Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Federal Department of the Environment, Transport,
Energy and Communications DETEC

Swiss Federal Office of Energy SFOE
Energy Research and Cleantech Division

Final report dated December 22, 2022

ISOTHERM

Intelligent Software for Thermal Energy Management

Source: © Carrosserie HESS AG, 2017

Date: December 22, 2022

Subsidy recipients:
ETH Zurich
Institute for Dynamic Systems and Control (IDSC)
ML K 32.3, Sonneggstrasse 3, CH-8092 Zürich
https://www.idsc.ethz.ch

Carrosserie HESS AG
Bielstrasse 7, CH-4512 Bellach
https://www.hess-ag.ch

Authors:
Fabio Widmer, ETH Zurich, fawidmer@idsc.mavt.ethz.ch
Prof. Dr. Christopher Onder, ETH Zurich, onder@idsc.mavt.ethz.ch

Nicolas Amacker, Carrosserie HESS AG, nicolas.amacker@hess-ag.ch
Alexander Böhm, Carrosserie HESS AG, alexander.boehm@hess-ag.ch

SFOE project coordinator:
Dr. Luca Castiglioni, luca.castiglioni@bfe.admin.ch

SFOE contract number: SI/501979-01

Zusammenfassung

Dieser Abschlussbericht hält die wichtigsten Erkenntnisse des Forschungsprojekts ISOTHERM fest. In diesem Projekt werden die thermischen Systeme von Fahrzeugen im elektrifizierten öffentlichen Verkehr im Detail analysiert. Dazu wird ein modellbasierter Ansatz verfolgt. Die entwickelten Modelle werden auf modulare Weise kombiniert, was die Analyse verschiedener Fahrzeugtypen und verschiedener Systemkonfigurationen ermöglicht. Ziel ist es, das Systemverständnis so zu verbessern, dass die inhärenten Zielkonflikte zwischen Energieeffizienz, thermischem Komfort und Batteriedegradation einerseits durch die Optimierung des Designs (d.h. der Hardware) des Systems für zukünftige Fahrzeuge und andererseits durch die Optimierung der Steuerung (d.h. des Betriebs durch die Software) eines gegebenen Systems in Echtzeit besser gelöst werden können. Wir erachten die folgenden drei Punkte als die wichtigsten Erkenntnisse.

Erstens werden für die Entwicklung der Modellstruktur die Paradigmen der “graphbasierten Modellierung” angewandt, um die Modularität der Modelle zu sicherzustellen. Die entwickelte Modellstruktur bietet viele nützliche Funktionen, wie die Visualisierung des Modells in Form eines gerichteten Graphen, die Bestimmung der Anzahl der Freiheitsgrade, die Simulation mit kausalen Reglern und die automatische Formulierung und Lösung von Problemen der optimalen Steuerung. Diese Modellstruktur bildet somit eine ideale Grundlage für die weitere Arbeit im nachfolgenden Projekt Swiss eBus Plus.

Zweitens wird eine eingehende Analyse des Energieverbrauchs verschiedener Fahrzeuge des öffentlichen Verkehrs unter der Annahme eines stationären Zustands durchgeführt. Dieser rechnerisch effiziente Ansatz ermöglicht es, mehrere tausend Szenarien unter verschiedenen Umgebungsbedingungen zu betrachten, und somit die ganzjährige Leistung des Systems bei unterschiedlichen Komfortanforderungen abzuschätzen. Den Kompromiss zwischen dem erreichten thermischen Komfort und dem Energieverbrauch können wir schliesslich visualisieren, indem wir beide Ziele in einer Pareto-Front darstellen. Der Vergleich dieser Ergebnisse für verschiedene Konfigurationen zeigt, dass Wärmepumpen mit einem Energieeinsparungspotenzial von 40–60% die vielversprechendste Verbesserung gegenüber Systemen auf der Basis von rein elektrischen Heizungen darstellen. Luftschleier können den Energieverbrauch in Bussen um etwa 25% senken, während der Effekt bei Zügen weniger ausgeprägt ist. Auf der anderen Seite kann eine um 20% erhöhte Isolation den jährlichen Energieverbrauch in Zügen um etwa 20% senken, während sie in Bussen weniger wirksam ist. Heizstrahler können in Bussen ohne Wärmepumpen zu einer Verringerung des Energieverbrauchs um 6–7% führen. Dieses Potenzial verschwindet jedoch, wenn eine Wärmepumpe installiert ist. Bei anderen Fahrzeugtypen bieten Heizstrahler kein Einsparungspotenzial.

Drittens zeigen die Untersuchungen im Rahmen dieses Projekts den erheblichen Einfluss des Wärme- und Traktionsenergiemanagements auf die Batteriealterung. Insbesondere wird das Potenzial der Einbeziehung eines thermischen Puffers in das Traktionsenergiemanagement gründlich analysiert. Dabei zeigt sich, dass es wichtig ist, eine hohe Pufferleistung von mindestens $50\,\mathrm{kW}$ anzubieten, um das volle Potenzial des Puffers auszuschöpfen. Andererseits muss die Energiekapazität des Wärmespeichers nicht allzu gross sein: Bereits $500\,\mathrm{W\,h}$ reichen aus. Basierend auf diesen Erkenntnissen wird eine intelligente Heizstrategie auf Grundlage des Pontrjaginschen Minimumprinzips für das Warmwassersystem entwickelt. Diese Strategie wird auf einem von den VBZ betriebenen Trolleybus-Prototypen umgesetzt, der über ein steuerbares Widerstandsheizelement verfügt. Nach über 1000 Betriebsstunden zeigt eine umfangreiche Datenanalyse, dass die neue Strategie effektiv in der Lage ist, die Batteriealterung um etwa 12% zu reduzieren. Schliesslich wird ein optimiertes Traktionsenergiemanagement vorgeschlagen, das auf periodischen Messungen der Batteriegesundheit basiert. Der dreistufige Regelalgorithmus kann diese Messungen nutzen, um eine bestimmte Lebensdauer der Batterie zu gewährleisten, selbst wenn die internen Modelle nicht perfekt sind oder nicht gemessene Störungen auftreten.

Summary

This final report records the most important findings of the research project ISOTHERM. In this project, the thermal systems of vehicles in electrified public transport are analyzed in detail. Therefore, a model-based approach is followed. The developed models are combined in a modular fashion allowing the analysis of different vehicle types and various systems configurations. The goal is to improve the understanding of the systems such that the inherent trade-offs between energy efficiency, thermal comfort, and battery degradation can be resolved more effectively, on the one hand, by optimizing the design (i.e., the hardware) of the system for future vehicles, and, on the other hand, by optimizing the control (i.e., the manner of operation through software) of the a given system in real time. We consider the following three points to be the main contributions.

First, for the development of the modeling framework, the paradigms of “graph-based modeling” are adopted to ensure the modularity of the models “by design”. The developed framework offers many useful features, such as model graph visualization, determination of the number of degrees of freedom, simulation using causal controllers and automatic formulation and solution of optimal control problems (OCPs). This framework hence forms an ideal basis for the continued work in the subsequent project Swiss eBus Plus.

Second, an in-depth analysis of the energy consumption of different public transport vehicles is conducted based on a steady-state assumption. This computationally efficient approach allows to evaluate several thousand scenarios in various ambient conditions to estimate the year-round performance of the system subject to different comfort requirements. Thus, we can quantify the trade-off between the achieved thermal comfort and the energy consumption by visualizing both objectives in a Pareto front. Comparing these results for various configurations, we show that heat pumps are the most promising improvement over systems based on positive temperature coefficient (PTC) heaters, reducing the energy consumption by 40–60%. Air curtains can reduce the energy consumption in buses by about 25%, while the effect is less pronounced for trains. On the other hand, insulation improvements by 20% can reduce the annual energy consumption in trains by about 20%, while being less effective in buses. Radiant heaters can lead to a reduction in energy consumption of 6–7% in buses that are heated with PTC heaters. However, this potential vanishes when a heat pump is installed. For other types of vehicles, radiant heaters offer no savings potential.

Third, the research conducted in this project reveals the significant influence of the thermal and traction energy management on battery aging. The potential of including a thermal buffer into the traction energy management is analyzed thoroughly, revealing the fact that it is important to offer a high buffering power of at least $50\,\mathrm{kW}$ to harness the buffer’s full potential. On the other hand, the energy capacity does not have to be large: $500\,\mathrm{W\,h}$ is already sufficient. On the basis of these findings, an intelligent heating strategy is developed based on Pontryagin’s minimum principle (PMP). This strategy is implemented on a prototype trolley bus, which features a controllable resistive heating element. After over $1000\,\mathrm{h}$ of operation in the public transport grid of Verkehrsbetriebe Zürich (VBZ), extensive data analysis suggests that it is indeed able to significantly reduce battery degradation by about 12%. Finally, an optimized traction energy management system (EMS) is suggested based on periodic state of health (SOH) measurements. This three-level control algorithm can use these measurements to ensure a certain lifetime of the battery even if its internal models are not perfect or if unmeasured disturbances occur.

Contents

Abbreviations

AC	air-conditioning
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning
BEV	battery electric vehicle
BFH	Bern University of Applied Sciences
BT	Bachelor's thesis
BVB	Basler Verkehrs-Betriebe
CAN	controller area network
CFD	computational fluid dynamics
COP	coefficient of performance
CSFM	Center for Sustainable Future Mobility
DC	direct current
DHI	diffuse horizontal irradiance
DNI	direct normal irradiance
DOD	depth of discharge
DOF	degree of freedom
ECMS	equivalent consumption minimization strategy
ECTS	European Credit Transfer and Accumulation System
EMS	energy management system
EOL	end-of-life
FD	final drive
FVM	finite volume method
GNSS	global navigation satellite system
GRU	gated recurrent unit
GVM	gross vehicle mass
GWP	greenhouse warming potential
HEV	hybrid electric vehicle
HP	heat pump
HVAC	heating, ventilation, and air-conditioning
IDSC	Institute for Dynamic Systems and Control
IFAC	International Federation of Automatic Control
ISO	International Organization for Standardization
ISOTHERM	Intelligent Software for Thermal Energy Management
IWB	Industrielle Werke Basel
LTO	lithium-titanate-oxide
LVD	longitudinal vehicle dynamics
MT	Master's thesis

NLP	nonlinear program
NTU	number of transfer units
OCP	optimal control problem
ODE	ordinary differential equation
PCM	phase change material
PI	proportional-integral
PMP	Pontryagin's minimum principle
PMV	predicted mean vote
PPD	predicted percentage dissatisfied
PTC	positive temperature coefficient
PWM	pulse-width modulation
SFOE	Swiss Federal Office of Energy
SIL	software in the loop
SOC	state of charge
SOH	state of health
SOM	studies on mechatronics
SP	semester project
SSH	secure shell
TCP	transmission control protocol
UMAP	uniform manifold approximation and projection
VBZ	Verkehrsbetriebe Zürich
VCC	vapor compression cycle
VCU	vehicle control unit
WP	work package

0 Introduction

0.1 Motivation

In public transportation, the heating, ventilation, and air-conditioning (HVAC) systems make up for a large portion of the overall energy demand. This is due to multiple reasons, e.g., the large number of passengers or the frequently opening large doors. For electrically driven public transport, this is even more pronounced, as these vehicles do not have waste heat abundantly available, as e.g. combustion-based vehicles do. Additionally, range considerations are a key aspect when investigating battery electric vehicles (BEVs). Therefore, not only ecological, but also practical reasons are motivating the research in energy efficient HVAC systems.

Apart from the aforementioned HVAC system, the thermal systems also include the thermal management of the traction components, such as motors, inverters, and battery. The combination of these systems—e.g., introducing a heat pump—allows for waste heat recovery, thus reducing the energy demand for heating. It therefore makes sense to examine these systems holistically.

Of course, the coupling of all of these systems introduces various interdependencies, ultimately increasing complexity, both in the hardware setup, as well as in the control system. It therefore becomes harder to make informed decision about hardware selection and component dimensioning. Additionally, the necessary tuning of the involved controllers for the various actuators becomes increasingly difficult.

A model-based approach in conjunction with optimal control theory can simultaneously tackle all of these challenges.

Therefore, the research project Intelligent Software for Thermal Energy Management (ISOTHERM) has been initiated by the industrial partners Carrosserie HESS AG and VBZ together with the Institute for Dynamic Systems and Control (IDSC) at ETH Zürich.

0.2 Objectives

The following two objectives are formulated for the project ISOTHERM.

- The mathematical models that are to be developed are used in **simulations** to gain an improved understanding of the system. Simulation studies conducted to quantitatively assess the importance of various influencing and potentially competing parameters.

  The same procedure is used to examine the optimal combination of components, including their respective sizing. Specific questions can be evaluated in detail, such as for example the energetic influence of radiant heaters or air curtains to lower door losses.

- On the other hand, the mathematical models are used to improve the **control strategy** of the various components. In a central, integral thermal management system, setpoints are generated either for underlying controllers or directly for the components. In this way, the trade-off between energy demand and comfort is analyzed with the goal of Pareto-optimality.

  In addition to the presumed energy saving potential, model-based control approaches offers options for simpler tuning, by introducing intuitive parameters and target values.

0.3 Project Breakdown

Based on the project objectives, four work packages (WPs) are introduced in the project proposal. The work packages are graphically represented in figure 1 and explained in the following.

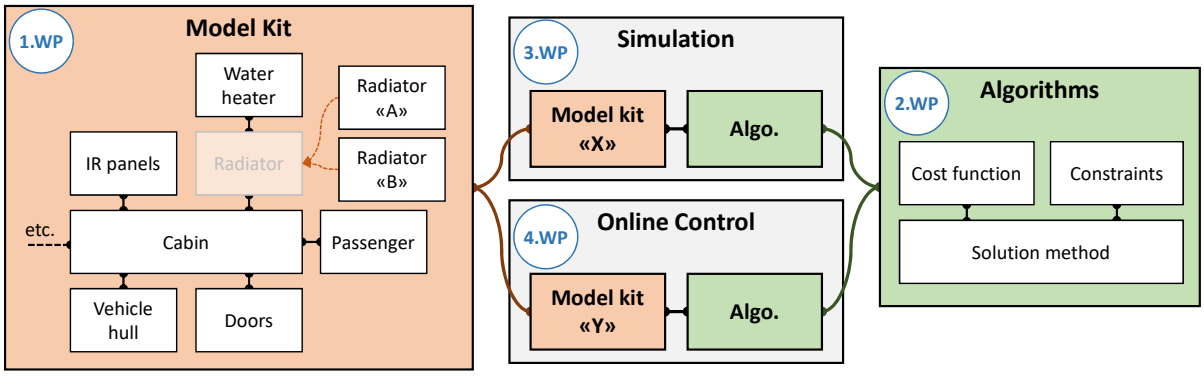

Figure 1: Project breakdown into four work packages (WPs).

