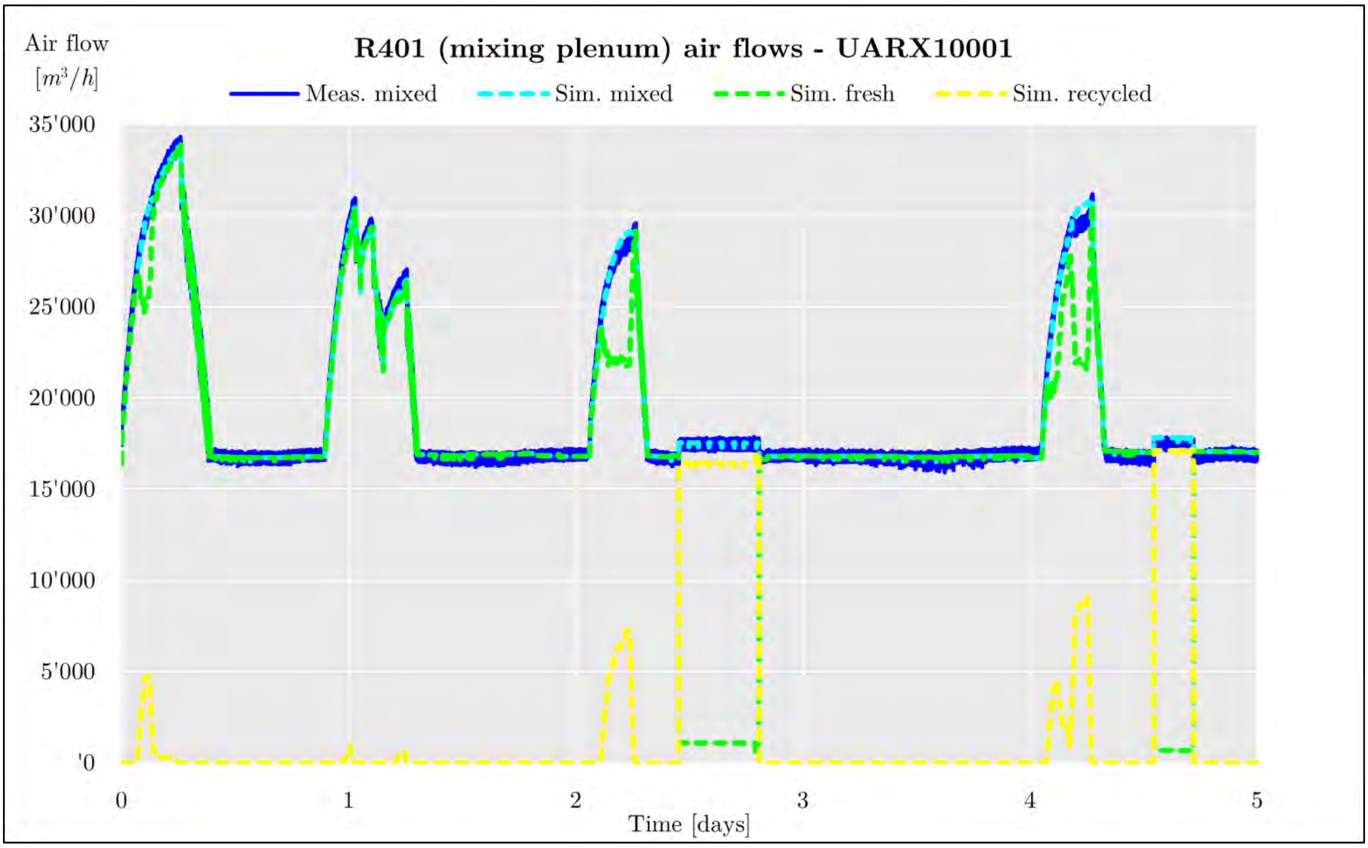

(b)

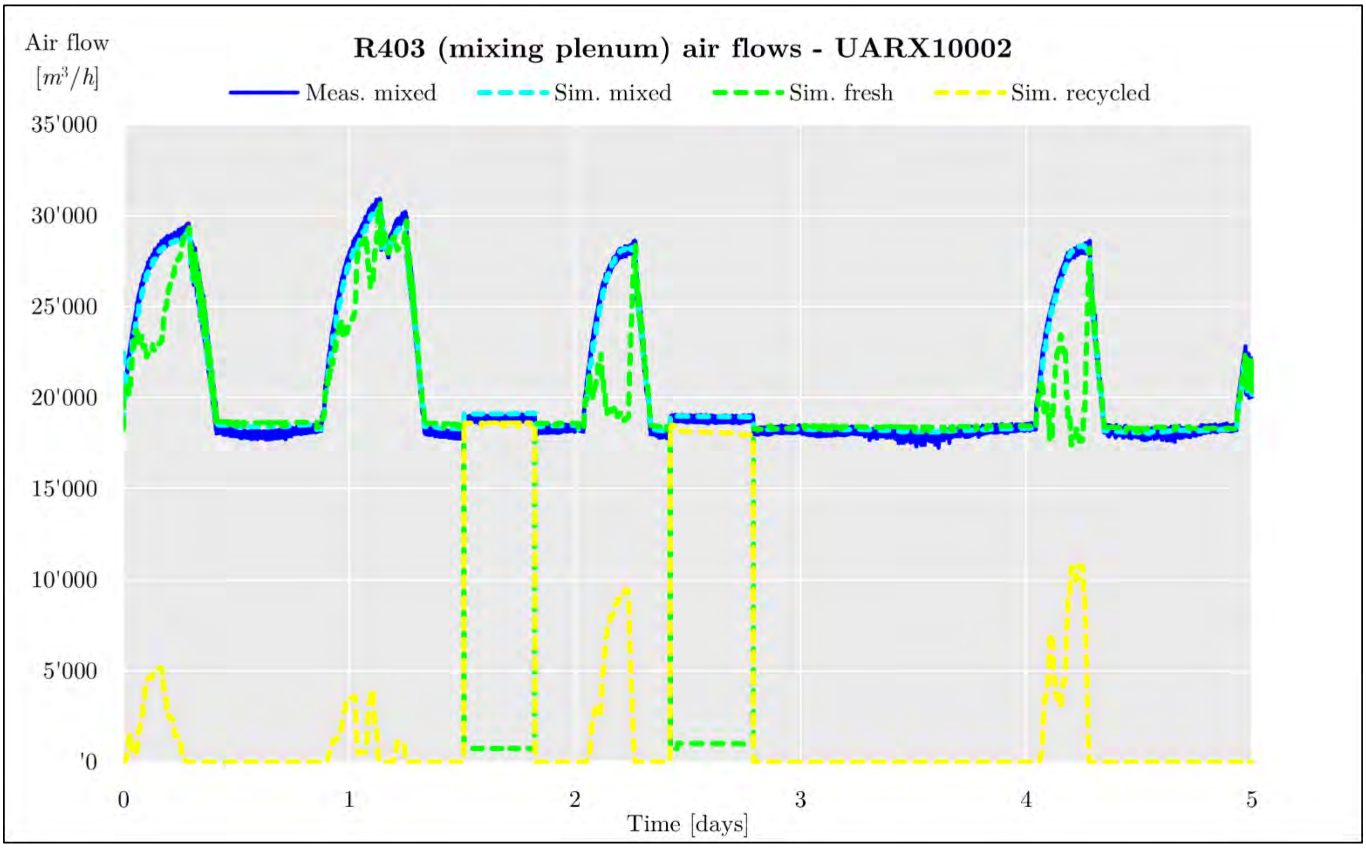

Table 7: Plots of the mixing plenum airflows for both AHU's in **(a)**, **(b)** (Own archive, 2022).

5.1.2 RMSE & MAPE

Measured & simulated variables	AHU	RMSE [$°C$]	MAPE [%]
Recycled/main hall air temperature	#	0.991	3.608
Supply air temperature	UARX10001	0.933	3.551
	UARX10002	0.959	3.666
Mixed air temperature	UARX10001	0.987	4.241
	UARX10002	1.176	4.639
After coils temperature	UARX10001	1.117	4.825
	UARX10002	1.108	4.515

Table 8: RMSE and MAPE on various temperature variables (Own archive, 2022).

5.1.3 Interpretation of the open-loop simulation

Already presented in **4.3.2**, a good fit for airflows and pressure drops in the system networked was seen. In **Table 7 (a)** and **(b)**, the results from the validation simulation can be seen with the measured and simulated mixed air flow for both AHUs having a nearly perfect fit. In **Table 6 (a)** and **(b)**, similar results can be seen for pressure drops across the fan and AHU filter, with the simulated values tending to have a slightly less peak value than the measured data. The evolution of the absolute pressure in the main hall is depicted but cannot be directly compared to measured values, even if available. The main hall pressure was of interest as it was impossible to model the exhaust air dampers, so other measures were applied. The main hall absolute pressure from the simulation is considered within realistic values.

Air temperature measurements

The model's ability to predict the four different temperature measurements is most notable.

In **Table 2**, **Table 3 (a)**, **(b)**, and **Table 4 (a)**, **(b)** is the recycled/main hall, supply and mixed air temperature presented, respectively. Common for all is that the model agrees with the measured data with a slight offset and some discrepancies.

The recycled air temperature is a mean value of two measurements on the actual plant as the room model cannot consider the local evolution of temperature. The model can produce one recycled air temperature, the same as the temperature within the main hall, and does so with an RMSE of slightly less than 1 $°C$ and a MAPE of 3.6 %. The heat load is set at a fixed value of 50 kW. Besides the time-varying heat load (heat transfer through walls), which is being calculated in the room model, indications can be seen that the room model heat load is not entirely correct. A possible dynamic heat load profile can be necessary if more precise results are to be achieved. In the first through the third day in **Table 2**, the simulated temperature has an increasing tendency with an undershooting on the first day, a near-perfect fit on the second day, and some overshooting on the third day with the measured value being almost consistent. In **Table 4 (a)**, **(b)** where the fresh air temperature can be seen to have a positive evolution towards higher peak air temperatures through these three days and influence the simulated recycled/main hall air temperature.

In **Table 3 (a)** and **(b)**, the supply air temperatures are depicted for each AHUs. The mentioned oscillations/peaks are easy to observe, and were impossible to eliminate completely, as argued in **4.3.4**. While it looks disturbing, the number of peaks is not that great, but these will possibly affect the RMSE negatively. The RMSE is calculated to be slightly under 1 $°C$ for UARX10001 and slightly above 1 $°C$ for UARX10002, respectively, with a MAPE of 3.551 % and 3.666 %.

Else are only minor discrepancies observed. Similarly, the after coils air temperature results are not depicted as they are near-identical to the supply air temperature as the measurement is before the fan instead of the supply air, which is measured shortly after the fan. An RMSE of ~1.1 $°C$ and a MAPE, respectively, for UARX- 10001 and 10002 is calculated to be 4.825 % and 4.515 %.