1. Modeling:
   As a basis for the rest of the project, this first work package describes mathematically all thermal subsystems and the associated disturbances. The model is designed to be modular and extendable.
2. Algorithm:
   In the second work package the basics for the simulations and optimizations are laid. This includes the definition of a cost function and constraints as well as the development of one or more suitable solution algorithms.
3. Simulation:
   In simulation studies, statements about the optimal design of the thermal systems in a city bus are made. This includes statements about technology combination and dimensioning as well as positioning of components.
4. Online control:
   In the last work package, findings from the previous work packages are used to develop an integral control software for use on an electric city bus.

0.4 Project Delimitation

In this subsection, we provide the delimitation of the project ISOTHERM by commenting on previous and subsequent projects, as well as on the premature termination of the project ISOTHERM.

- SwissTrolley plus:
  The project ISOTHERM can be considered a direct successor of the project “SwissTrolley plus ”, which has been finished in 2019 [1]. While an analysis of the thermal systems has already been envisaged in the project SwissTrolley plus, some initial work revealed the high complexity of this topic. This was the initial impetus for the initiation of the research project ISOTHERM. Some of the findings presented in this report are based on the foundation that has been laid within the project SwissTrolley plus.
- Premature project termination:
  As it was not possible to assess the potential of the optimization of the control algorithms at the outset of the project ISOTHERM, an option to stop the project prematurely in a “Go/No-Go decision” after three years was added to the project plan. Should the energetic potential be found to be small, the project could be stopped prematurely, without spending too much time on the implementation of something with small energetic potential. At the same time, the plan in the description of the project ISOTHERM was foreseeing that the vehicle in the focus might change. A premature termination thus enables to prevent the focus from locking in on a vehicle with a sub-optimal component selection,

for instance. As we show in this report, potential improvements by optimized thermal management are clearly present. Thus, in our eyes, it is not necessary to stop the further research on the topic. On the contrary, a follow-up project was initiated by Carrosserie HESS AG, which is outlined in the following.

- Swiss eBus Plus:
  During the course of the project ISOTHERM, Carrosserie HESS AG initiated a new P+D project, i.e., the project “Swiss eBus Plus”. The main goal of this project is to develop a novel, purely battery-electric bus that overcomes the range issues with the current generation of battery-electric buses by combining latest technology components and intelligent control. Hence, in addition to the industrial partners Carrosserie HESS AG and VBZ, research groups from ETH and Bern University of Applied Sciences (BFH) are also involved in the project. Instead of running the ISOTHERM and the Swiss eBus Plus projects in parallel, the Swiss Federal Office of Energy (SFOE) and all project partners decided to stop the project ISOTHERM prematurely after three years and merge the research endeavors of the two projects instead. Thus, some of the tasks from the project ISOTHERM are moved to the project Swiss eBus Plus. These are listed in detail in section 5.2.

0.5 Report Structure

Due to the project breakdown into the four work packages as introduced above, we decided to use an unorthodox section numbering style, starting at zero for this current section. The following sections 1 to 4 then correspond to the four work packages, where each of these sections lists the most relevant results of the respective work package. Finally, section 5 provides a conclusion to this project and provides an outlook on the remaining work in the project Swiss eBus Plus on the one hand and lists additional ideas for future research on the other.

For convenience, appendix A provides a list of all scientific contributions and all student reports that have been generated in the course of this project.

1 Work Package 1: Modeling

Since we follow a model-based approach, this first section forms the basis of all subsequent work. We therefore aim to give a solid overview of the work performed and go into some detail in selected parts.

Section 1.1 provides insights into the modeling of certain components that offer particular challenges in their mathematical descriptions and were thus not straight-forward to model. Section 1.2 covers some work that was put into estimation and prediction of parameters that are hard to measure in real time. Section 1.3 shows a complete vehicle model that is based on the assumption of steady state. Section 1.4 introduces a modular framework capable of encapsulating dynamic component models and combining them into a complete vehicle model. Finally, section 1.5 concludes this section with a brief outlook.

1.1 Model Components

In this first section, we present a subset of the modeling work. Thereby, we focus on model components for which the modeling procedure was either not clear at the beginning or particularly challenging. In the following section 1.1.1, we first provide a system overview of all involved components, which helps to follow the remainder of this text. After that, some details of the modeling work for thermal comfort, the (stratified) cabin air, and the vapor compression cycle (VCC) are given in sections 1.1.2 to 1.1.4, respectively.

1.1.1 System Overview

In a first step, we provide a detailed overview of how the thermal systems of a current trolley bus by Carrosserie HESS AG are interconnected. The hydraulic circuit diagram in figure 2 shows the physical connections present on the real system and serves as the basis of the following modeling processes.

The following list explains the general purpose of the installed components:

- Water Heater and Hot Water Circuit
  In “conventional”, i.e., combustion-driven, buses, hot water is provided to the heating system from the engine cooling system. BEVs do not have such a “free” energy source. Therefore, an electric water heating coil is installed. The hot water pump, which is located directly in front of the water heater, is used to pump the hot water through the roof boxes and the front box near the driver.
- Roof Boxes
  Single-articulated trolley buses have two air conditioning boxes on their roof (one on the front carriage, i.e., in front of the articulation, and one on the rear carriage). They can be operated in either air-conditioning (AC) or HP mode and serve a multitude of purposes:
  - Ventilation of the cabin using outside air or recirculating cabin air. This includes, if necessary, heating (using the heat pump as well as the hot water provided by the water heater) or cooling of the ventilated air.
  - Only the rear roof box: Cooling of the battery cooling water.
  - Only the front roof box: Providing pressurized liquid refrigerant to the driver front box (for air cooling).
- Driver Front Box
  The “front box” near the driver seat has the purpose of providing a comfortable environment for the driver as well as to provide defrosting and dehazing functionality for the front windshield. It is operated by the driver, very similar to how the air conditioning system in a regular passenger car works. It can ventilate outside air or recirculate internal air, heating (using hot water) or cooling (using the refrigerant) the air in the process.

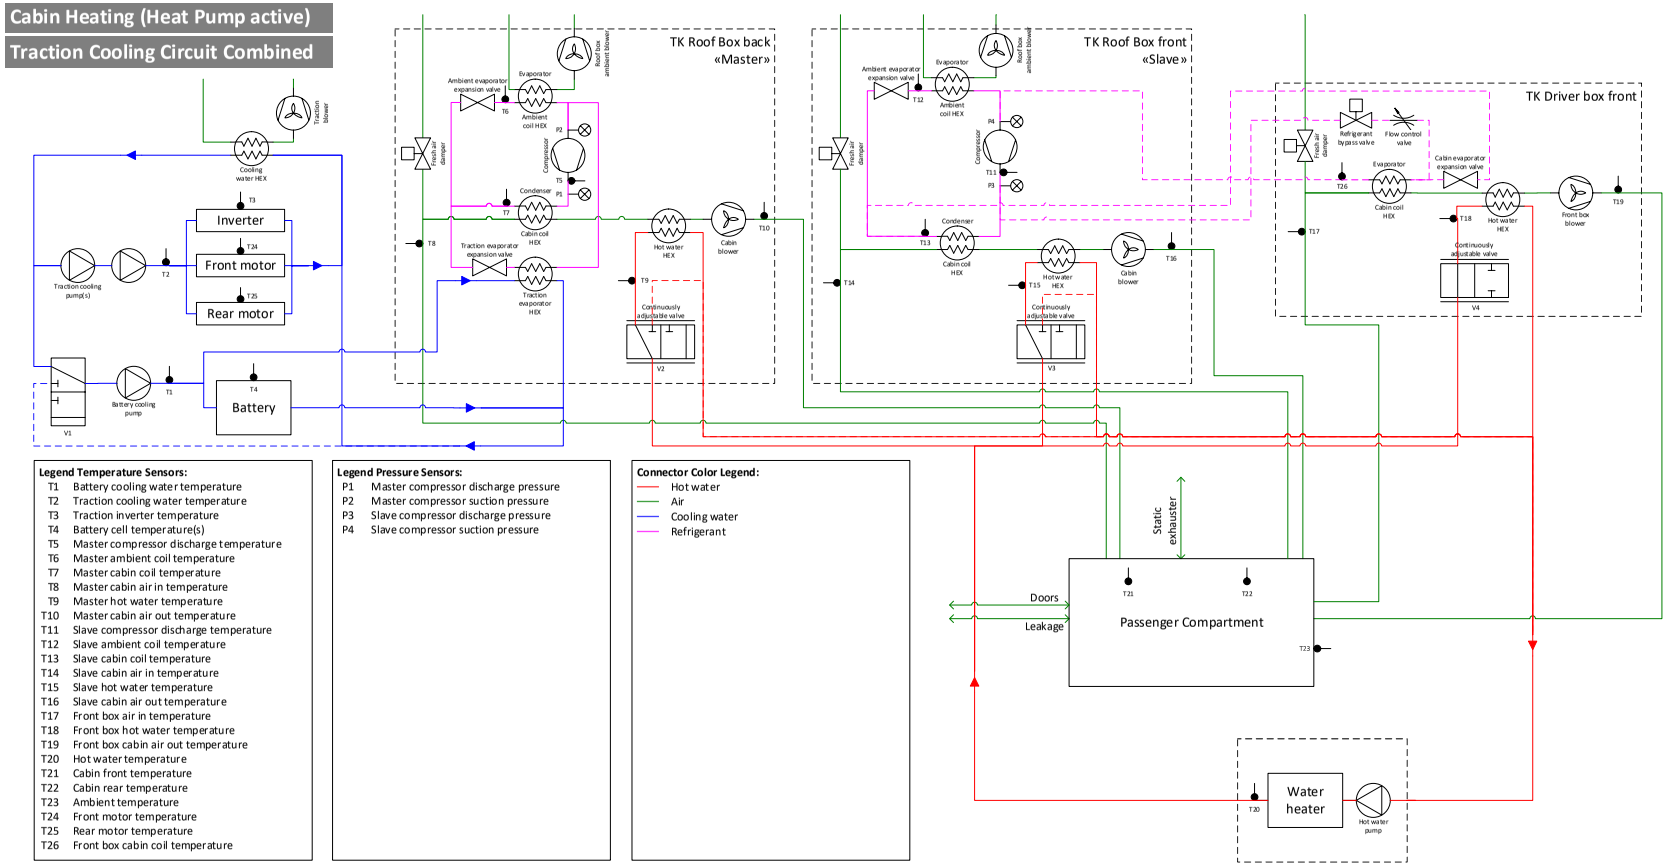

Figure 2: Hydraulic scheme of a trolley bus. This scheme depicts a typical heating scenario with both the heat pumps as well as the water heater being active. The two traction cooling circuits are combined in order to allow the heat pump to use the waste heat of all components. The two “roof boxes” and the “front box” at the driver’s compartment are both standard components with a simple control interface. For this reason, internal data for many of the sensors installed in these boxes is unfortunately not available on the bus-internal communication systems.

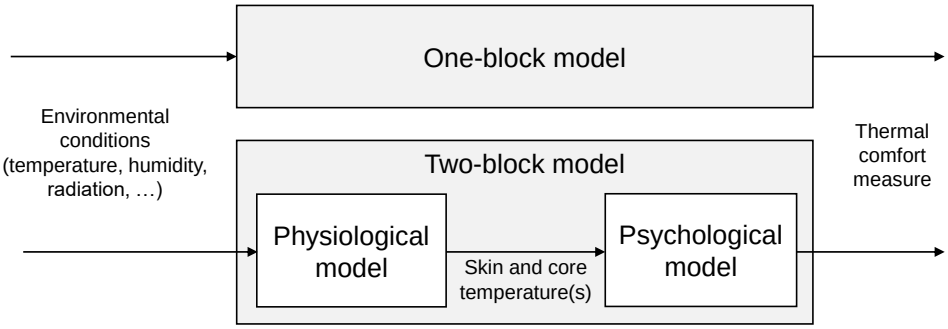

Figure 3: Visualization of the main categories of thermal comfort models.

- Traction Cooling
  The traction components of the buses are cooled in water-based cooling circuits: One circuit cools the battery, while the other circuit cools both motors as well as the power electronics. If the two circuits are operated independently, the heat from the battery is extracted by the rear roof box, while the heat from the motors and the power electronics is rejected directly to the atmosphere. The two circuits can also be combined in a single circuit, with a mixture of components in parallel and in serial.

It is noteworthy that the passenger compartment is heated and cooled by the air that is provided by the two roof boxes only. In particular, radiators that are typically installed in older-generation trolley or diesel buses are not installed on this generation of trolley buses to reduce weight and increase floor space.

1.1.2 Thermal Comfort

As already stated in the ISOTHERM funding application, the energy consumption of the HVAC system is inherently linked with the thermal comfort. In fact, the two objectives of minimum energy consumption and maximum thermal comfort are conflicting. To be able to quantitatively treat these conflicting objectives, it is important to have a mathematical description of not only the technical components of the HVAC system, but also the human thermal comfort.

The obvious challenges in this regard is that, on the one hand, the human organism and its metabolism is a very complex system, where the heat production and dissipation is influenced by many internal and external factors. On the other hand, the thermal comfort is ultimately not an objective measure that can be completely traced down to quantitative measures like temperatures.

In a literature survey, Huber found that most thermal comfort models can be divided in either of two categories [2], as visualized in figure 3:

One-block models directly link all the inputs to a thermal comfort output.

Two-block models attempt to manage the complexity of the modeling part by separating it into two processes. First, the environmental conditions are used as inputs into a physiological model, where intermediate physical quantities like skin and core temperatures are calculated based on a thermodynamic model of the human organism. As shown in figure 4, such models can become arbitrarily complicated. Second, these intermediate values are used in a psychological model to predict how a person might feel under these circumstances. Such models can either predict a “global” thermal comfort or a “local” comfort sensation.