In **Table 4 (a)** and **(b)**, is the mixed air temperature plotted for each AHU and the measured fresh and simulated recycled/main hall air temperature. A clear example is these two plots of an ongoing problem/distraction, which always had to be in mind during the model development work. Sensor inaccuracies, or what presumably is, have been spotted on several accounts. No extraordinary measures have been invested in confirming this or if prominent disturbances affected the measurements, but this is, of course, something that hindered the work. The model does not take this into account, and with estimating the model parameters against this operational data, some values either did not make sense or were counter-intuitive to try and reach. For example, at the one-day mark on plot **(a)**, the measured mixed air temperature is lower than the measured fresh air temperature even though in **Table 5 (a)**, the fresh air damper is observed to be fully open. The coils are after this measurement, so these will not affect it. The simulated mixed air temperature is thus the same as the measured fresh air temperature and can only be seen (dotted pink) deviating when dampers change position.

In addition to the offset possibly caused by sensor inaccuracies an else good agreement between the measured and simulated variables is seen. The most considerable discrepancy is between the two- and three-day mark applying to both AHUs, with UARX10002 from **Table 4 (b)** having the more significant deviation of the two.

It can be observed to be caused by the overshooting of the recycled/main hall air temperature at the same time. The mixed air temperature RMSE was calculated under 1 $°C$ for UARX10001 and 1.176 $°C$ for UARX10002 with a MAPE of 4.241 % and 4.639 %, respectively.

5.2 Closed-loop simulation for model validation

The closed-loop simulations on the model were done by creating a setup with a production-ready version of the existing controls loaded into two PLCs, each one controlling one AHU. The model can replace the actual plant with some limitations, e.g., the model cannot differentiate between the main hall- and recycled air temperature, which effectively will affect the regulation. The signals to the exhaust air dampers which is not present can just be overwritten with no adverse effect on the simulation.

The differences are relatively minor but will affect the regulation and the results, which will be discussed in **5.2.2**.

5.2.1 Plots

See **Table 9** through **Table 12** on pages **60-63**.

(a)

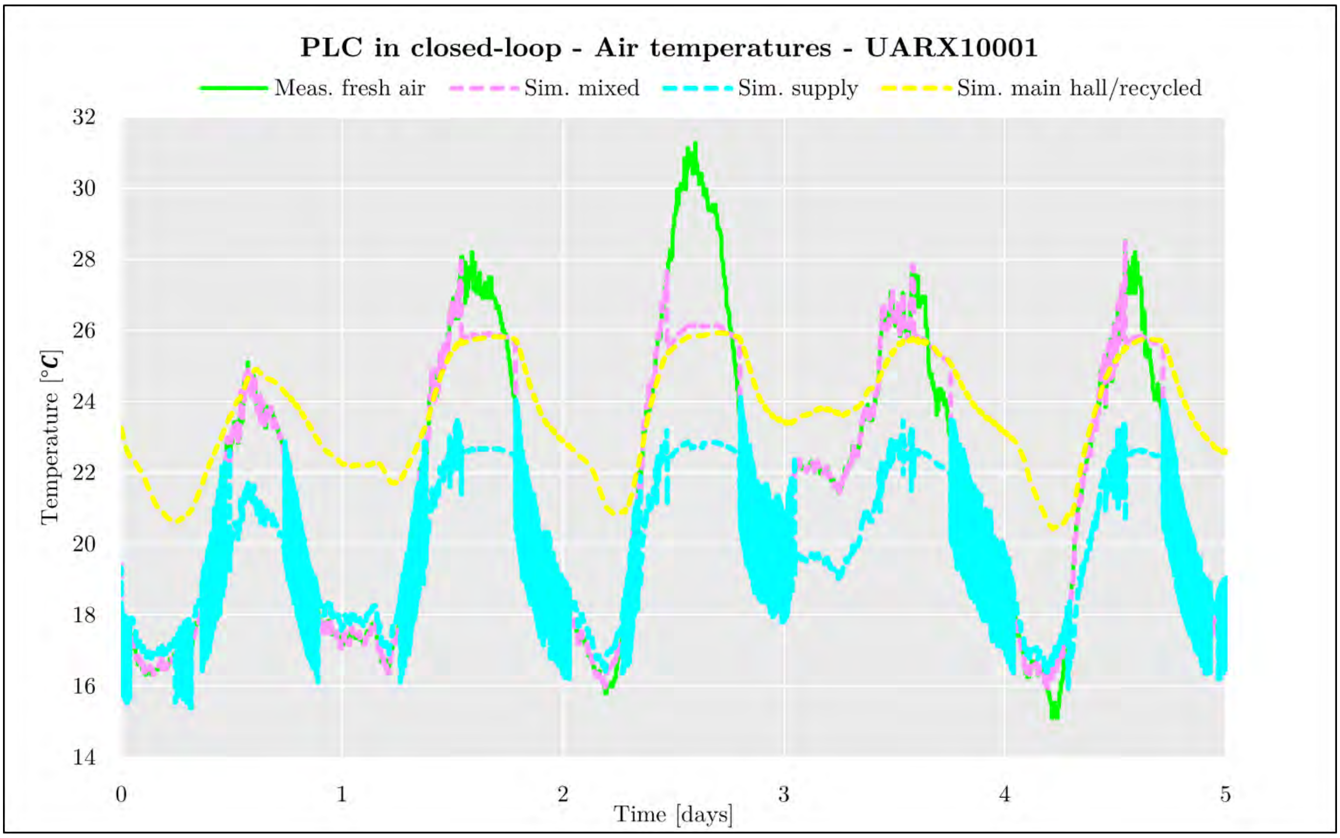

(b)

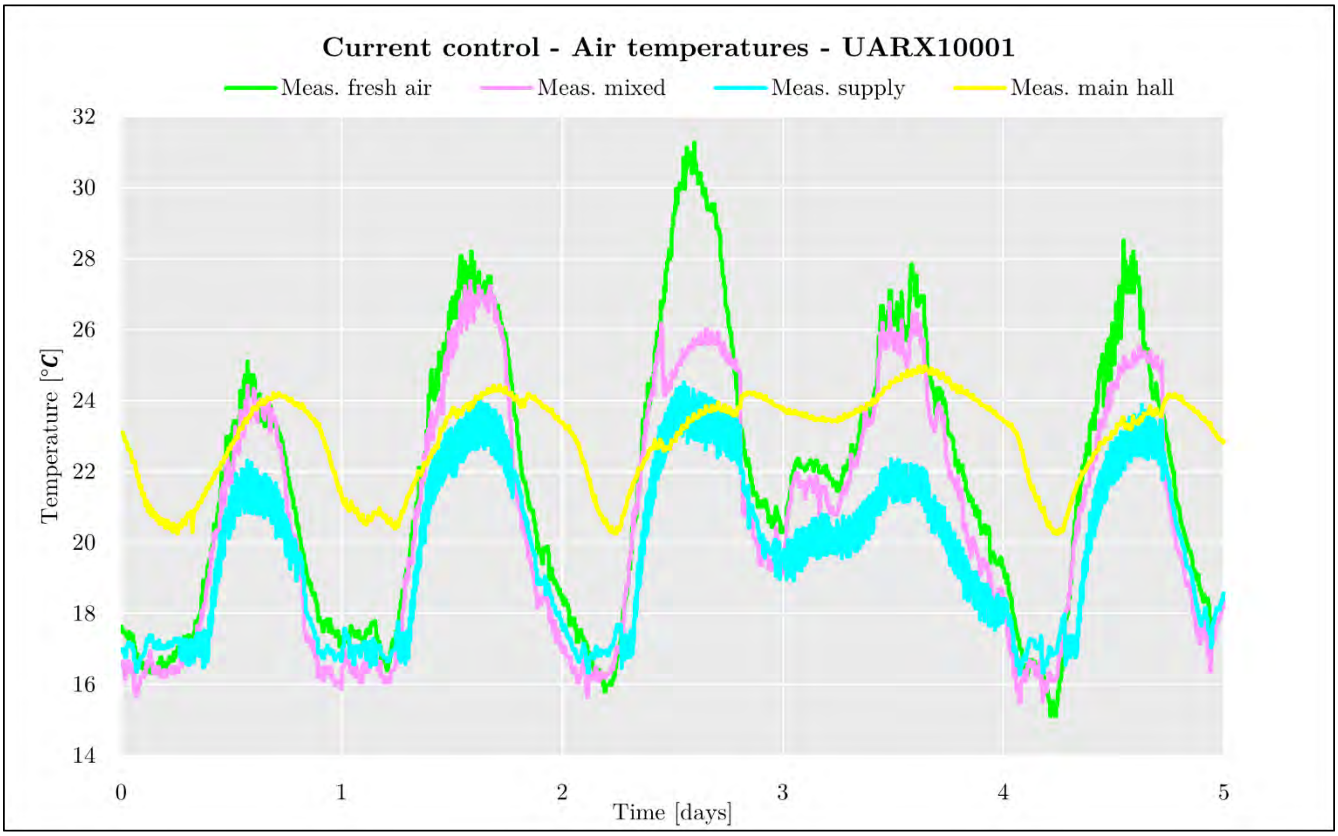

Table 9: PLCs running the model and the current control on temperature measurements for UARX10001 (Own archive, 2022).

(a)

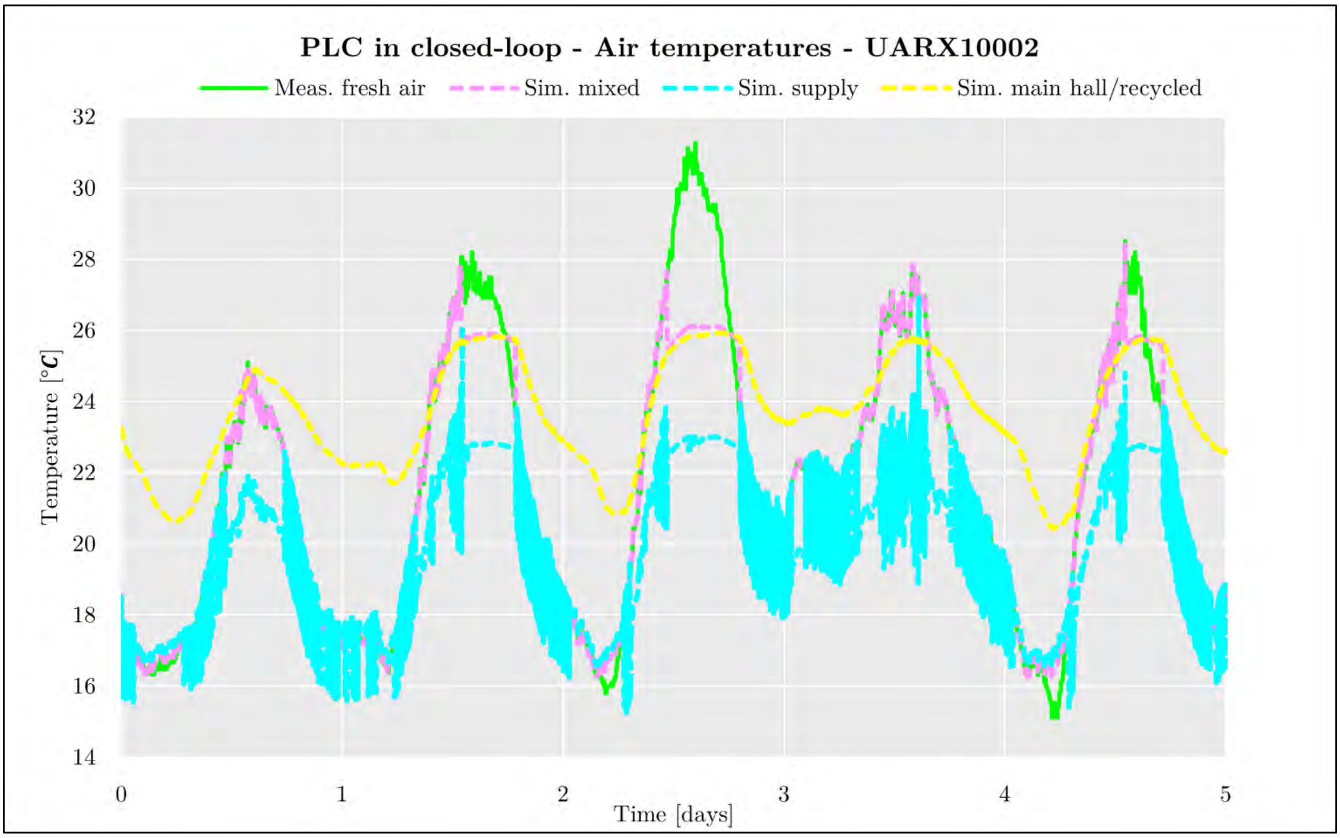

(b)

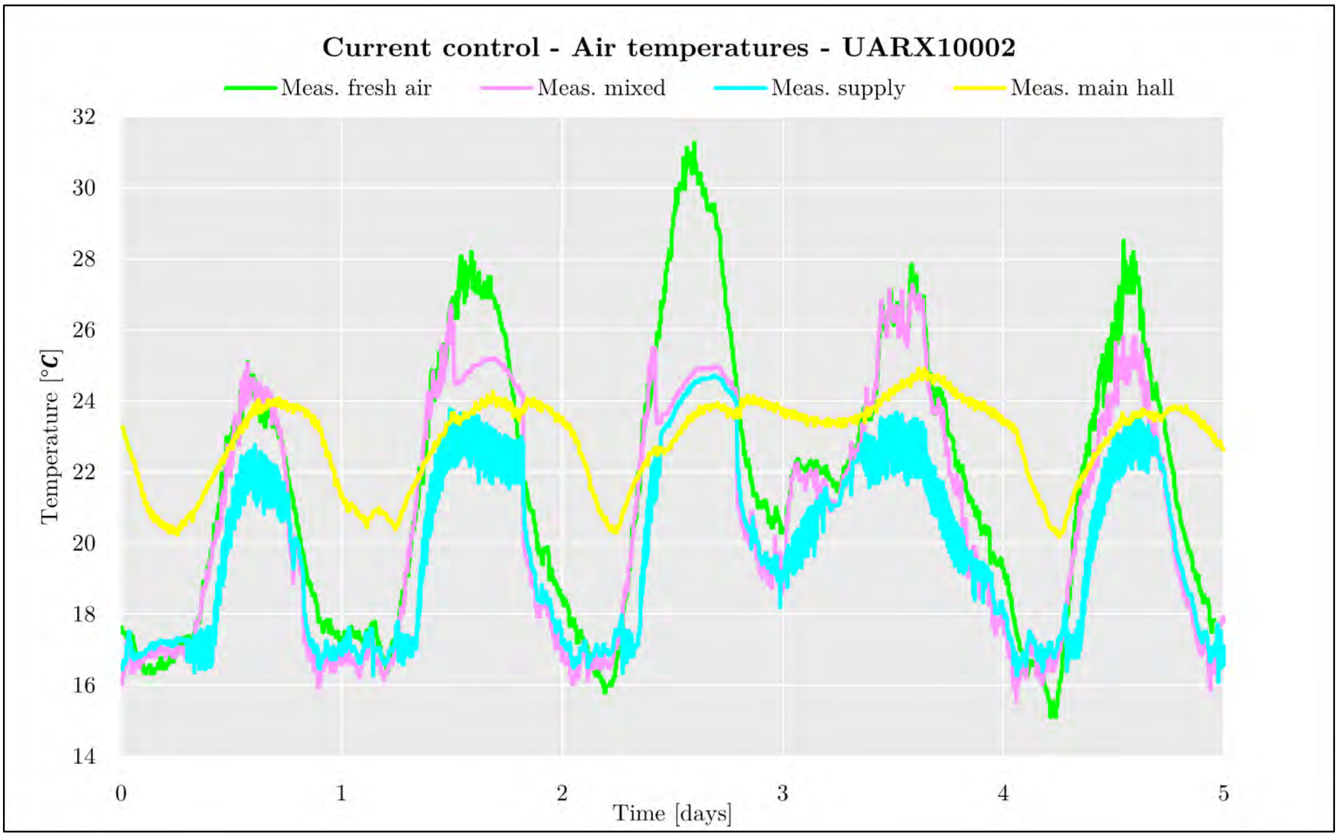

Table 10: PLCs running the model and the current control on temperature measurements for UARX10002 (Own archive, 2022).

(a)

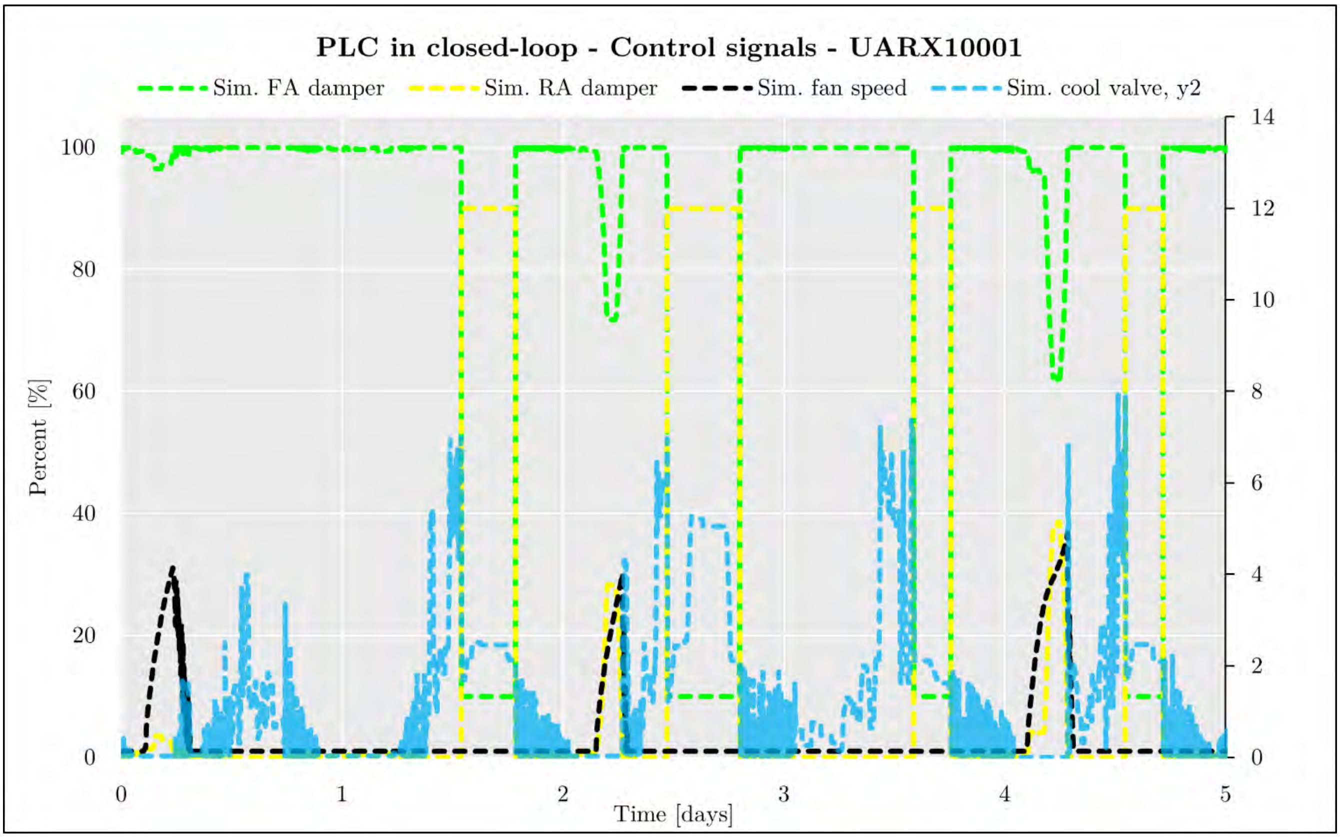

(b)