For the project ISOTHERM, we chose to use the one-block comfort model by Fanger [4]. This model is broadly used for thermal comfort, both in buildings and vehicles and it has been used in the definition of the ISO 7730 norm [5]. It consists of an empirical formula for the predicted mean vote (PMV) $\Psi \in [-3, 3]$ as shown in table 2. The calculation is based on the surrounding air temperature $T_{\mathbf{air}}$, the mean radiant

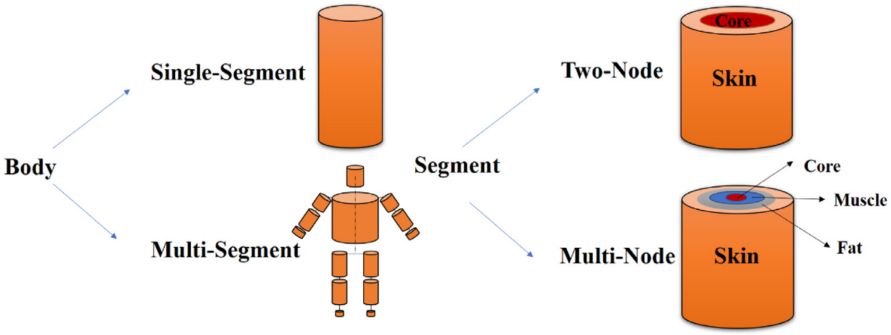

Figure 4: Possibilities of subdividing the human body into one or more segments using two or more nodes [3].

Table 2: PMV scale [5].

PMV	Description
$+3$	Hot
$+2$	Warm
$+1$	Slightly warm
0	Neutral
-1	Slightly cold
-2	Cool
-3	Cold

temperature T_{MR}, the metabolic rate M, the mechanical power performed by the person W, the clothing insulation I_{cl}, the air velocity v_{air}, and the relative humidity ϕ_{air}. It can thus be interpreted as a function

$$\Psi = f_{\mathrm{EN-ISO\,7730}}\left(T_{\mathrm{air}},\, T_{\mathrm{MR}},\, M,\, W,\, I_{\mathrm{cl}},\, v_{\mathrm{air}},\, \phi_{\mathrm{air}}\right). \tag{1}$$

Apart from the PMV, the norm additionally provides a quantity to further analyze thermal comfort. The predicted percentage dissatisfied (PPD) can tell us, depending on the PMV, how “uncomfortable” certain ambient conditions are.

1.1.3 Stratified Cabin Model Air

Measurements performed on a trolley bus in a climate chamber have shown a very significant temperature stratification within the cabin [6]. The largest temperature differences of up to 24 K between top and bottom measurements have been observed during periodic door openings at very low ambient temperatures. However, these measurements have been conducted without passengers, which are expected to influence both the air exchange around the doors as well as the temperature stratification.

To be able to assess the factors influencing the stratification quantitatively, a Bachelor’s thesis has been conducted by Solèr [7]. During the initial phase of the work, some single publications in the scientific literature could be identified dealing with temperature distribution in vehicles, but they mainly present high-resolution measurements or computational fluid dynamics (CFD) simulations and no simplified low-dimensional models for fast and efficient parameter studies. In the building domain, such simplified models are more common and some of them have been used and optimized for more than three decades. However, the environmental conditions in buildings are very different from those in moving buses and thus, the influences of e.g., passengers, door openings and driving speed are typically not modeled.

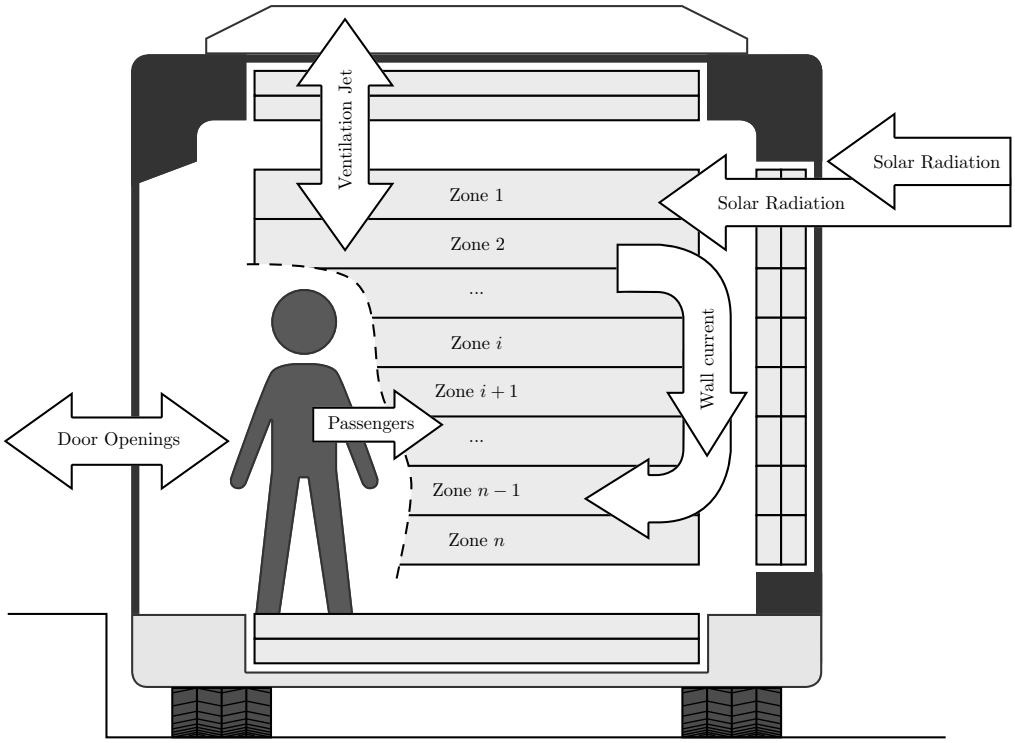

Figure 5: Overview of the developed model for the thermal stratification in the cabin air [7].

Based on basic thermodynamic equations and existing models from the scientific literature in the field of building simulation, a model specific for the trolleybus was developed in this Bachelor's thesis. The developed air temperature model is based on the finite volume method (FVM) approach and was extended with models for air mass exchange through door openings, the HVAC system, heat input by passengers, thermal buoyancy along walls and windows and solar radiation (see figure 5). Importantly, no a priori assumptions are made for the air circulation, as is often the case with simplified models in the building domain. This feature allows the model to be applied to both heating and cooling scenarios, although the air circulation setting is typically different in these two cases. The data from a the climate chamber experiments was used for model calibration.

To validate the developed model, additional temperature stratification measurements in regular bus operation have been performed. Small temperature sensors were temporarily installed in a VBZ trolley bus at five different heights during a cooling and a heating scenario. Figures 6 and 7 show the resulting recorded temperature profiles for the cooling and heating scenario, respectively. Figure 8 shows the mean stratification over the height of the bus cabin, along with the simulated trajectories, for both the cooling and heating scenario.

These measurements, as well as the results of the model, indicate that the temperature stratification at moderate winter temperatures is in the range of $3\,^{\circ}\mathrm{C}$ to $4\,^{\circ}\mathrm{C}$, which is significantly lower than the values measured in the climate chamber. Such a discrepancy has also been observed for Trams by Strebel, who observed a stratification of up to $9\,^{\circ}\mathrm{C}$ during standstill in the garage, but only about $4\,^{\circ}\mathrm{C}$ during regular operation [8]. As expected, stratification is less relevant in summer. Consulting literature, this seems like an acceptable range: In the American Society of Heating, Refrigerating and Air-Conditioning (ASHRAE)

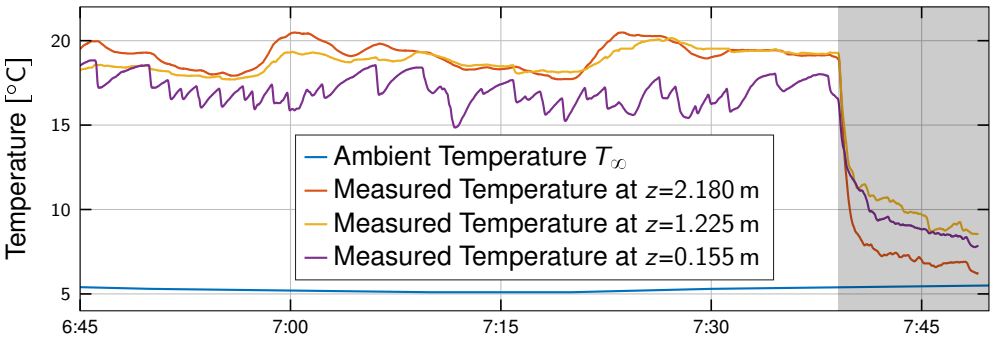

Figure 6: Visualization of the temperature recordings on the 27.04.2022 [7]. The grey shaded area represents the time during which the sensors were outside of the bus. The ambient temperature is measured by MeteoSwiss at Kaserne Zurich.

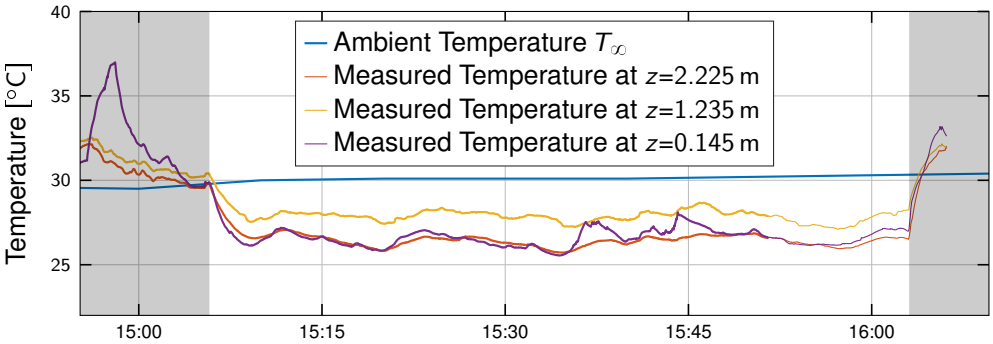

Figure 7: Visualization of the temperature recordings on the 17.06.2022 [7]. The grey shaded area represents the time during which the sensors were outside of the bus. The ambient temperature is measured by MeteoSwiss at Kaserne Zurich.

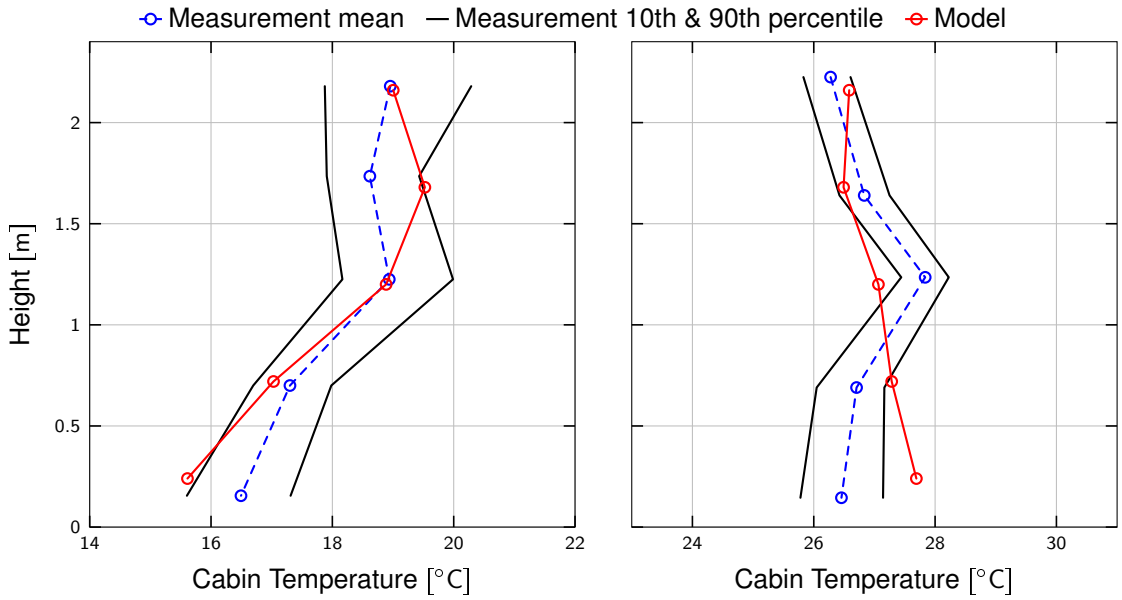

Figure 8: Measured and simulated stratification for an ambient temperature of $T_\infty = 5\,°C$ (left) and $T_\infty = 30\,°C$ (right) in real world operation [7].

Standard 113, a stratification of $3\,°C$ is allowed [9]. Zhang et al. conclude that if the mean temperature is near the comfort zone, a stratification of up to $7\,°C$ can be accepted [10].

The results presented in this Bachelor's thesis have to be taken with a grain of salt. The measurements have been conducted at only one place in the passenger cabin and only during short periods of less than one hour each. More data is necessary to validate the model, particularly for the air jets which have been massively simplified due to lack of measurement data. However, with respect to the overall stratification, we think that these results indicate that this can be neglected for most considerations in buses.

1.1.4 VCC Model

As the HVAC consumption is a large factor of the energy consumption of electrically propelled urban public transport vehicles, more expensive concepts to provide the necessary heat can be explored. Among the most promising candidates is the use of heat pumps, which typically use ambient heat to increase the heat provided to the passenger cabin per amount of electric power used. More complex concepts additionally allow to use waste heat from the traction cooling circuit. Heat pumps are in principle "reversed" air conditioning devices. Thus, they are not a separate new component, but are rather a more complex air conditioning device. As these devices work by evaporating and condensing a refrigerant in a vapor compression cycle, we will use the term VCC for the remainder of this text for this system, where both air conditioning and heat pump mode are considered.

The VCC forms the central component of the HVAC system of an electric city bus. It is the only component able to provide cooling in summer. Due to its limited capacity, it is typically not the only component to provide heating in winter, but is the most efficient one. In terms of energy management, it is also the key component, as it offers the most degrees of freedom to control the magnitude and direction of heat flows among the various components.