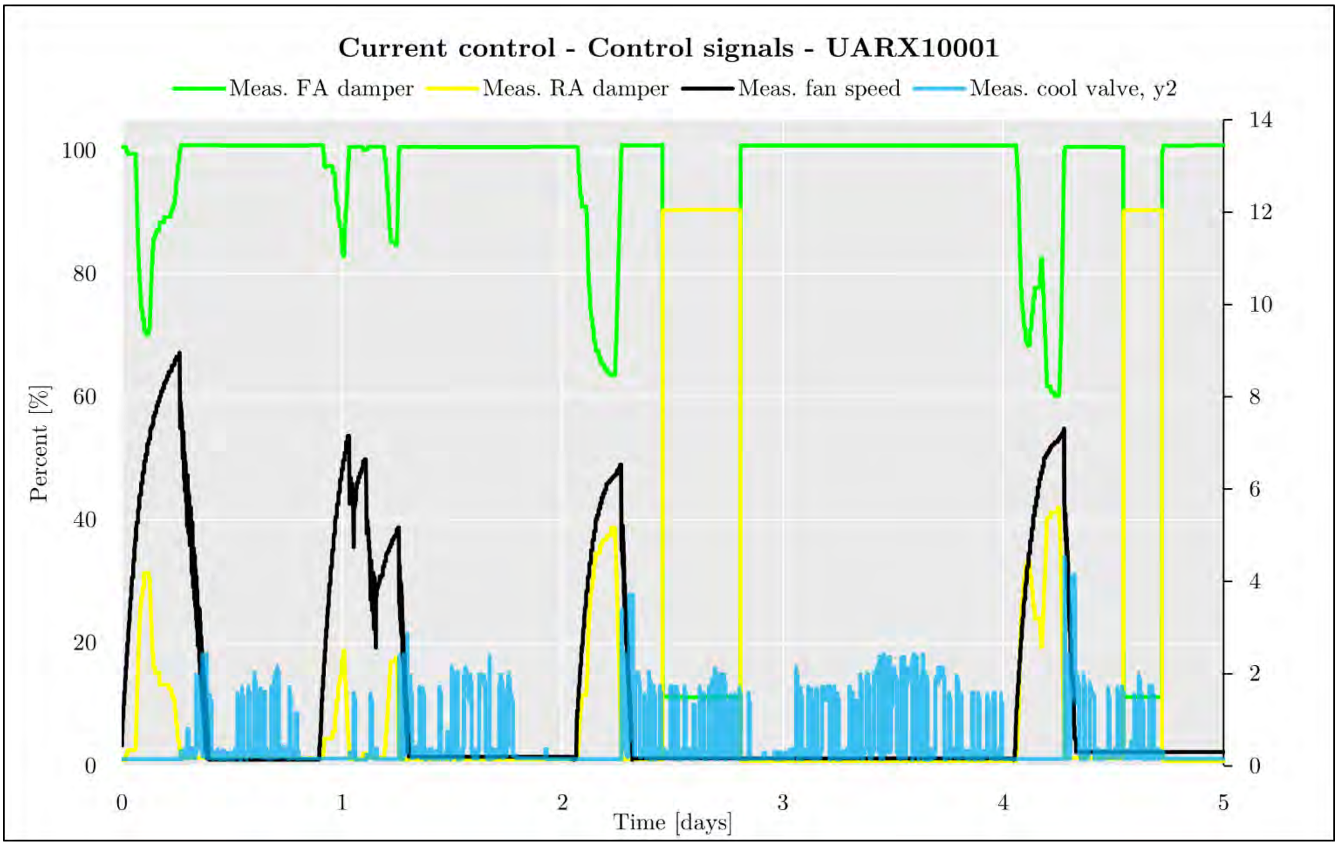

Table 11: PLCs running the model and the current control on control signals applied for UARX10001 (Own archive, 2022).

(a)

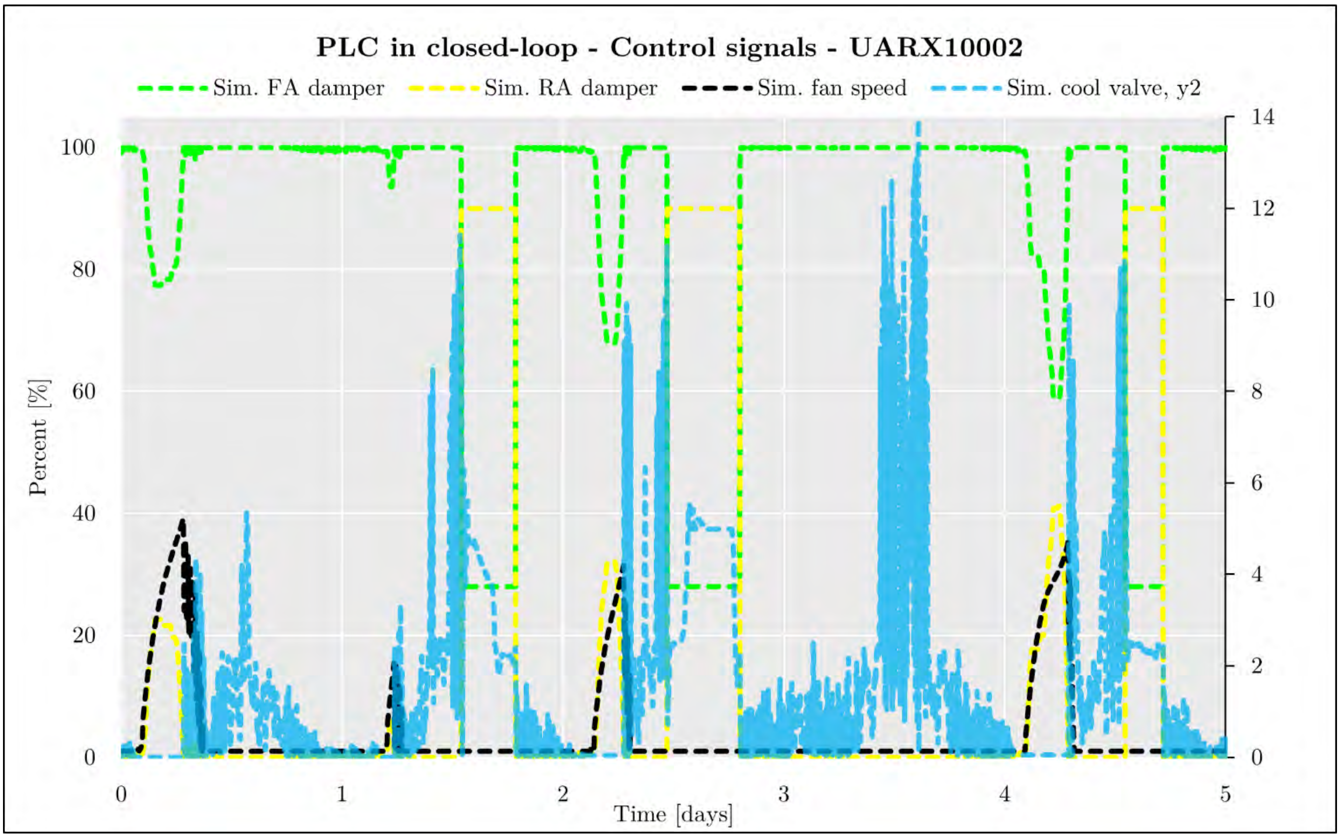

(b)

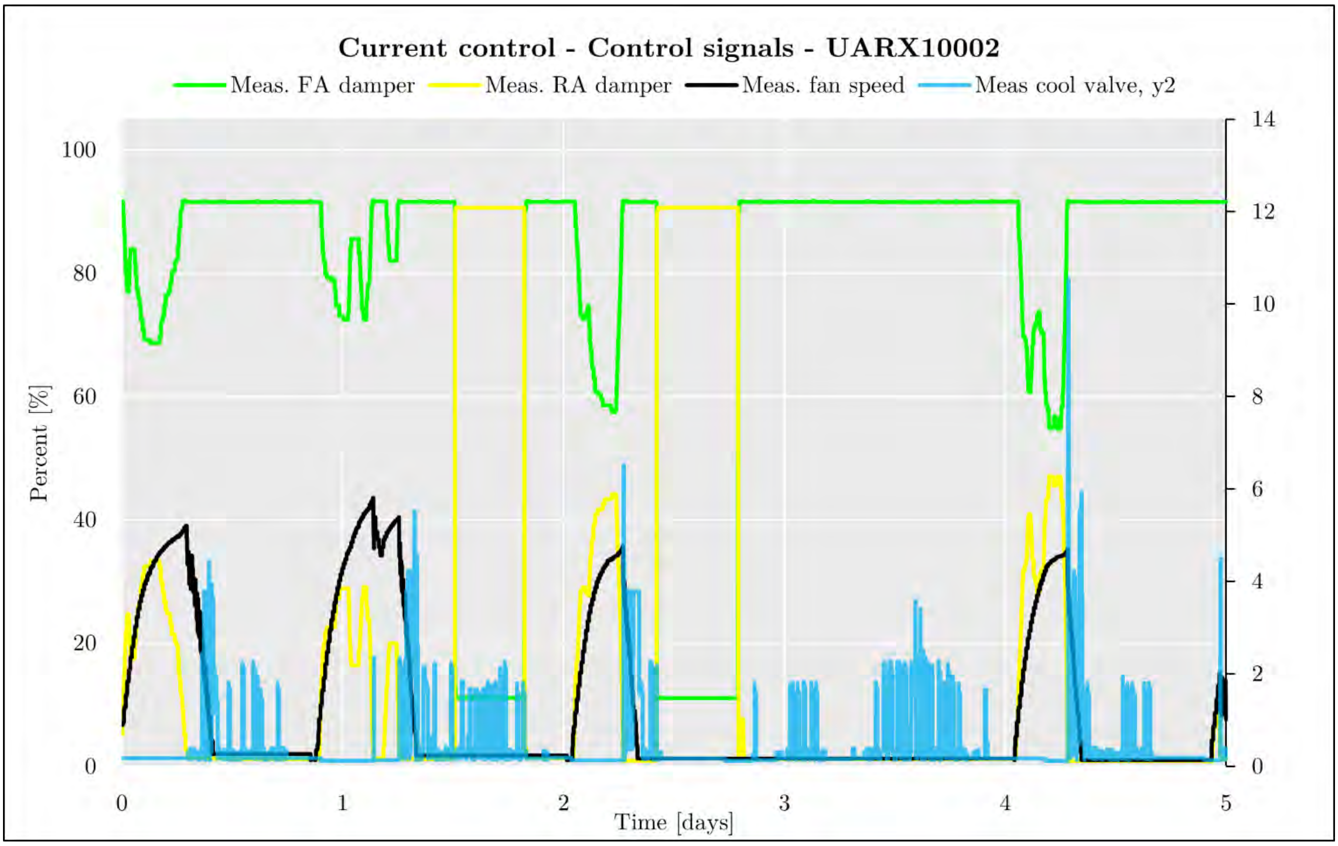

Table 12: PLCs running the model and the current control on control signals applied for UARX10002 (Own archive, 2022).

5.2.2 Interpretation of the closed-loop simulation

The closed-loop simulation results aimed to evaluate if the existing control configuration running on two PLCs could run the model and produce a good consistent regulation. The model differences from the actual plant were more likely to be expressed in differences in the control signals applied to the model from the PLCs.

Table 9 and Table 10 show the various temperature variables with (a) being the PLCs running the model in closed-loop and (b) the measurements from the plant in the same period, respectively, for UARX10001- and 10002. The most significant offset is seen on the main hall/recycled air temperature, with offsets up to 2 $°C$ in peak temperatures during the day. An offset of this caliper is to be expected given some of the limitations in the model; here are the two variables main hall- and recycled air temperature used in the regulation but are identical values, which is not the case on the actual plant. It will mean differences in applied control signals.

Compared to the measurements, a consistent dynamic profile can be observed with no discrepancies or offsets present. In the same plots, the supply air temperature can be seen, and oscillations are present throughout the five-day simulation. Already explained the reason for those, these are of less concern in the closed-loop simulation as the control system will react accordingly to provide the necessary amount of cooling. The oscillations were impossible to get completely rid of but are considered of less importance in else good and consistent simulation results.

It can be seen that the dynamic profile for the supply air temperature is consistent with the measurements and, apart from the oscillations, have no offsets or discrepancies.

The PLCs running the existing controls on the model produce consistent results on the supply air temperature and with a peak offset on the main hall/recycled air temperature Therefore, the main thing is to look at the applied control signals, as the closed-loop simulation will try to reach the setpoints and apply the needed control signals on actuators.

A good regulation can be the case but with a variance in the control signals compared to the measurements.

Table 11 and Table 12 shows the control signals applied by the PLCs on the model in (a) and the measurements in (b) for UARX-10001 and 10002, respectively.