Apart from its importance, the VCC is also a highly complex system. It involves heat transfer among various different fluids. Additionally, the working fluid undergoes a phase change in the process. In a literature review, Furrer found that there are two main ways how VCCs can be modeled [11]:

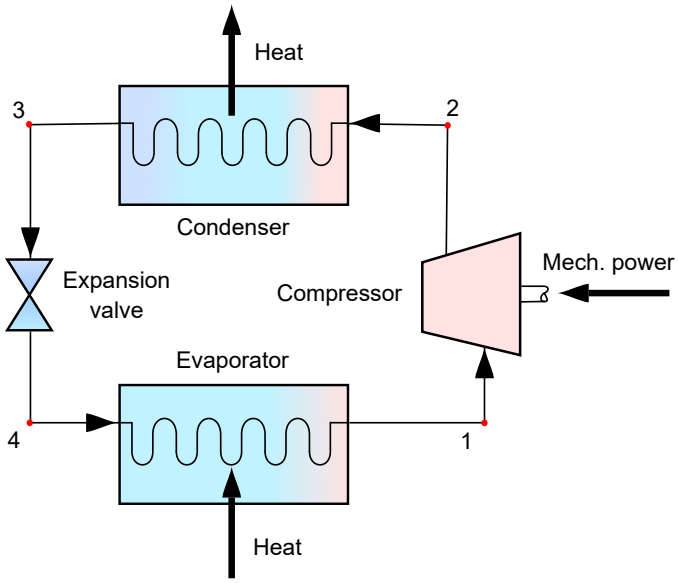

Figure 9: Refrigerant flow through the four main components of a VCC [11]. Depending on the use case (cooling or heating), the cabin evaporator or condenser serves as the cabin-side heat exchanger. A schematic pressure-enthalpy diagram is given in figure 13.

One-block models are typically used when a static model formulation can be used.

Four-block models are typically used for dynamic models of the VCC. In such models, the refrigerant flow through the four most important components of a VCC is modeled, see figure 9. The refrigerant state is then given at locations 1 to 4 between the components.

For the analyses presented in this report, both a static and two dynamic models of a VCC based on the R407c refrigerant have been developed, which are introduced below. The static model in section 1.1.4.1 can be used for steady state conditions. However, if the dynamic operation of the VCC should be analyzed (e.g. switching on or off), dynamic models are necessary, such as the ones shown in sections 1.1.4.2 and 1.1.4.3. Finally, section 1.1.4.4 comments on the choice refrigerant, in part as a outlook on the following project Swiss eBus Plus.

1.1.4.1 Static VCC Model

In a static VCC model, the energy inputs and outputs are then linked by the coefficient of performance (COP), which is defined as follows:

$$\eta_{\text{cop,cool}} = \frac{\dot{Q}_{\text{cool}}}{P_{\text{compr}}}, \tag{2}$$

$$\eta_{\text{cop,heat}} = \frac{\dot{Q}_{\text{heat}}}{P_{\text{compr}}}, \tag{3}$$

where the first formulation is given for the cooling mode, and the second formulation is given for the heating mode. Our model for this COP is shown in figure 10. The regressions are based on practical experience by engineers of Carrosserie HESS AG and information from the manufacturer about steady-state operating points. This model is thus able to predict the performance of the VCC under steady-state conditions.

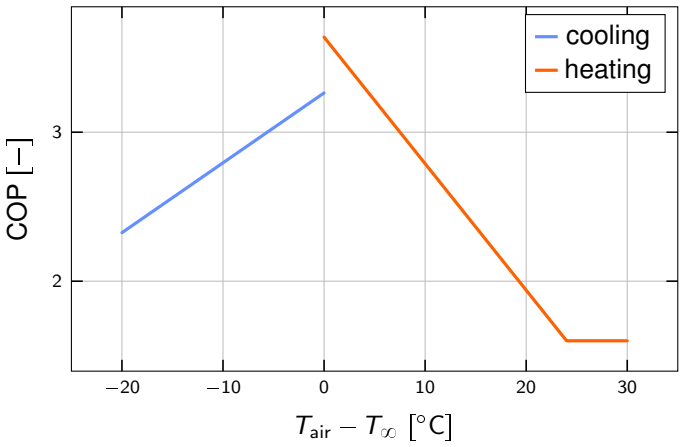

Figure 10: COP for heating and cooling, depending on the temperature difference between cabin and ambient.

1.1.4.2 Dynamic High-Dimensional VCC Model

A complete model of the VCC was developed in a Bachelor's thesis by Krause [12]. For this purpose, we use the Simulink extension Simscape, which allows to model components based on the underlying physics. Different model components are already available as templates to get started. Among those, there are also models for components that work with two-phase flow[1]. The different components are then connected in Simscape, where each connection resembles a physical connection in the real world and is characterized by "through" and "across" variables. For instance, a connection of electrical components ensures that all connected components share the same voltage ("across variable"), i.e., Kirchhoffs first law. For the current ("through variable"), a connection ensures that the sum of all inflowing and outflowing currents is equal, i.e., Kirchhoffs second law.

The resulting top-level overview of the Simscape model of the heat pump is shown in figure 11. Apart from the four main blocks already shown in figure 9, the model additionally includes block for the pipes between those components. These pipes add some transport delay and thermal inertia to the system. Additionally, the receiver tank and accumulator are also modeled. These systems, which are also present in reality, add further dynamics to the system.

The parameters of the models were tuned mainly using steady-state data of the COP, as shown in figure 10. For compressor power and delivered heat, the mean relative error of the model is around 3%.

While this data allows to tune the steady-state performance of the model, it of course does not yield any information about the dynamic behavior of the system. For this purpose, mobile temperature sensors were installed on a bus in operation by VBZ. With these sensors, the surface temperature of the refrigerant pipes was measured in multiple locations around the VCC. Some comparisons of measured and simulated pipe temperatures are shown in figure 12. This visualization clearly shows that the most relevant dynamics of the VCC can be well approximated in simulation.

Finally, the Simscape VCC model was re-arranged to be able to simulate the cooling mode using the exact same components as in the heating mode. These simulations were used as model validation by comparing them again to steady state performance data from the manufacturer. For compressor power and heat removal, the mean relative error of the model is around 4%.

[1] `https://ch.mathworks.com/help/simscape/two-phase-fluid-models.html`

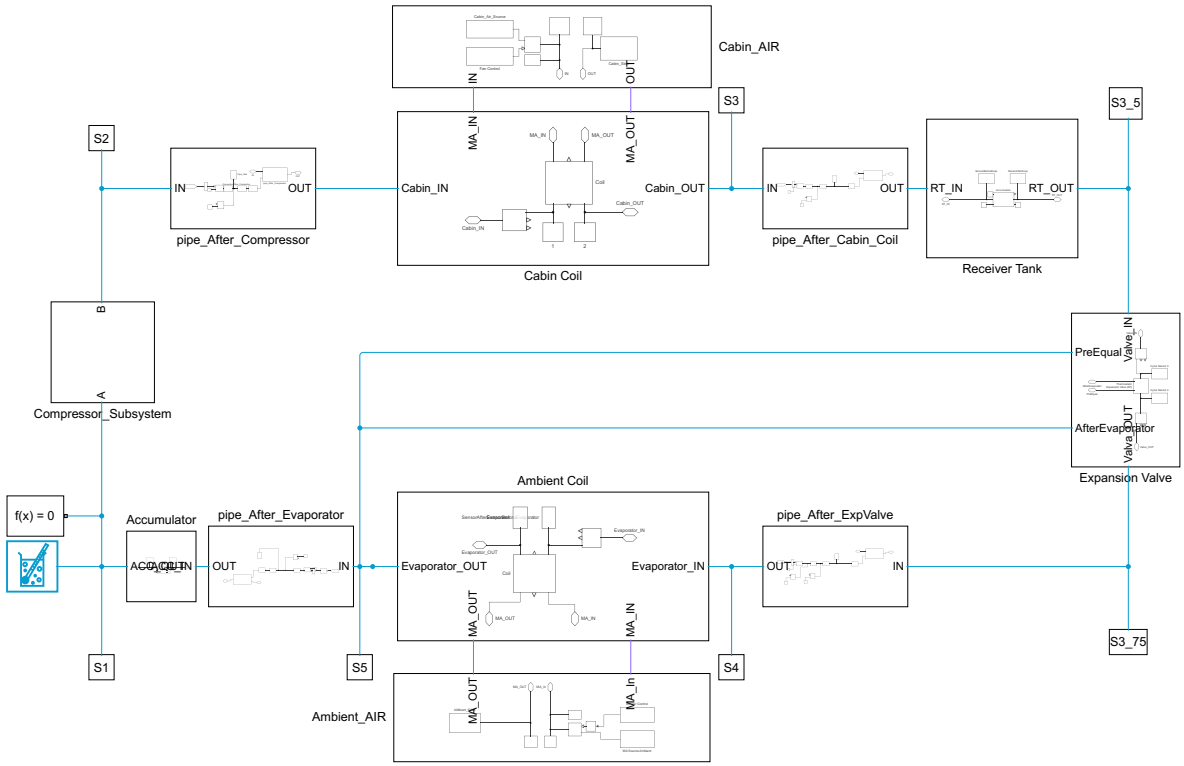

Figure 11: Top-level overview of the Simscape model [12]. The blue lines represent refrigerant connections, where the states of the refrigerant “S1” to “S4” match the ones shown in figures 9 and 13. The purple lines connecting the two heat exchangers to the cabin air and the ambient air, respectively, represent “moist air” fluid connections.

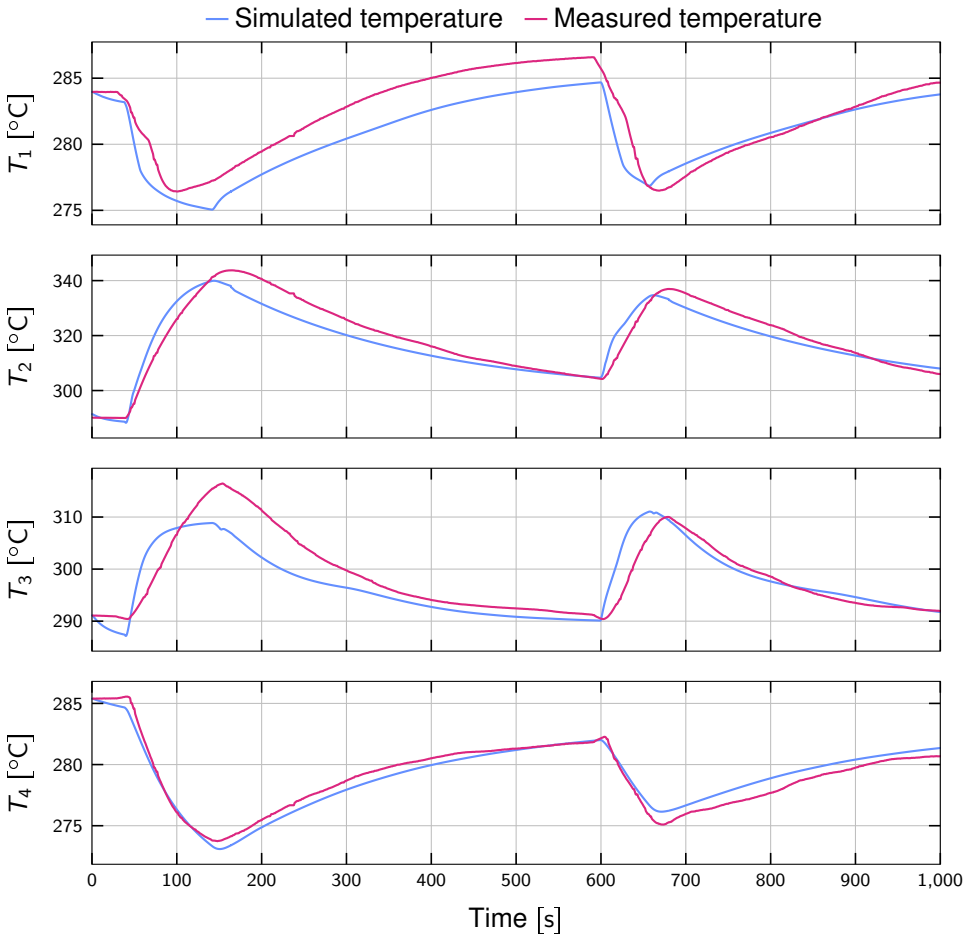

Figure 12: Comparison between measured and simulated pipe temperatures in various locations in the VCC, using the high-dimensional Simscape model [12]. The numbers of the locations (1 – 4) correspond to the ones shown in figure 9.

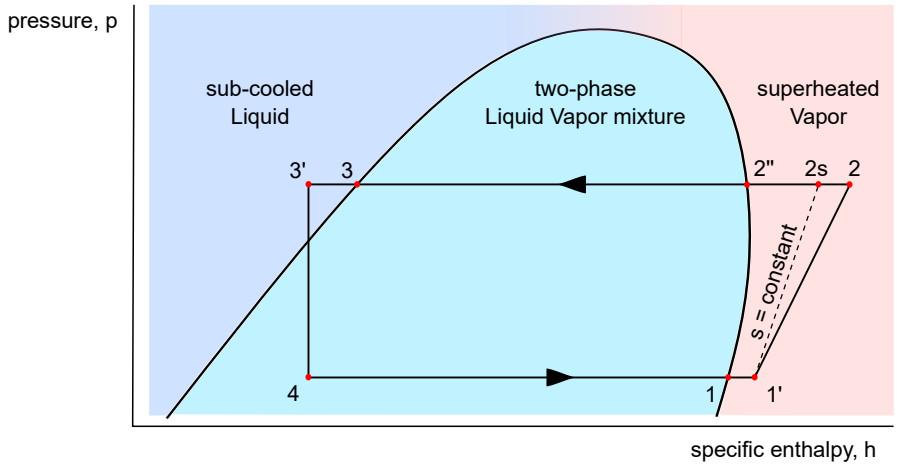

Figure 13: Pressure-enthalpy diagram of a classical VCC [11]. The locations (1 – 4) correspond to the ones shown in figure 9.

1.1.4.3 Dynamic Control-Oriented VCC Model

The high-dimensional Simscape model shown in section 1.1.4.2 is not suitable for optimizations due to its high complexity. Therefore, a third model of the VCC has been developed, which attempts to bridge the gap between the two models presented above. Specifically, this model should be able to capture the most relevant dynamics of the VCC system with only a low number of dynamic state variables, such that it can be used in optimization frameworks. In his Mater's thesis [13], Stähelin initially adapted the model presented by Zhang et al. [14] to the specific case of our trolley bus VCC.

The model initially proposed by Zhang et al. has two dynamic states, one for the pressure in each of the two heat exchangers. It is based on the following simplifying assumptions on the VCC [14] (see figure 13 for a pressure-enthalpy diagram of a VCC):

- The pressure in the two heat exchangers is assumed constant (negligible pressure losses). The refrigerant is assumed to be present only in two-phase state. The refrigerant properties are averaged over the heat exchanger volume and thus, over the evaporation or condensation process.
- The mass flow of the refrigerant is uniform throughout the VCC.
- The thermal mass of the heat exchangers themselves is lumped into the thermal mass of the refrigerant.
- The compressor is modeled with a set of algebraic equations, where isentropic and volumetric efficiency depend on the pressure ratio and compressor frequency via empirical formulas.
- The expansion process is perfectly isenthalp, hence $h_3 = h_4$.
- The state after the condenser (state 3) is exactly saturated liquid, i.e., fully condensed and no sub-cooling.
- The expansion valve is assumed to perfectly control the superheat before the compressor to $\Delta T_{\text{superheat}} = 5\,\text{K}$.
- Heat exchange to the air can be modeled with the ε-NTU method.

Figure 14 shows a simulation comparison between the high-dimensional and the control-oriented VCC models for a switch-on of the VCC at 300 s. While this results shows some clear deviations between the two models in the steady-state performance, it shows that the low-dimensional model is able to predict the

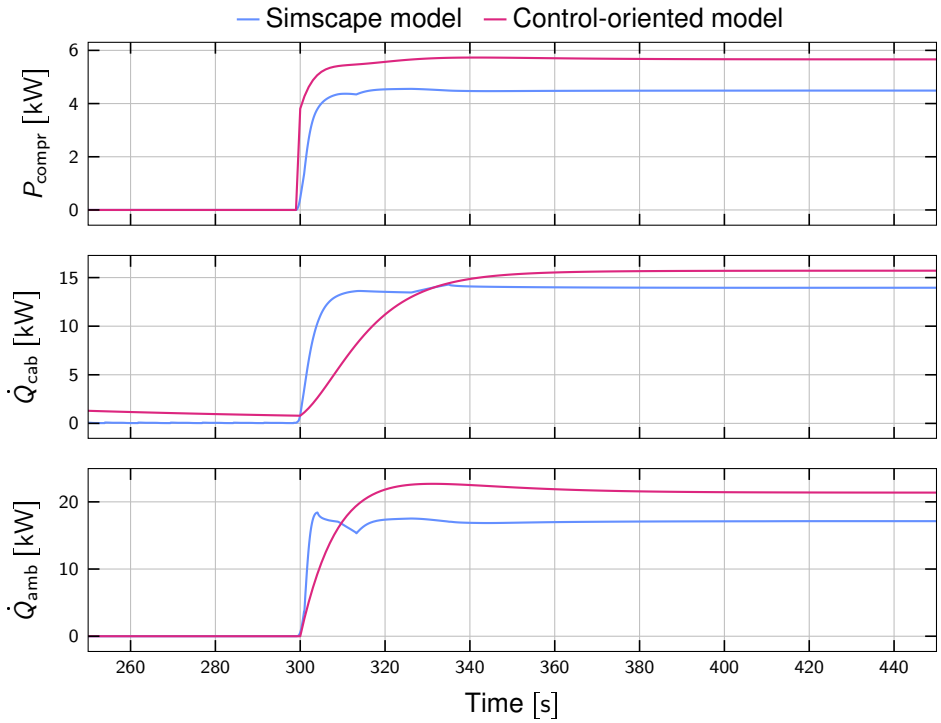

Figure 14: Comparison between the high-dimensional Simscape and the low-dimensional (control-oriented) VCC model based on [14].

general behavior of the VCC quite well, both in terms of its dynamic response as well as its low-frequency behavior. With some tuning, we are convinced that the fit could be improved further.

The refrigerant properties of R407c are evaluated using the open-source thermophysical property library CoolProp [15]. However, this library cannot be used with derivative-based optimization tools. For this purpose, the fluid properties in the relevant state regions are approximated using polynomial fits.

Since this model is based on physical considerations, it is relatively easy to extend it to a dual-source VCC, i.e., a VCC with two evaporators. In this case, the model still only has two dynamic state variables, but one more heat flow.

1.1.4.4 Refrigerants

The models shown in the sections above are all based on an R407c refrigerant, which is used in many current automotive VCCs. For the new project Swiss eBus Plus, Carrosserie HESS AG is planning to include a VCC based on CO_2 (R744). Apart from being environmentally friendly (being non-toxic and with a greenhouse warming potential (GWP) of one), R744 can offer a higher heating capacity and COP at lower ambient temperatures [16].

Compared to the subcritical VCCs modeled in this section, VCCs based on R744 are typically operated transcritically, i.e., with supercritical heat rejection and subcritical heat absorption processes [16]. Therefore, at the time of writing this report, we are unsure which parts of the modeling process can be transferred to R744 VCCs. We therefore decided not to put more work in these models before any data of the new VCC is available.

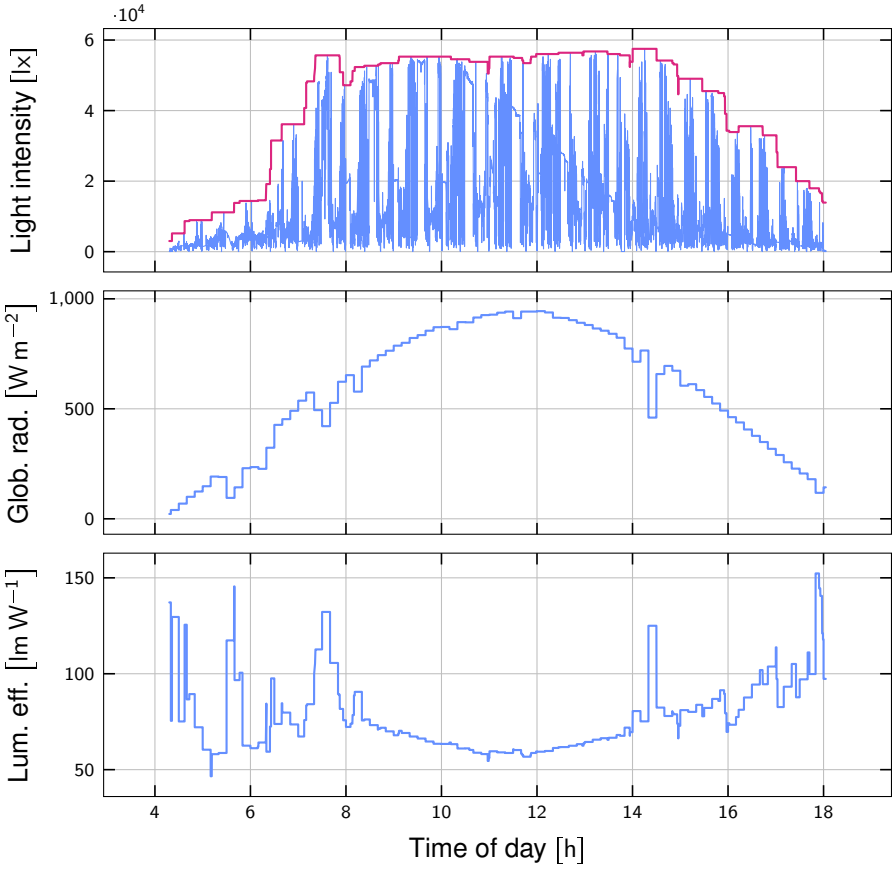

Figure 15: Analysis of light intensity measurements on 2022-06-20. The top graph shows the raw light intensity measurements (blue) along with a moving maximum value with a 30 min window (magenta). The second graph shows global irradiation measurements. The third graph represents the estimated luminous efficacy of the sensor.

1.2 Estimation and Prediction

Estimation algorithms are a useful tools to determine the values of parameters that are hard to measure in practice. In this section, we show how the solar irradiation could be estimated based on simple luminance sensors in section 1.2.1. Additionally, we briefly introduce an algorithm for the simultaneous estimation of road grade and gross vehicle mass (GVM) in section 1.2.2. Since the number of passengers in the bus has a substantial influence on the thermal energy balance, a method for predicting the upcoming passenger load is shown in section 1.2.3.

1.2.1 Estimation of Solar Irradiation

In a side project, we attempt to estimate the solar irradiation based on a sensor measuring visible light intensity. Such sensors have the advantage that they are much cheaper than sensor measuring global radiation directly. For this purpose, we compare the visible light intensity measured in lx to measurements of the global radiation on a horizontal surface by MeteoSwiss in $\mathrm{W\,m^{-2}}$. The sensor by MSR used for this purpose[2] is installed on the roof of the bus.

[2] https://www.msr.ch/de/produkt/wlan-datenlogger-msr145w2d/

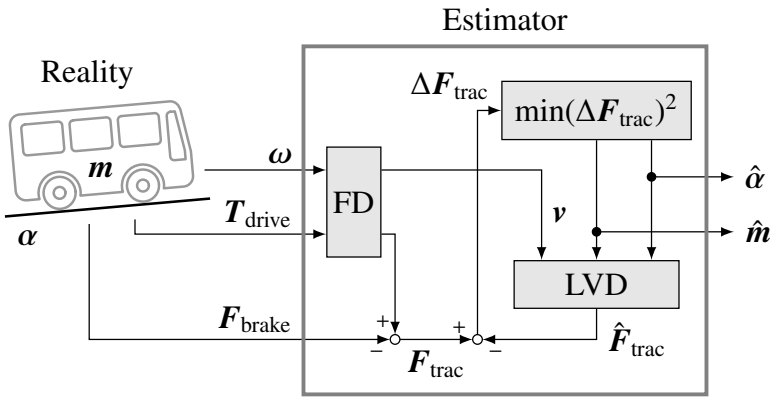

Figure 16: Method overview for simulataneous estimation of GVM $\hat{m}$ and road grade $\hat{\alpha}$ [18].

Figure 15 shows the result of this analysis. To account for short-term shadowing effects on the moving bus, we use a centered moving maximum value of the light intensity measurement. Using the the signals shown in the top two graphs of figure 15, we can estimate the luminous efficacy, which is defined as the ratio of daylight illuminance to solar irradiance expressed in $\mathrm{lm\,W^{-1}}$. The global horizontal luminous efficacy for daylight is expected to be around $90\,\mathrm{lm\,W^{-1}}$ to $100\,\mathrm{lm\,W^{-1}}$ [17]. The range calculated and shown in figure 15 is around this value, with a mean of $78\,\mathrm{lm\,W^{-1}}$.

This initial experiment is a promising result: With this approach, it might be possible to estimate the solar irradiation online, without the use of expensive additional sensor equipment.

1.2.2 Simultaneous Estimation of Vehicle Mass and Road Grade

In public transportation, GVM typically varies greatly with the number of passengers on board. Both parameters influence the operation of the vehicle:

- For a smooth acceleration from a standstill, it is important to have an estimation of the current GVM and the road grade.
- The number of passengers influence the operation of the HVAC system due to metabolic heat.

However, neither vehicle mass nor the number of passengers is typically available online, motivating the development of suitable estimation algorithms.

Therefore, an algorithm that simultaneously estimates GVM and road grade has been developed and published [18]. In contrary to previous estimators, this new approach exploits entire sequences of powertrain measurements at once and is formulated as a nonlinear program (NLP). A graphical overview of the approach is given in figure 16. The estimator is provided with time-resolved drivetrain measurements, i.e., the rotational speed $\boldsymbol{\omega}$, the torque $\boldsymbol{T}_{\text{drive}}$ at the driveshaft, and the force produced by the friction brake $\boldsymbol{F}_{\text{brake}}$. Based on this data and models of the final drive (FD) and the longitudinal vehicle dynamics (LVD), the estimator searches for a combination of the grade profile $\hat{\boldsymbol{\alpha}}$ and the GVM $\hat{\boldsymbol{m}}$ that minimizes the difference $\Delta\boldsymbol{F}_{\text{trac}}$ between “measured” and calculated traction force.

1.2.3 Passenger Count Prediction through Machine Learning

Predictive HVAC control algorithms might benefit from a prediction of the passenger load. However, the estimation algorithm for the GVM presented in section 1.2.2 can only provide estimates of the current passenger load. Thus, Stokholm attempted to predict the passenger load using machine learning algorithms in his semester project [19].

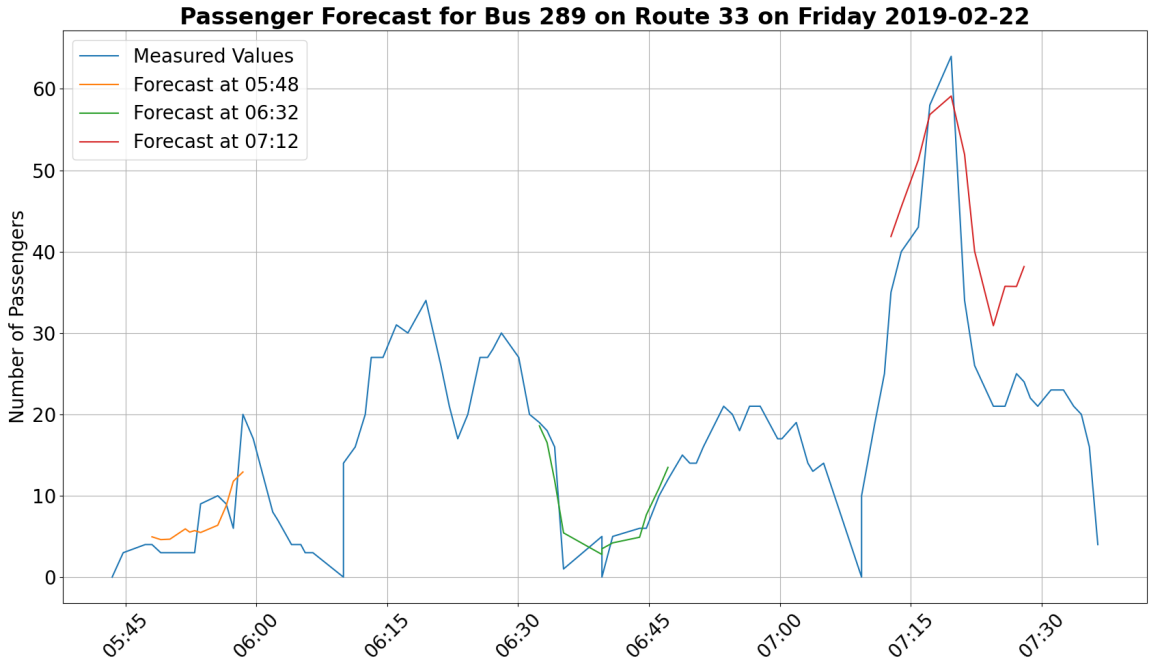

Figure 17: Passenger load forecast on a horizon of ten stops on a sample run in the test set (this data was not used for training), along with the ground truth labels [19]. A forecast is usually generated at every stop, however, in this graphic, only three forecasts are shown to improve clarity.