In the cooling coil valve comparison, it is seen that they both introduce oscillatory behavior, but the simulated value has larger mean- and peak values. It was expected given the accepted limitation in the valve model, as earlier explained. Furthermore, knowing that the model's main hall and recycled air temperature are the same, differences were expected as the simulated value had higher peak values.

For UARX10001 can be seen the most significant difference. Over the five days, the fan has been ramping four times, and the model is missing one of these. Furthermore, the amplitudes are generally slightly lower. The operation of the fresh air damper reproduces similar behavior. However, with the model limitation mentioned above for the main hall/recycled air temperature, the controls introduce more frequent switching of damper positions. Else a good agreement is seen on the dynamics where dampers are modulating to keep the setpoint for the supply air temperature except between the zero-to-two-day mark, where sensor inconsistencies on the actual plant can be explained as a reason.

5.3 Applicability of the developed plant model

With the simulation results for model validation presented and interpreted in the former two sections, the applicability of the developed plant model can be discussed.

The main point of investigating the model is to assess its ability to replace the actual plant in a virtual environment. The model validation is a crucial step in determining whether the model can be used or not. The model's validity is essential if decisions and conclusions are to be made using the developed model.

A two-step validation study was deemed ideal for determining if the model can be considered validated or not. First, the open-loop simulation produced good results with an RMSE of ~1 °C and an MAPE of < 5 % in all cases on the four temperature variables used, which is believed near the model's potential. Initially, RMSE values of less than 1 °C were thought unrealistic, which then can be considered a satisfying result.

The open-loop simulation is the best measure of model precision, and the result from the simulation is satisfactory.

The closed-loop stimulation on top of the closed-loop simulation aimed to cement the model's ability to reproduce actual plant behavior with two different approaches, adding valuable information to the model state.

The results from the closed-loop simulation were better than expected, with the PLCs running the existing controls being able to achieve a good regulation with some offsets accounted for in the former subsection. The control signals applied are where the most significant discrepancies and offsets could be expected, but they were minimal. The dynamics were captured well, and some model limitations were accounted for, resulting in differences in the damper operation and offsets on the fan.

Validation

The model can be considered valid but needs to be addressed in some areas based on the simulation results. As argued in **1.4**, the model was parametrized to fit a summer scenario with expected active cooling.

An effort to ensure the model applicability during an entire year was considered but eventually delimited for the short time frame of the project. It is because two or more sets of model parameters can be necessary to account for operational differences caused by seasonal changes in temperature, which are characteristic.

Therefore, with the model's current state, validation will also have to be performed using different datasets, e.g., winter and mid-season, to fully validate the model. It will likely mean that further effort needs to be put into the model with parameters fitted to data from other seasons.

For the validation performed, a five-day dataset was used. The dataset was extracted outside of the data used to parametrize the model, but as the summer season data is similar, one must be more critical of the results.

The MPC simulations only ran five times faster than in real-time, and the PLC closed-loop simulation was initially simulated in real-time before a solution to speed it up to 10 times real-time was successfully implemented.

Due to these simulation time concerns, a relatively short time frame of five days was simulated.

More extended simulations with different datasets still applied for a summer scenario should be considered. In addition, more validation simulations will potentially exploit errors or mismatches in the model, adding to a less valid model, which must be considered when using the model.

Lastly, a further point related to the plant model validation and applicability is the MPC model that forecasts system behavior. It is different from the EcosimPro model developed in the thesis but also derived from operational data. That model can be considered low-fidelity based on relatively few sets of differential equations compared to the EcosimPro model with more than $1'000$ equations. However, with the review in the following subsection, a good correlation is the case with the MPC able to utilize its model on the EcosimPro plant model, which is a further point to add to the validation.

5.4 Virtual commissioning simulation

With the plant model validated for a summer scenario, tests and studies were carried out in the Ph.D. project. First, specific scenarios were designed and then simulated to evaluate whether the MPC controller behaved as expected or if corrections were needed. Eventually, a full five-day virtual commissioning simulation was performed using the MPC as control. It enables a thorough look at how the MPC would behave in a realistic scenario with data available to compare with the existing controller.

The simulation results are visualized and presented in between **Table 14** through **Table 17** on pages **68**-**71**.

In subsection **5.4.2**, a brief review of the simulation results will be performed. However, as the MPC control development is not part of the thesis but merely the ultimate purpose of using the developed plant model for validating the MPC controller, its performance will not be evaluated. Instead, the review will focus on some of the identified problems, how the MPC can solve them, and at the same time, utilize a more energy-efficient operation.

The constraints are significant on the MPC, and the current control also has constraints on the setpoint calculation for the main hall- and supply air temperature loops. The following table will explain them as they are used in the plots.

Temperature constraints	Value [°C]	Indicators
Supply air maximum	30 °C	
Main hall air maximum	26 °C	
Main hall air ideal setpoint	20 °C	
Supply- and room air minimum	17 °C	
Supply air minimum (only for MPC)	14 °C	

Table 13: Temperature constraints on both the MPC controller and current control (Own archive, 2022).

5.4.1 Plots

See **Table 14** through **Table 17** on pages **68**-**71**.

(a)

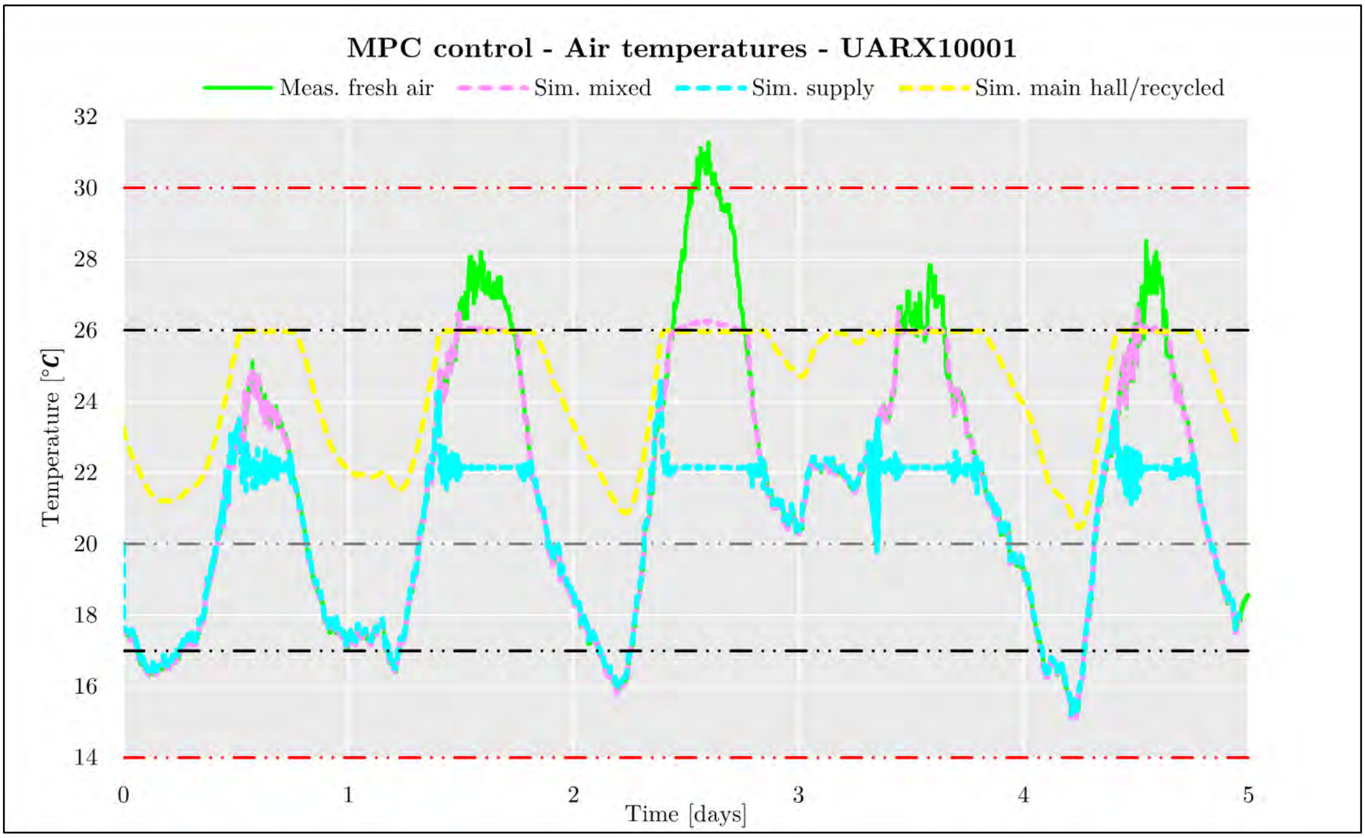

(b)

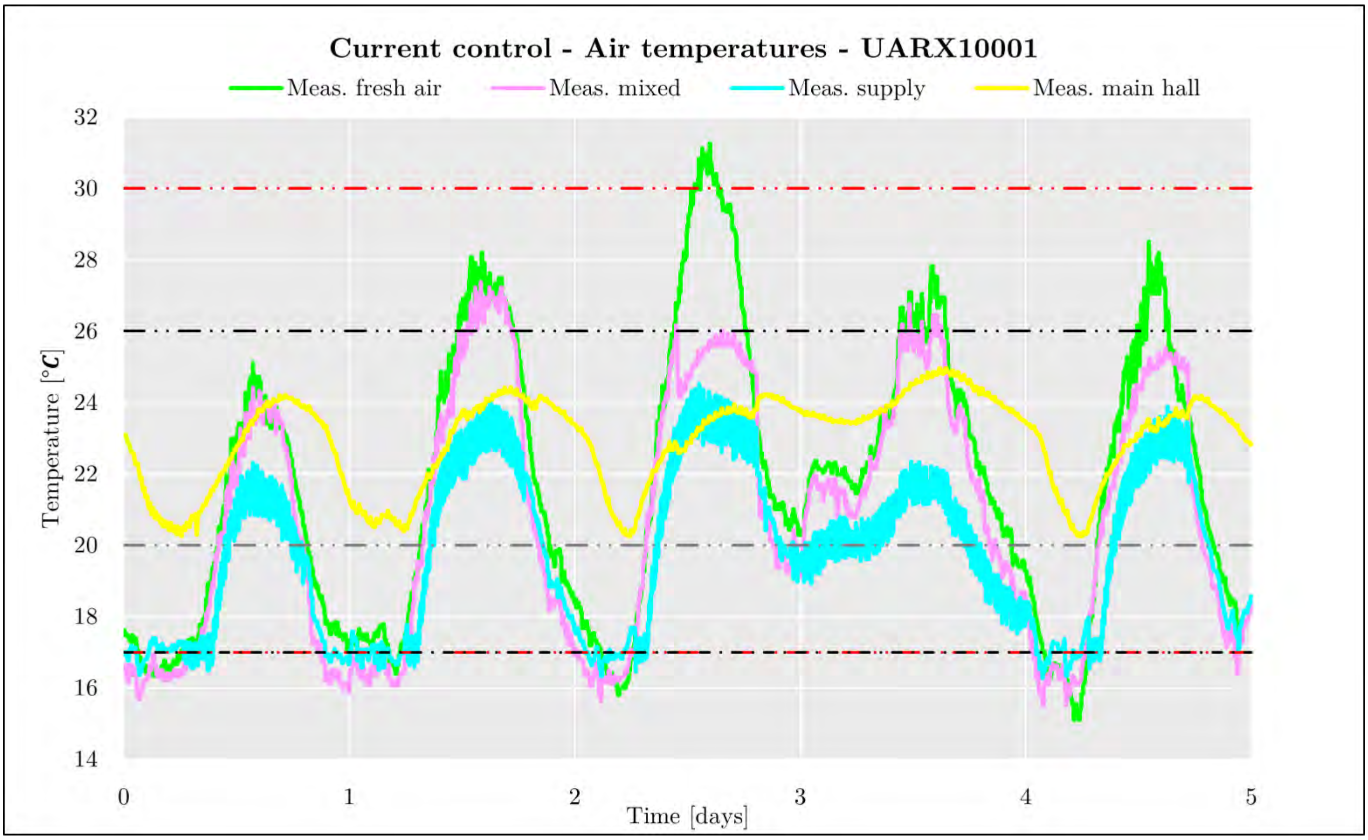

Table 14: Current control compared to MPC control on four different temperature measures for UARX10001 (Own archive, 2022).