The training data used in this thesis was provided by VBZ. The labels (passenger counts) in the data is given by the infrared-based passenger counting system that is installed in about 20% of VBZ's buses. The data set used comprises data between December 2018 and December 2019 from all buses equipped with a passenger counting system. Overall, the data consists of 2.6 million data points.

The thesis allows the following conclusions:

- A reliable prediction with a mean absolute error of 5.5 passengers (for one-stop prediction) up to around 6.5 passengers (for a ten-stop prediction horizon) is achievable with a neural net with gated recurrent units (GRUs). Figure 17 shows some samples of such predictions.
- The most important group of features is the "time data", which includes the time of the day, the weekday, holidays, vacations, etc. Second most important is positional data (i.e., location from the global navigation satellite system (GNSS) sensor). If the line information from the dispatch system (i.e., current bus line, current and next stops, etc.) is available during prediction, this slightly improves the results. However, this data was not included for the results shown here, since it currently is not available on the bus-internal communication systems.
- Surprisingly, incorporating weather data into the features is counterproductive and decreases the performance in our case. Apparently, the variance in this data is quite large. This is expected to change if even more data were available.

1.3 Steady State Model

In this section, we briefly present a model that is based on the assumption of steady state. This model, along with some of the results generated, is presented in detail in a paper that has been submitted to be presented at the International Federation of Automatic Control (IFAC) world congress in 2023 [20].

This model has been built on the basis of three student theses:

- Bagajo developed the initial version of the model [21]. The focus was on analyzing radiant heaters. However, only one radiant heater and one passenger that is placed directly underneath the radiant heater are considered.
- First model extension by Büeler [22]. In this version, situations with several heaters and several passengers were investigated. Additionally, the analysis was extended to a “typical” winter.
- Further model extension by Achermann [23]. In this version, not only heating, but also cooling scenarios can be analyzed. Therefore, the comfort criterion is defined as an allowed band instead of a specific value. Based on this extension, the annual energy consumption can be analyzed.

While dynamic models might be a necessity for simulation of passenger vehicles, which are often operated in transient mode (e.g., cabin heat-up), we think that for the application in city buses, steady state models can be used for high-level approximations. The reason for this assumption is that the time constant of the relevant system dynamics in buses is in the order of tens of minutes. The disturbances, on the other hand, have either significantly faster time constants in the order of tens of seconds (door openings, change of driving speed) or significantly slower in the range of hours (ambient conditions like temperature). We can thus approximate the HVAC system performance in steady state.

Compared to dynamic simulations a system in steady state allows to simplify the analysis significantly, as it allows “invert” causality: For instance, instead of applying a certain heating power, for which the model returns a temperature, which can then be used to evaluate a thermal comfort model, we can use such a model to prescribe a certain thermal comfort, which will then determine the heating power necessary.

Figure 18 provides a graphical overview of all the four modeled thermal reservoirs, i.e., the radiant heater, the cabin air, the inside wall surface, and the outside wall surface. The model includes the following heat flows:

- Convection at the radiant heaters, the inside, and the outside wall,
- radiation between the radiant heaters and the interior, and on the outside wall,
- conduction in the wall,
- solar irradiation on the outside wall, the interior, and the cabin air,
- door losses to the ambient (through air exchange),
- metabolic heat by the passengers, and
- heating/cooling by the HVAC system.

The HVAC system consumes power to heat the radiant heaters on the one hand, and to provide heating (or cooling) to the cabin through the VCC. We use the model introduced in section 1.1.2 for thermal comfort.

The model shown in figure 18 can be used to formulate the energy balance in steady state:

$$0 = \begin{bmatrix} \dot{Q}_{\text{pass}} + \dot{Q}_{\text{h,rh}} - \dot{Q}_{\text{h,i}} - \dot{Q}_{\text{door}} + \dot{Q}_{\text{sol,a}} + \dot{Q}_{\text{hvac}} \\ P_{\text{rh}} - \dot{Q}_{\text{rad,rh}} - \dot{Q}_{\text{h,rh}} \\ \dot{Q}_{\text{h,rh}} + \dot{Q}_{\text{h,i}} + \dot{Q}_{\text{sol,w,i}} - \dot{Q}_{\text{k}} \\ \dot{Q}_{\text{k}} - \dot{Q}_{\text{h,o}} - \dot{Q}_{\text{rad,o}} + \dot{Q}_{\text{sol,w,o}} \end{bmatrix}, \tag{4}$$

We can then formulate a static optimization problem

$$\begin{aligned} \underset{\dot{Q}_{\text{hvac}}, P_{\text{rh}}}{\text{minimize}}\ & P_{\text{hvac}} + P_{\text{rh}}\,, \\ \text{s.t. energy balance } & (4)\,, \\ & T_{\text{rh}} = T_{\text{rh,target}}\,, \\ & \Psi \in [\Psi_{\text{min}}, \Psi_{\text{max}}]\,, \end{aligned} \tag{5}$$

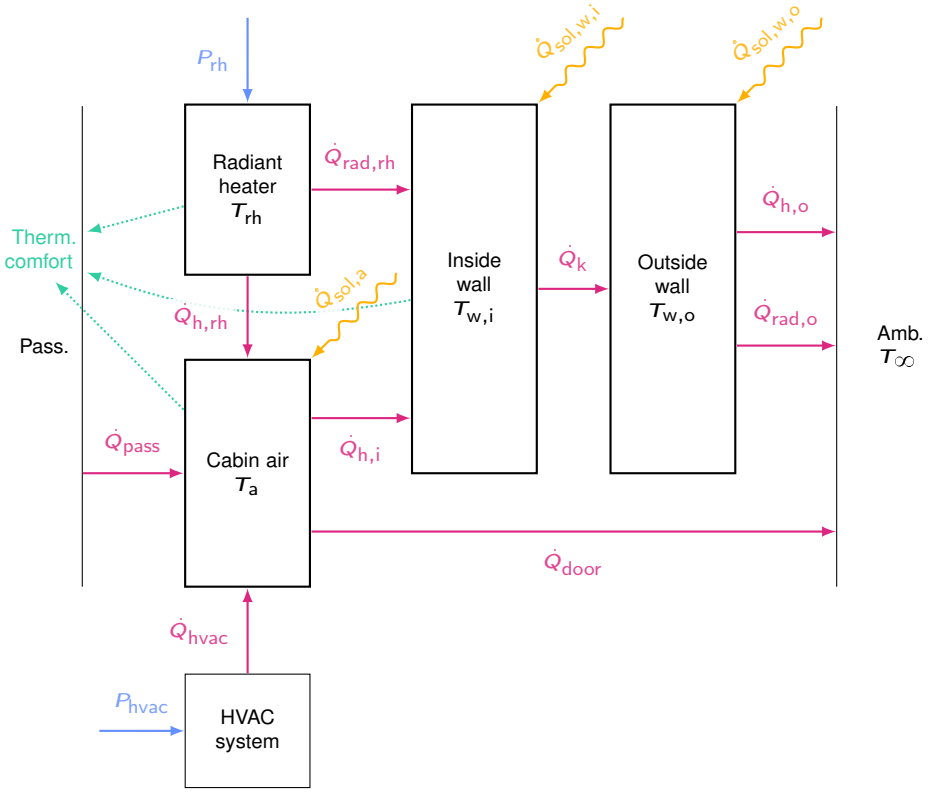

Figure 18: Model overview including the thermal reservoirs (bold blocks) and the connecting flows [20]. Red arrows denote heat flows, blue arrows denote electric power flows, and green dotted arrows denote the influence on thermal comfort.

Table 3: Model parameters.

Parameter	Value
Metabolic heat rate	1.20 met
Metabolic heat rate	125 W
Cabin air velocity	0.10 $\mathrm{m\,s^{-1}}$
Relative humidity	40.0 %
Convection inside	8.01 $\mathrm{W\,m^{-2}\,K^{-1}}$
Convection at radiant heaters	2.10 $\mathrm{W\,m^{-2}\,K^{-1}}$
Solar to air fraction	50.0 %
Window solar transmissivity	80.0 %
Paint solar absorptivity	30.0 %
Roof coverage	70.0 %
Door discharge coefficient	0.60
Radiant heater surface temperature	70.0 °C
COP of Heat pump / AC	Figure 10

where the two last conditions ensure that the radiant heaters are operated at their design temperature and that a thermal comfort within a desired comfort band is achieved, respectively.

We use the general model structure introduced above to perform simulations with three different public transport vehicle types, i.e., electric buses, regional, and intercity trains. For the trains, only one wagon is simulated, while for the bus, the whole vehicle is simulated. The model parameters are given in tables 3 and 4, where table 3 contains parameters that are common to all vehicles considered and table 4 contains the parameters that are different among the vehicles. For details of the model implementation and how the parameters are utilized, we refer to the original publication [20]. In the following, we will only briefly treat the most relevant parameters.

For the clothing insulation, we use a model presented by [24] and assume that passengers do not change clothing when entering a city bus. Furthermore, we introduce a lower limit of 0.3 in summer. For winter, we introduce maximum values for the two trains (see table 4), accounting for the fact that people might take off some clothing when entering these vehicles.

The values for the convective heat transfer coefficients are based on [25], adjusted for increased air movement inside the cabin and for an average bus velocity of $15\,\mathrm{km\,h^{-1}}$. For the regional and intercity trains, we use an estimated average velocity of $60\,\mathrm{km\,h^{-1}}$ and $100\,\mathrm{km\,h^{-1}}$, respectively.

For the fraction of time where the bus is in the shade (e.g., by buildings or bridges), we use season-dependent values as suggested by [26]. We use different surroundings for the bus and the trains, as listed in table 4.

The wall conduction values are tuned to match overall U-values. For the bus, this value has been experimentally determined to $2.91\,\mathrm{W\,m^{-2}\,K^{-1}}$ by Sidler et al. [6]. For the regional train, wagon, the U-value has been calculated as $2.42\,\mathrm{W\,m^{-2}\,K^{-1}}$ at standstill by Isenschmid [27]. Finally, for the intercity train wagon, the U-value has been calculated as $1.57\,\mathrm{W\,m^{-2}\,K^{-1}}$ by Sidler [28]. However, this last value has been calculated at “norm conditions”, i.e., not at standstill. This fact has been considered during the calculation of the conduction coefficient. The radiant heater panel size has been adapted to achieve a matching ceiling coverage ratio of 16.5% for all vehicles.

Section 3.2 shows how this model has been used in whole-year simulation studies.

Table 4: Model parameters that differ between the three vehicles.

Parameter		Trolley bus	Regional train	Intercity train
Hull geometry	m	18.7 × 2.6 × 2.4	14.3 × 2.8 × 2.8	27.0 × 2.8 × 2.8
Window fraction	%	35.4	16.9	16.9
Convection outside	$W\,m^{-2}\,K^{-1}$	20.7	70.7	115.1
Wall conduction	$W\,m^{-2}\,K^{-1}$	6.86	4.43	2.09
Door height	m	1.95	2.67	2.51
Total door width	m	4.42	1.66	2.04
Max. clothing insulation	clo	–	1.13	0.95
RH panel size	m	1.50 × 0.33	0.76 × 0.27	1.44 × 0.27
Shadow fraction		city	residential	residential
Mean no. of passengers	–	16.6	24.8	37.2
Mean door open fraction	%	21.8	10.0	6.7

1.4 Modular Model Framework

A modular framework was envisioned at the outset of the project ISOTHERM. Such a model is schematically shown already in the project breakdown in figure 1. Using such a modular principle allows individual components to be replaced, added or omitted at any time in a straight-forward manner. This is particularly useful, if different HVAC systems are to be compared in simulations, or if a model-based controller is to be used on vehicles with different HVAC configurations, as visualized in figure 1.

First, the Modelica modeling language has been used, as it was deemed a promising approach for exactly our purpose. This development is briefly shown in section 1.4.1. After some experience, we decided not to continue with Modelica, since it is not targeted for optimization-based modeling, which is what we intend to do. Therefore, in a second approach, a graph-based modeling framework was developed, which is explained in detail in section 1.4.2 further below.

1.4.1 Modelica-Based Toolchain

With the goal of a modular framework in mind, “object-oriented modeling” approaches quickly attracted our attention. Building on the concepts of object-oriented programming, it is particularly suited for large models. It builds on the idea that the governing equations of a physical component are the methods of the class, while the dynamic variables and the design parameters are the attributes of the class. The most commonly used object-oriented modeling language is Modelica [29]. “The Modelica Association” defines the language grammar and syntax while many third-party compilers can be used to perform simulations.

Modelica comes with a big availability of libraries for different physical domains, both commercial and open source. The Modelica Standard Library, maintained by The Modelica Community [29], already contains more than one thousand model components in different domains such as mechanical, electric, magnetic, thermal, and fluid. It also provides a standardized interface definition that improves the interoperability of all the other community libraries. Of specific interest are the libraries for buildings and automotive, sectors in which Modelica is mainly used.

The Modelica language is designed for describing the behavior of systems in a “forward” way, i.e., for simulations. Optimica, introduced in [30], extends Modelica language and allows the user to define optimization problems as well.

The research introduced in the above paragraphs motivated us build models on the basis of Modelica. Pelizzari initially developed a first complete vehicle model that represents all relevant thermal reservoirs in the simplest possible way using the Modelica language [31]. Later on, Imperatori developed a model for the traction cooling circuit using Modelica [32]. In this process, the difficulties of working with Modelica became clear. Unfortunately, no very good open-source tools are available and thus, debugging is an extremely hard task. While the finished models usually look very clean and logical, the way to get there is often not straightforward. Additionally, we found that one of the strenghts of Modelica is undoubtedly the availability of large plug-and-play model libraries. However, since these models are mainly optimized for simulation, they cannot be directly used for optimization purposes. Building a similar library with optimization-oriented models is an option, but is a very time-consuming task. Ultimately, we decided not to pursue Modelica for optimization-oriented modeling.

However, a first toolchain to be able to optimize models in Modelica has been built. This work is briefly presented in section 2.1.

1.4.2 Graph-Based Modeling Framework

The second approach to achieve flexibility and modularity within the developed models is to use a "graph-based" modeling approach. In such an approach, the model is represented in a graph, where energy is stored in nodes and power (or heat) can flow along edges. The following section 1.4.2.1 lays the theoretical foundation of graph-based modeling approaches. section 1.4.2.2 presents our framework, which is implemented in Matlab. Finally, section 1.4.2.3 explains the special treatment necessary to include VCCs in the graph-based models.

1.4.2.1 Theoretical Foundation

The approach we follow is very similar to the one presented in, e.g., [33, 34]. Hence, we also use a similar notation to those sources for the presentation of our approach.

Let $\mathcal{G}$ represent the directed graph corresponding to the mathematical model formulation. The directed graph consists of a set of nodes $\mathcal{N}$ and edges $\mathcal{E}$. This can be formalized as follows:

$$\mathcal{G} = (\mathcal{N}, \mathcal{E}) \,, \tag{6}$$

$$\mathcal{N} = \{n_i\}, \quad i \in \{1, \dots, N_n\} \,, \tag{7}$$

$$\mathcal{E} = \{e_j\}, \quad i \in \{1, \dots, N_e\} \,, \tag{8}$$

where n_i and e_j are used to represent the single nodes and edges, respectively. In total, the graph consists of N_n nodes and N_e edges.

There exist three different types of nodes

- **Dynamic** nodes have a dynamic state variable, which is related to a certain energy content. The energy content of a dynamic node $n_i \in \mathcal{N}^{\mathrm{d}}$ is represented by the state variable x_i. We use $\mathcal{N}^{\mathrm{d}}$ to denote the set of all dynamic nodes and $N_n^{\mathrm{d}} = |\mathcal{N}^{\mathrm{d}}|$ to denote its cardinality, i.e., the number of dynamic nodes.
- **Algebraic** nodes do not offer any energy storage capabilities. Thus, they can be interpreted as dynamic nodes with zero energy capacity. Similarly, we introduce $\mathcal{N}^{\mathrm{a}}$ to denote the set of all algebraic nodes and $N_n^{\mathrm{a}} = |\mathcal{N}^{\mathrm{a}}|$ to denote its cardinality.
- **Source** or **sink** nodes are used to represent connections to the world outside of the model. Mathematically, source and sink nodes can be treated identically. We therefore use the term "sink" from now on, but any energy "sources" are also represented with such nodes. Sink nodes can be interpreted as dynamic nodes with infinite energy capacity, i.e., nothing changes, no matter how much energy is drawn from or fed into a sink node. Again, $\mathcal{N}^{\mathrm{s}}$ represents the set of all sink nodes and $N_n^{\mathrm{s}} = |\mathcal{N}^{\mathrm{s}}|$ its cardinality.

It holds that

$$\mathcal{N} = \mathcal{N}^{\mathrm{d}} \cup \mathcal{N}^{\mathrm{a}} \cup \mathcal{N}^{\mathrm{s}}\,, \tag{9}$$

$$N_n = N_n^{\mathrm{d}} + N_n^{\mathrm{a}} + N_n^{\mathrm{s}}\,. \tag{10}$$

Each edge is defined by

$$e_j = (n_j^{\mathrm{tail}}, n_j^{\mathrm{head}})\,, \tag{11}$$

where n_j^{tail} represents the tail node and n_j^{head} represents the head node of edge j. The set of inlet and outlet edges for a node n_i are defined as

$$E_i^{\mathrm{in}} = \{e_j | n_j^{\mathrm{head}} = n_i\}\,, \tag{12}$$

$$E_i^{\mathrm{out}} = \{e_k | n_k^{\mathrm{tail}} = n_i\}\,, \tag{13}$$

respectively.

Each edge e_j represents a power flow P_j from its tail node to its head node. Thus, for a dynamic node n_i, the following differential equation holds, based on the conservation of energy:

$$\dot{x}_i = \sum_{\{j | e_j \in E_i^{\mathrm{in}}\}} P_j - \sum_{\{k | e_k \in E_i^{\mathrm{out}}\}} P_k\,. \tag{14}$$

Similarly, for an algebraic node n_i, the balance equation becomes

$$0 = \sum_{\{j | e_j \in E_i^{\mathrm{in}}\}} P_j - \sum_{\{k | e_k \in E_i^{\mathrm{out}}\}} P_k\,. \tag{15}$$

There is no need for a balance equation for sink nodes.

Generally, the achievable power values P_j in a certain edge e_j are not directly influenced by the energy content in adjacent nodes or the power values in another edge. For example, in a battery-assisted trolley bus, the amount of power provided by the battery can be chosen independently of the battery SOE. However, in some cases, power values cannot be chosen freely:

- An edge e_j can be considered a “disturbance edge”. In this case, the trajectory of the corresponding power value is completely pre-determined, i.e.,

  $$P_j = f(t)\,. \tag{16}$$

  In the context of control systems, this is typically referred to as a “disturbance” to the system.
- The power values of edges adjacent to the same node n_i can be linked by an equation of the form

  $$f(x_i, P_j, P_k, \ldots) = 0\,, \tag{17}$$

  assuming that the corresponding edges $e_j, e_k, \ldots$ are all connected to the node n_i. For instance, the power losses in a converter typically are linked to its output power via a conversion efficiency.
- The power value of an edge can be linked to the values of the adjacent node states using a function of the form

  $$P_j = f(x_j^{\mathrm{head}}, x_j^{\mathrm{tail}}, u_j)\,, \tag{18}$$

  where $f(\cdot)$ can be an arbitrary function of the two node states and an control input value u_j. For instance, the convective heat flow from a heating coil to the surrounding fluid depends on the temperature values, corresponding to the energy contents, of the heating coil and the fluid. The additional input u_j, in this case, could be a fan to increase the convective coefficient.

All equations introduced above, i.e., (14)–(18) only depend on a “local neighborhood” within the graph. This is an important property, as it ensures the flexibility of the developed models “by design”. For instance, if a certain node in the system corresponding to a specific component is to be removed, model adaptations are only necessary in the adjacent edges, while the rest of the model can remain the same.

1.4.2.2 Implementation in Matlab

A object-oriented implementation was pursued in Matlab. The following classes are provided to simplify the implementation of graph-based models:

- An overall class `GraphBasedModel` is used to combine both nodes and edges.
- To create nodes, the following classes can be instantiated. Equalities like (17) can be defined for any of these classes.
  - `AlgebraicNode`: A node for which (15) holds.
  - `DynamicNode`: A node for which (14) holds.
  - `ThermalNode`: A special version of a dynamic node for which the energy state can be linked to a temperature.
  - `SinkNode`: A node with no power balance equation.
  - `ThermalSinkNode`: A sink node with a temperature.
- Edges are always instantiated by directly connecting two nodes. The following classes can be used:
  - `GraphEdge`: A regular edge.
  - `DisturbanceEdge`: An edge with pre-determined flow variable according to (16).
  - `ThermalConductionEdge`: An edge where the flow variable is determined based on the temperature difference between the adjacent nodes (i.e., both nodes need to be thermal nodes or thermal sink nodes). Hence, the power flow is determined according to (18) without an additional input u_j.
  - `ControlledThermalConduction`: A thermal conduction edge, where the thermal conductance can be varied. Hence, the power flow is determined according to (18), including the additional input u_j.
- A utility class `SmartFunctionHandle` is used to define equations like (17), (18). This offers the advantage over standard Matlab function handle that it can parse its argument names.

An example model implementation is given in listing 1 below.

Listing 1: Example Matlab code to generate a graph-based model.

```
clc
clearvars
close all

%% Initialize model
mdl = GraphBasedModel('exampleModel');
dt = 1;
N = 900;

%% Nodes
% Algebraic nodes
dcLink = AlgebraicNode('dcLink');
mdl.addNode(dcLink);

% Dynamic nodes
batteryEnergy = DynamicNode('batteryEnergy');
batteryEnergy.nominalValue = 65 * 3.6e6; % 65 kWh
Rb = 0.09;
UbOC = 760;
batteryEnergy.addEquality(SmartFunctionHandle('battery_loss', ...
    @(P_b, P_bloss) (UbOC - sqrt(-4*P_b*Rb + UbOC^2)).^2/(4*Rb) - P_bloss));
batteryEnergy.initialValue = 0.7 * batteryEnergy.nominalValue;
mdl.addNode(batteryEnergy);

specificHeatWater = 4.15e3; % J/(kg K)
```

```
heatingWaterEnergyCapacity = specificHeatWater * 45;
heatingWater = ThermalNode('heatingWater', heatingWaterEnergyCapacity, 68);
heatingWater.upperBound = 7 * heatingWaterEnergyCapacity;
heatingWater.lowerBound = 0 * heatingWaterEnergyCapacity;
heatingWater.initialValue = heatingWater.upperBound;
mdl.addNode(heatingWater);

cpAir = 700; % J/(kg K)
volumeAir = 20*2.5*2.5; % m^3
rhoAir = 1.225; % kg/m^3
cabinEnergyCapacity = cpAir * rhoAir * volumeAir;
cabinNode = ThermalNode('cabin', cabinEnergyCapacity, 15);
cabinNode.initialValue = cabinEnergyCapacity * 3;
mdl.addNode(cabinNode);

% Sink nodes
ambient = ThermalSinkNode('ambient', 0);
mdl.addNode(ambient);

mission = SinkNode('mission');
mdl.addNode(mission);

sinkBat = SinkNode('sinkBat');
mdl.addNode(sinkBat);

%% Edges
% Regular edges
batteryPower = GraphEdge('P_b', batteryEnergy, dcLink);
batteryPower.upperBound = 350e3;
batteryPower.lowerBound = -350e3;
batteryPower.nominalValue = 150e3;

batteryLoss = GraphEdge('P_bloss', batteryEnergy, sinkBat);
batteryLoss.nominalValue = 5e3;

heatingPower = GraphEdge('P_h', dcLink, heatingWater);
heatingPower.lowerBound = 0;
heatingPower.upperBound = 50e3;
heatingPower.nominalValue = 50e3;

insulationConductivity = 10; % W/K
heatLoss = ThermalConductionEdge('Q_loss', heatingWater, ambient, insulationConductivity);
heatLoss.nominalValue = 1e3;

heatConsumption = ControlledThermalConduction('Q_cons', heatingWater, cabinNode, 0, 1000);
heatConsumption.nominalValue = 1e3;

heatLossCabin = ThermalConductionEdge('Q_cLoss', cabinNode, ambient, 2.9*200);
heatLossCabin.nominalValue = 1e3;

% Disturbances
data = load('powerRequestSignal.mat');
requestedPower = DisturbanceEdge('P_rest', dcLink, mission, data.powerRequestSignal);

%% Visualize graph model
figure
mdl.plot()
```

Note the following:

- An equation of the form (17) is used on lines 18–19 to define the battery losses.
- The “nominal values” of all nodes and edges is a useful property for numerical optimization tools, which we will introduce in section 2.2.

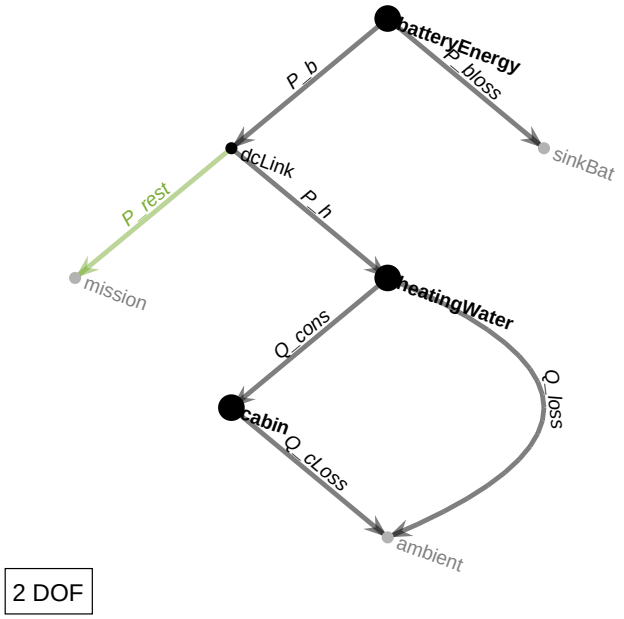

Figure 19: Automatically generated visualization of the graph-based model created in listing 1. Algebraic and dynamic nodes (including thermal nodes) are marked with small black dots and a large black dots, respectively. Sink nodes are marked with gray dots. Disturbance edges are drawn in green. The number of degrees of freedom (DOF) is determined based on the structure of the model automatically.

- For a thermal node, an offset temperature is specified. After that, all energy levels are with respect to this offset. For example, on lines 25–27, the heating water energy is defined relative to its offset temperature of $68\,^{\circ}\mathrm{C}$. The lower and upper bounds on lines 29–29 are defined relative to this energy, i.e., $0\,^{\circ}\mathrm{C}$ and $7\,^{\circ}\mathrm{C}$ higher.
- Finally, on lines 81–82, the model is automatically visualized, the result of which is shown in figure 19.

The example in listing 1 and figure 19 has 6 non-disturbance edges. The following 4 equations reduce the number of degrees of freedom to 2:

- Power balance of the form (15) at the direct current (DC)-link node
- Battery loss equation (line 20–21 in listing 1) of the form (17)
- Water and cabin loss conduction equations, which are of the form (18)

The two remaining degrees of freedom are the

- the heating power, and
- control variable of consumed heat.

1.4.2.3 Including VCCs in the Graph

To be able to include the VCC model introduced in section 1.1.4.3 in the graph-based modeling framework, separate node types are required. The reason for this separate treatment is that the dynamic state variables used in the dynamic VCC model represent the condenser and evaporator pressure levels. Hence, they do not directly represent an energy content, as is the case with the rest of this modeling approach. Thus, the following additional nodes are implemented:

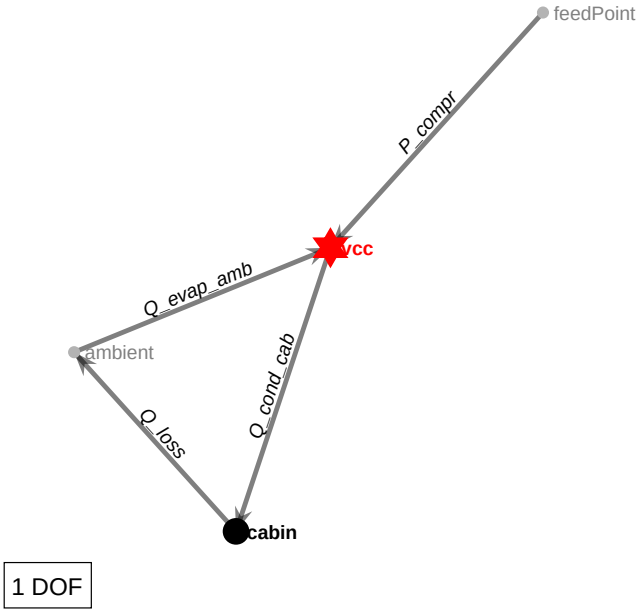

Figure 20: Automatically generated visualization of the graph-based model example in with a VCC node (red star). The VCC in this case represents a HP. It is directly powered by a “feed point” and can heat up the cabin by extracting heat from the ambient. The cabin, on the other hand, loses heat to the ambient. The VCC is assumed to have a variable compressor frequency, hence one degree of freedom (DOF).