(a)

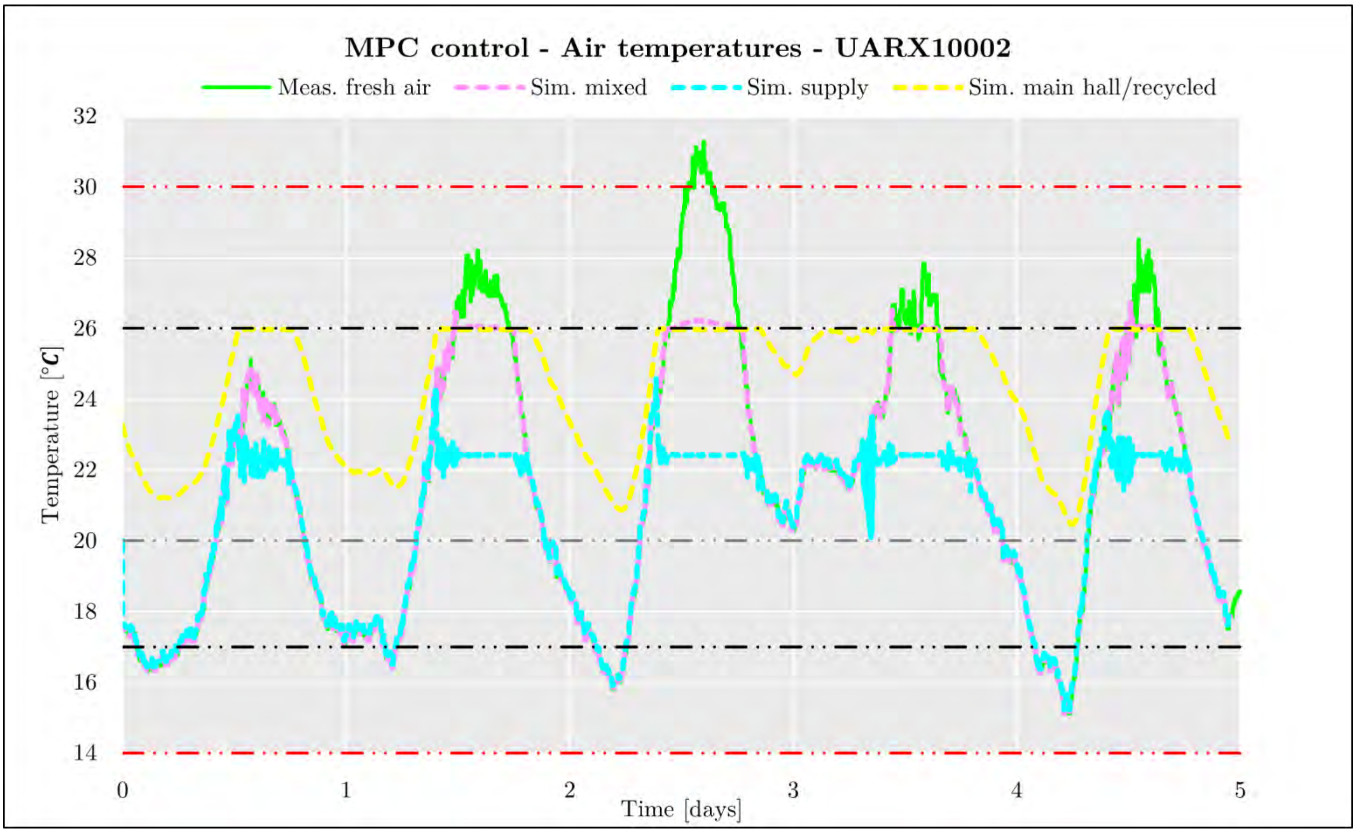

(b)

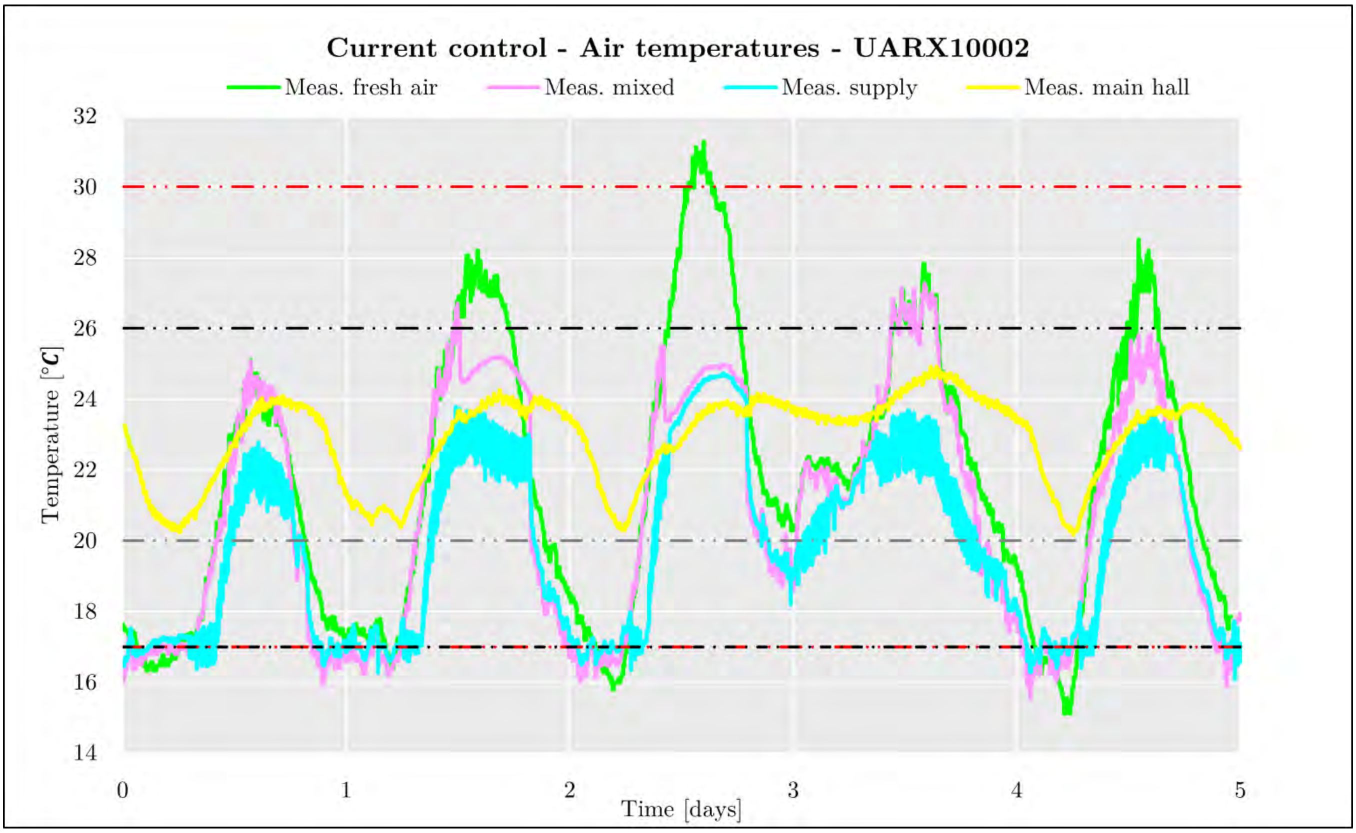

Table 15: Current control compared to MPC control on four different temperature measures for UARX10002 (Own archive, 2022).

(a)

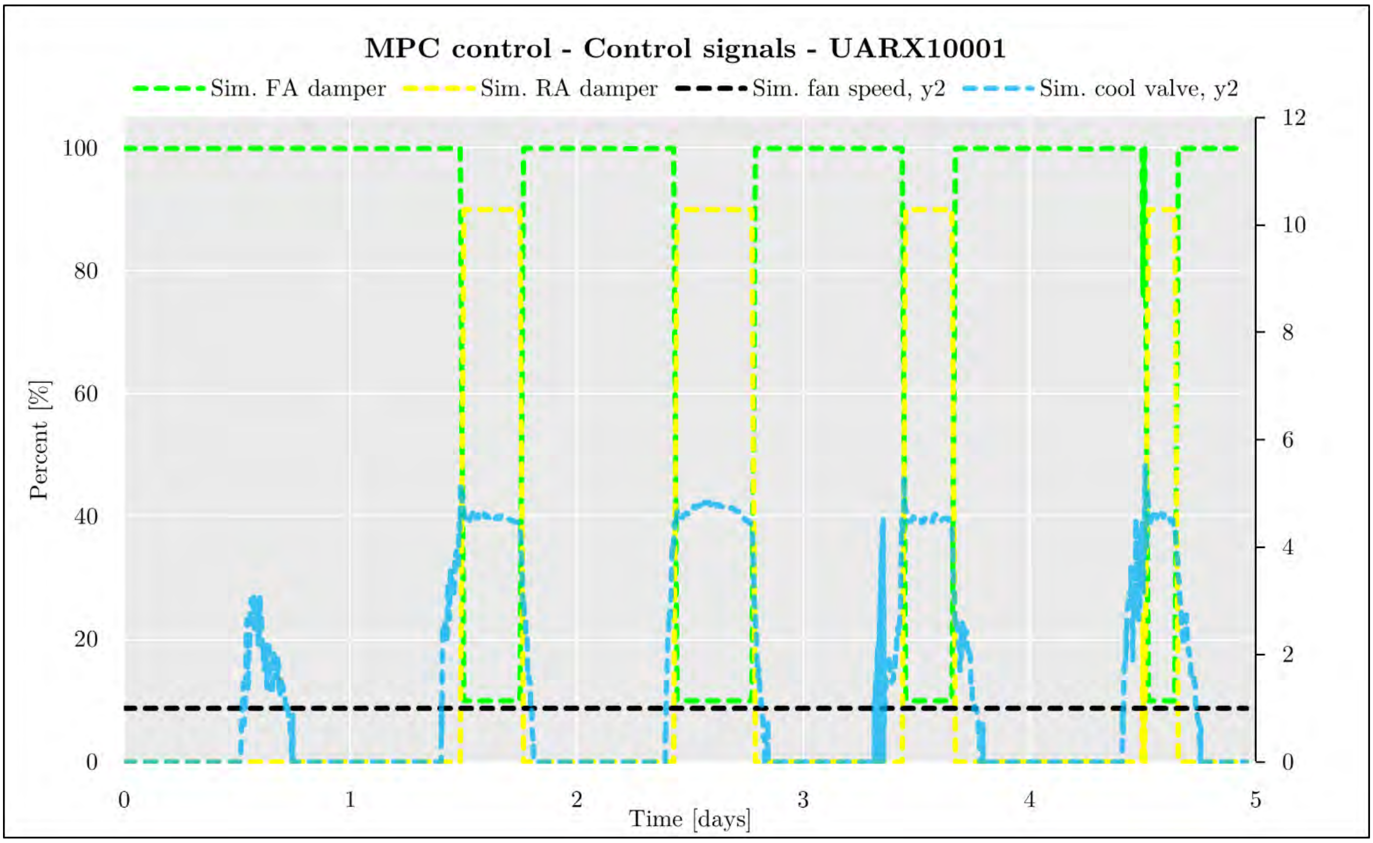

(b)

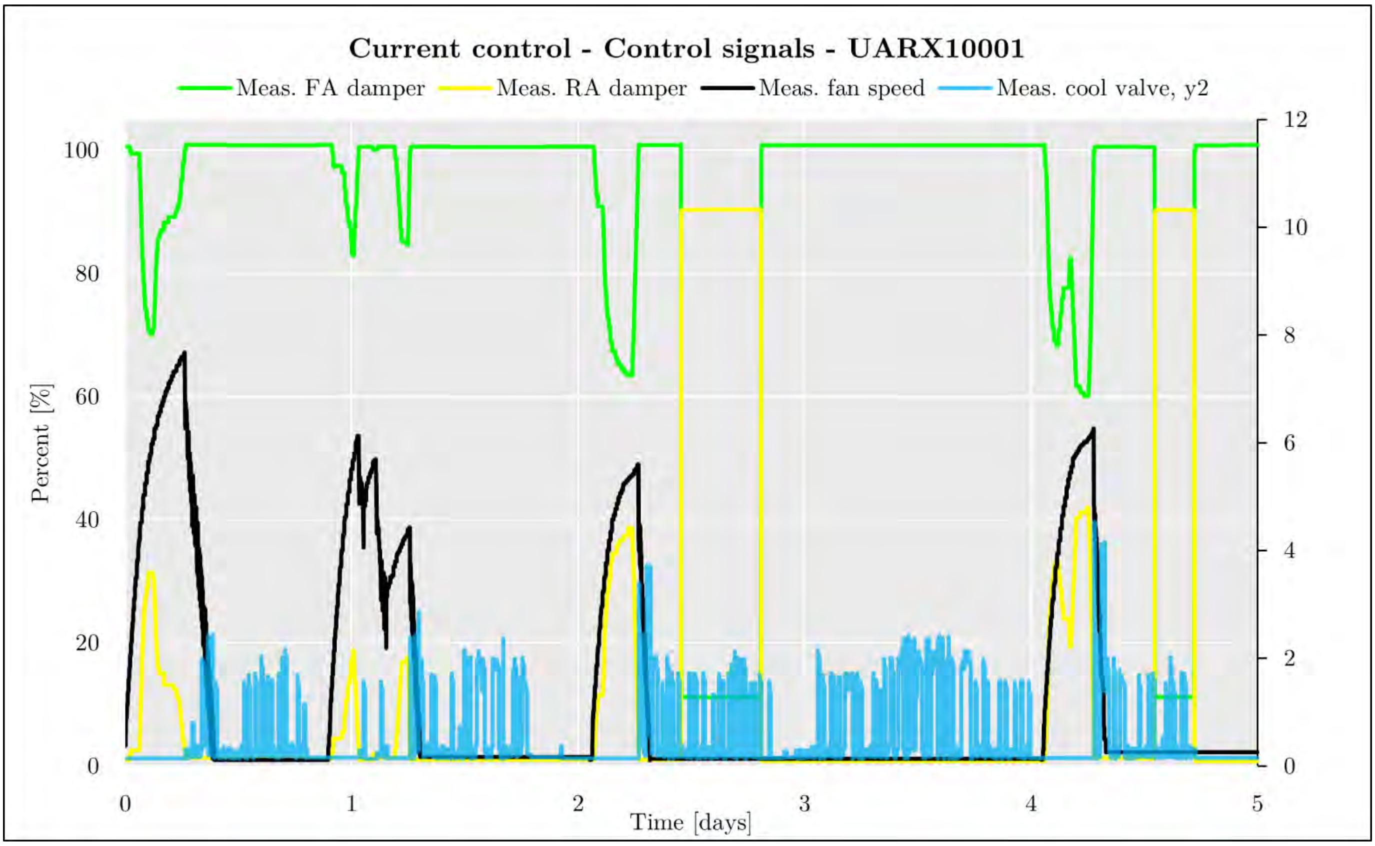

Table 16: Current control compared to MPC control on the applied control signal for UARX10001 (Own archive, 2022).

(a)

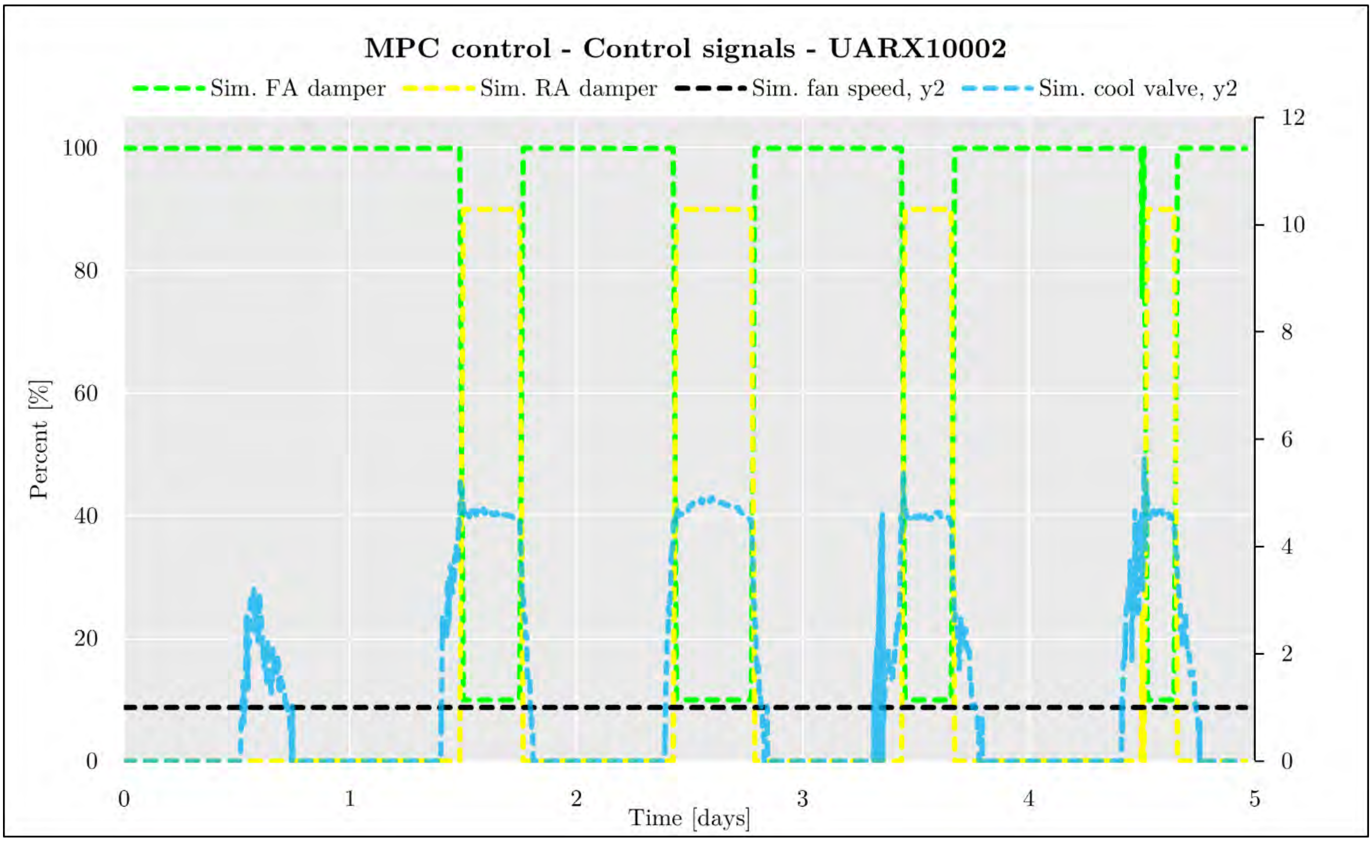

(b)

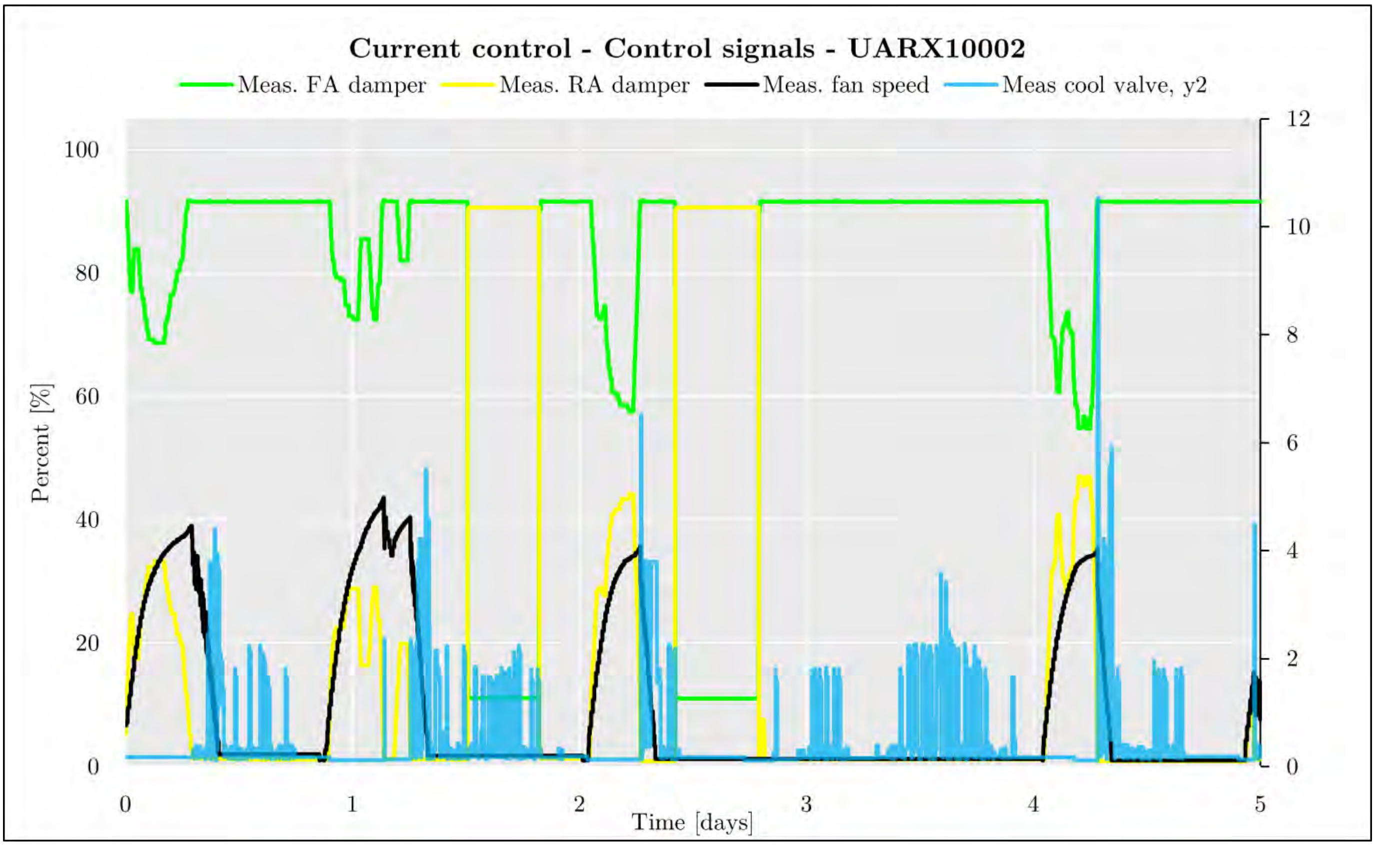

Table 17: Current control compared to MPC control on the applied control signal for UARX10002 (Own archive, 2022).

5.4.2 Review on the virtual commissioning simulation results

As explained in section 5.1 for the model outputs, the recycled and main hall air temperatures are the same. In the open-loop simulation, it can have minor or no impact at all indifferent from this simulation.

The recycled air temperature is permanently higher in the order of ~1 $°C$, but not entirely consistent as the value deviates to lower and even higher values. It affects the MPC controller as it uses the recycled air temperature input as if it was the main hall temperature, thus making a direct comparison slightly biased. Furthermore, because the model cannot reflect different zones horizontally- and vertically, the consequence is that the local evolution of temperature present in these zones, which is used in the regulation, will not be different. Therefore, the same control signals are applied to the two AHUs. Because there is a slight difference in the parameters of the two, it can be observed in the various temperature variables because there is a slight difference. The UARX10002 has a slightly higher airflow rate and a slightly more considerable influence over the other AHU.

The MPC development reached a state where a first controller version was considered done. It aimed to fix some known problems of the current control, introduce more stability, and, most importantly, utilize a more energy-efficient operation.

Some known problems:

- Undesirable coupling between the calculated setpoint and fan speed
- Significant oscillations in the valve control signal
- Conservative lower bound constraint on supply air temperature
- Unrealized energy optimization potential

In **Table 14** and **Table 15 (a)**, **(b)**, the MPC- and current control both satisfy the main hall air temperature bounds. However, it is essential to note that the main hall air temperature is slightly higher for the MPC control than realistically because the recycled/main hall air temperature is identical in the model and is parametrized using the recycled air temperature as a reference. In this case, the controlled variable, the main hall air temperature, is slightly higher as it is identical to the recycled air temperature. With that noted, it can still be observed that the main hall air temperature with the MPC as control tends to be closer to the upper limit bound compared to the current control, which is the main point to observe in the scenario.

In **Table 16** and **Table 17 (a)**, **(b)** both can be seen that while the current control ramps up the fan for a large portion of a day caused by the coupling between calculated setpoint and fan speed, the MPC control does not. Observed when the fan is ramping, the cooling coil valve is also operating. That being the case, even though the main hall air temperature is well under the upper bound. From the energy-optimal perspective, the plant should ideally spend as much time in free cooling mode, which is not the case with the current control.

The MPC controller can be seen over the full five days to not ramp up the fan and only start operating the cooling coil valve as the main hall air temperature upper bound can no longer be satisfied. Furthermore, looking at the cooling coil valve operating with MPC as control, it is seen that significantly less actuation is applied, which is ideal as components wear down less.

The valve openings can be seen is more consistent but also have higher values. With the known limitation accepted in the model that the valve model could not account for the large cooling effect at small openings as the flow characteristics were not entirely correct, the openings are presumably larger than would be the case if applied to the actual plant.

The MPC can be seen to utilize the energy-optimal operation in free cooling mode until it no longer can satisfy the bounds. Whenever, e.g., active cooling is required, the main hall air temperature should be precisely on the upper limit bound as it is the most efficient, which is not the case with the current control but can easily be seen as the case with the MPC control. At the same time, the current control both uses the cooling coil and the fan, which can be considered over actuating as the main hall air temperature is well below the upper limit bound. In between the three- and four-day mark, the MPC control can be seen to introduce some oscillatory behavior of the cooling coil valve. As MPC is an optimal-control methodology, an optimal control strategy can be oscillatory, which was deemed unwanted, so measures to make changes to prevent it are currently being investigated.

The lower bound constraint on supply air temperature is set conservatively at 17 $°C$ in the current control. It is mainly to prevent condensation, but as the dewpoint temperature changes throughout the day, month, and year an approach to dynamically calculate it and set the lower bound to a conservative value above was proposed as a solution.

In this exact MPC control simulation, the value is set fixed at 14 $°C$. It is interesting from the energy-optimal perspective as it is seen in the current control that because of this minimum setpoint, the recycled air dampers are modulating when the fresh air temperature goes below 17 $°C$ to keep the minimum setpoint. At the same time, the main hall air temperature is well situated in the middle of the upper- and lower bounds. Therefore, the main hall air temperature can reach lower values by allowing a lower supply air temperature, thus having a lower initial starting point at the beginning of the next day. As a result, the outside air temperature rises once again, consequently allowing the plant to operate in free cooling mode for a longer period before entering active cooling mode, hence ensuring a more energy-efficient operation.

Currently, the Cooling and Ventilation group assumes that the most energy-efficient operation, e.g., in active cooling mode, is modulating the cooling coil valve at minimum fan speed. However, there is no formal proof of this. Therefore, the balance of contribution from the fan and cooling coil is currently being investigated, potentially exploiting an energy-optimal split contribution of the two.

Summary

With the performed virtual commissioning simulation on the validated plant model, the known problems aimed to eliminate with the MPC showed promising results. It was done using the plant model and can realistically show expected results if implemented on the actual plant. With data available for the same period, with the same dataset used (fresh air temperature), an almost direct comparison can be made to the existing controls and is, therefore, a solid foundation from which future decisions can be made.

6

Conclusion

This thesis addresses the research question: ***"How is it possible to create a high-fidelity digital twin of a cooling system for use in advanced controls?"*** and related objective, which ultimately was to provide a model of the SR1 pilot plant, validate it, and enable the possibility to perform virtual commissioning.

At first, it was necessary to examine previous cases where experts' have done similar projects in the past. It provided a great insight into the working principles needed and emphasized the need to include all accessible data. EcosimPro and MPC were presented as EcosimPro was the primary tool used. MPC was the process control methodology with a controller developed in the Ph.D. project to be validated using the digital twin performing virtual commissioning.

The methodology and data used were assembled and put into practice. Most of the project work was in the iterative process of running multiple experiments and estimating model parameters based on the investigations into the experiment results. Therefore, it was essential to access and process historic control system data to estimate parameters after using all available as-built documentation to learn about the plant and set initial model parameters.

The SR1 plant was investigated as an instrumental basis for conducting the work. It was shown how it was possible to build and drive the model in EcosimPro to perform simulations. With the monitor developed, an efficient platform to conduct the work and investigate the experiment results was in place.

Several essential aspects of the concrete model development work were shown how it was possible to parametrize different components for all of the work done. Model limitations were addressed and turned into possibilities, and instead, implementing solutions that were better than those if no action was taken. For example, deficiencies in the heat exchanger model were found, which proved problematic. However, a solution limiting the negatives followed using the model could be implemented to damage-limit the problems.

With the model development naturally coming to a conclusion, an important next step was instrumental. Model validation. It was possible to present good results from both open-loop and closed-loop simulations with a high-degree of fidelity and an interpretation of results and discussions on them. The model is considered valid under certain conditions, with input data corresponding to a summer scenario. Further potential problems concerning the model validity are addressed, which is also vital to know when using the model.

Ultimately, the model was provided to the Ph.D. project, where good results were achieved, some included in this thesis. In addition, the model and its results are already being used in advanced control studies and virtual commissioning of the MPC to show its ability to eliminate known problems compared to the current controller on the SR1 plant while also utilizing a more energy-efficient operation.

It can be definitively concluded that this thesis shows how it was possible to create a high-fidelity digital twin of a cooling system for use in advanced controls.

6.1 Perspectives & future work

HVAC libraries in EcosimPro

A side project was initiated alongside and in parallel with the work presented in this thesis. Empresarios Agrupados, the creator of EcosimPro, has made a HVAC and buildings library commercially available. The Industrial Control Systems group's interest is to evaluate and compare it to the internally developed HVAC library already in use. Some problems, deficiencies, and irregularities were discovered during the model development with the existing CERN HVAC library. The evaluation was not concluded partly due to the problems mentioned above, which consumed much time to find viable solutions but did not fix the core problem. Therefore, it is recommended to continue the evaluation and consider the discoveries made in the existing library to decide which library to use in the future eventually.

Diagnostics module for AHUs

With a substantial amount of effort put into modeling the plant, validating it, and using it for its purpose, future perspectives could be subject to exploitation. For example, a side product of the modeling effort caused by the need for in-depth knowledge about the plant led to discovering a leaking recycled air damper on another plant. The control system can cope with anomalous behavior, and therefore, it can easily go unnoticed and still reach the control objective. However, from the energy-optimal perspective, it has negative consequences. The scenario mentioned above would cause the cooling coil or fan to operate more than necessary and waste energy. A diagnostics module for the AHU can potentially indicate if a fault is present.

A model of each possible fault in normal operating conditions is necessary and can be derived from operational data. A significant deviation from the model output would indicate anomalous behavior and possibly be caused by a fault. The effort into creating such a module is relatively easy, but the potential benefits could be more prominent. The models will run parallel with the existing system and require no maintenance. However, they will indicate anomalous behavior if detected, enabling the people in charge to decide if an investigation into the matter should be carried out or not.