- `VccNode`: As mentioned above, this node represents the two dynamic pressure level variables. It offers three connection points and requires that all three points are connected to edges, which is another difference to all other nodes, which in principle can have an arbitrary number of connections:
  - A compressor power edge, where the electrical power to the compressor is delivered.
  - A condenser and evaporator edge, which represent the heat rejected and absorbed in the condenser and evaporator, respectively. The nodes on the other side of the connecting edge need to represent a thermal reservoir, i.e., be either a `ThermalNode` or a `ThermalSinkNode`.

  A simple example of a graph-based model with a `VccNode` is given in figure 20. A `VccNode` typically has one degree of freedom, i.e., the compressor frequency. However, it is also possible to add degrees of freedom that scale the amount of heat exchanged in the condenser and compressor (representing variable blower or fan speeds, for instance).
- `DualSourceVccNode`: This node is similar to the `VccNode` with the difference that it has two evaporators. Thus, it requires two separate edges where heat can be absorbed by the two evaporators.

1.5 Outlook

In terms of component modeling and modeling framework, we consider the tasks outlined in the description of the project ISOTHERM to be fulfilled. The only exception in that regard is the CO_2 VCC (R744). For this kind of VCC, the existing models shown in section 1.1.4 need to be either adapted, if possible, or rebuilt, if necessary. This process will be started within the project Swiss eBus Plus as soon as possible, i.e., as soon as measurement or modeling data from the VCC manufacturer is available.

2 Work Package 2: Algorithms

This section deals with the development of general-purpose algorithms, as outlined in figure 1. A complete separation of models and algorithms, as indicated in figure 1, is desirable for several reasons.

- For example, the further development of the models is independent of the selected optimization algorithm and vice versa. Thus, different models can exist which can be optimized with different algorithms without having to know the respective link at the start of the modeling process.
- Furthermore, such a separation makes it possible to perform both optimizations and simulations (for example with classical control approaches) with the same model, as shown in figure 1.

Two approaches have been followed over the course of the project ISOTHERM. Initially, a pipeline based on the Modelica modeling language was developed, which is described in section 2.1. Later, the concepts of this toolchain have been adapted to work with the graph-base modeling approach. These algorithms are described in section 2.2. Finally, section 2.3 concludes this section with a brief outlook.

2.1 Modelica-Based Toolchain

In this section, we will briefly describe the prototype toolchain that has been developed based on Modelica and Optimica languages introduced in section 1.4.1.

Figure 21 shows the envisioned optimization approach. As this pipeline is per design independent of the models it can be applied to, it can not only be applied for the OCPs of the thermal management of a trolley bus, but is also applicable in projects of other domains. A first working prototype of the described pipeline has been developed in Python by Pelizzari in his Master's thesis [31]. The pipeline is able to process object-oriented Modelica models and automatically generate an OCP that can be further processed and ultimately solved by general purpose numerical optimization tools.

Following the completion of Pelizzari's thesis and some further development of the tool, we have written a paper proposal for the developed pipeline, which we were able to discuss with a few experts in the field. Ultimately, we decided not to pursue the Modelica approach, as already discussed in section 1.4.1, since we think the advantages of Modelica as a modeling framework are counter-weighed by the increased complexity when it is used for optimization-oriented models.

2.2 Algorithms for Graph-Based Models

The graph-based modeling framework introduced in section 1.4.2 is merely focused on the setup and connection of models. To generate useful results from these models, we need simulation and optimization

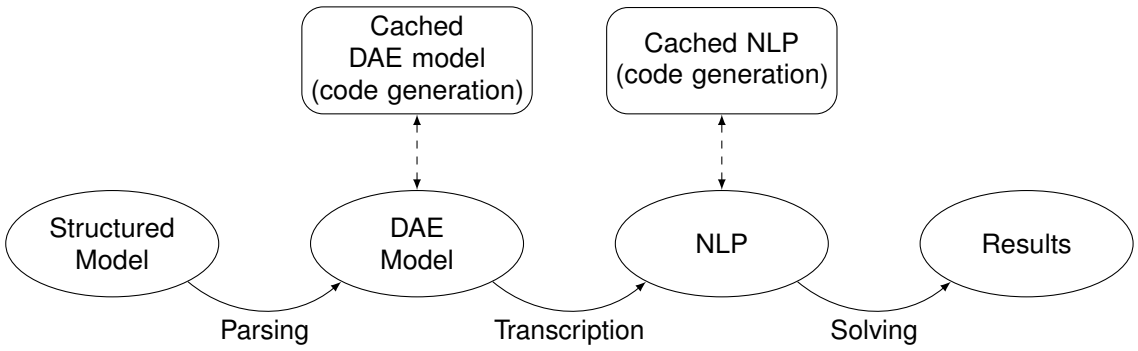

Figure 21: Suggested automated toolchain.

Listing 2: Example Matlab code to perform a simulation based on the model defined in listing 1.

```
1 %% Simulate
2 waterHeaterCtrl = HysteresisGraphController('P_h', 'heatingWater', ...
3     [heatingWater.lowerBound, heatingWater.upperBound], ...
4     heatingPower.upperBound, heatingPower.lowerBound, true);
5 cabinHeatingCtrl = PIGraphController('Q_cons_control', 'cabin', ...
6     cabinNode.initialValue, -1e-3, 100, 100, dt);
7
8 simResults = mdl.simulate(N, dt, {waterHeaterCtrl, cabinHeatingCtrl});
```

capabilities. Therefore, in this subsection, we first explain the steps necessary to perform simulations in section 2.2.1. Further below in section 2.2.2, we explain how the framework is used to formulate optimization problems.

2.2.1 Performing Simulations Based on Graph-Based Models

In simulations, the degrees of freedom in the model are controlled by one or more controllers. For instance, as mentioned previously, the power used to heat the heating water and the heat consumed for cabin heating represent the two degrees of freedom in the example shown in figure 19. Hence, these two flows need to be determined by one or two controllers. Listing 2 shows how this can be set up in the code.

In general, new controllers can be implemented for every application considered by extending the abstract class `AbstractGraphController`. However, for convenience, the following general-purpose controllers are available:

- `HysteresisGraphController`: This implements a simple bang-bang controller. To instantiate such a controller, as can be seen on lines 2–4 in listing 2, we need to specify which variable is used (`P_h` in this case) and which level variable is to be controlled (`heatingWater` in this case). Furthermore, we need to specify the lower and upper switching points as well as the power values in the “on” and “off” states. Finally, the last input determines the initial state.
- `PIGraphController`: Similarly to the hysteresis controller, we need to specify the variables involved when instantiating a proportional-integral (PI) controller, as we see on lines 5–6 in listing 2. Furthermore, the controller parameters need to be specified, i.e., the gain, the time constant, and the initial value, respectively. Lastly, the update time is specified.
- `FeedforwardGraphController`: This controller simply uses predetermined outputs for a given signal.

The simulation procedure is then schematically shown in figure 22. For every instance in time t, all time-based signals, such as disturbances, upper, and lower bounds can be evaluated by performing linear interpolation. Then, the desired control action can be obtained from all controllers.

The key component in the simulation procedure, which is highlighted in red in figure 22, then calculates all the flows in the graph, based on the graph structure and the desired control actions, while respecting all flow constraints. It does so by automatically formulating and solving an optimization problem. For the

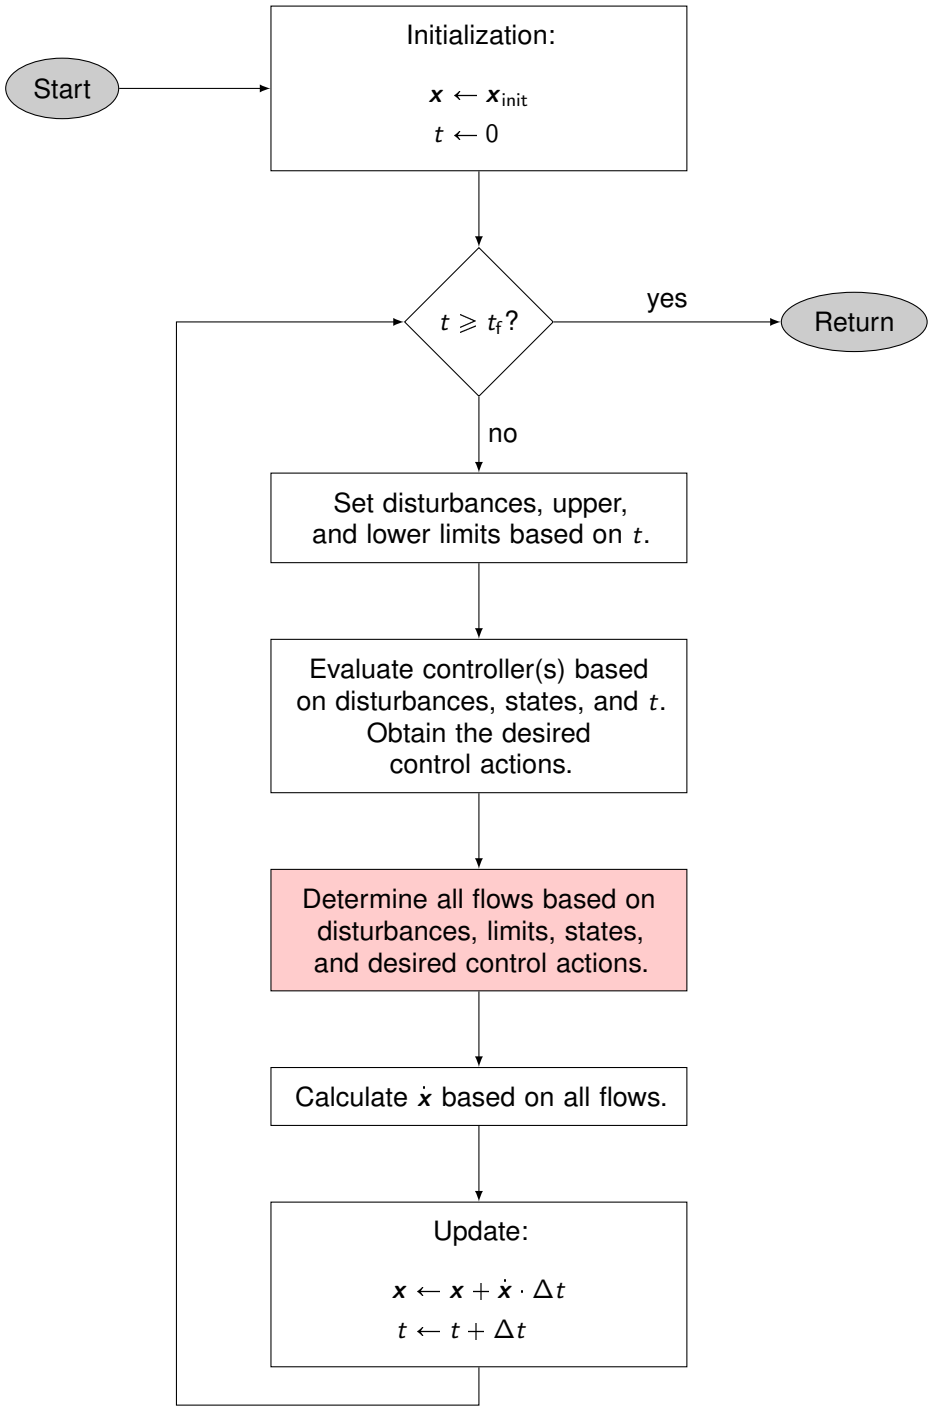

Figure 22: Flowchart of the simulation procedure based on the graph-based modeling approach.

Listing 3: Example Matlab code to perform an optimization based on the model defined in listing 1.

```
1 %% Optimize
2 heatingWater.chargeSustaining = true;
3 cabinNode.chargeSustaining = true;
4
5 mdl.objective = SmartFunctionHandle('objective', @(batteryEnergy, cabin) -batteryEnergy(end) ...
6     + sum((cabin - cabinNode.initialValue).^2 / 2e6));
7
8 optResults = mdl.solve(N, dt);
```

example in figure 19, the optimization problem looks as follows:

$$\begin{aligned}
\min_{\mathcal{U}} \quad & \left(\frac{P_\mathrm{h}-P_\mathrm{h,des}}{50000}\right)^2 + \left(\frac{u_\mathrm{heat}-u_\mathrm{heat,des}}{1000}\right)^2 \\
\text{s.t.} \quad & \text{DC link node:} \quad P_\mathrm{b} = P_\mathrm{h} + P_\mathrm{rest} \\
& \text{battery node:} \quad P_\mathrm{b,loss} = \frac{(U_\mathrm{b}-\sqrt{-4\cdot P_\mathrm{b}\cdot R_\mathrm{b}+U_\mathrm{b}^2})^2}{4\cdot R_\mathrm{b}} \\
& \text{edge constr.:} \quad \dot{Q}_\mathrm{loss} = k_\mathrm{w}\cdot(T_\mathrm{w}-T_\infty) \\
& \qquad\qquad\quad \dot{Q}_\mathrm{cons} = u_\mathrm{heat}\cdot(T_\mathrm{w}-T_\mathrm{cab}) \\
& \qquad\qquad\quad \dot{Q}_\mathrm{cab,loss} = k_\mathrm{cab}\cdot(T_\mathrm{cab}-T_\infty) \\
& \text{variable bounds,}
\end{aligned} \tag{19}$$

where the set of the optimization variables is given by

$$\mathcal{U} = \left\{P_\mathrm{b}, P_\mathrm{h}, P_\mathrm{b,loss}, \dot{Q}_\mathrm{loss}, \dot{Q}_\mathrm{cons}, u_\mathrm{heat}, \dot{Q}_\mathrm{cab,loss}\right\}. \tag{20}$$

The value of the numerators in the definition of the objective are taken from the nominal value of the respective variables. The reason why control actions are treated as “desired” actions and included in the objective, is that some constraints might make certain actions impossible. In the example in figure 19, for instance, a desired control action might not be realizable if the battery power is limited.

Finally, the resulting flows from the solution of (19) can be used to calculate the derivative of all states, which is finally used to perform state integration. In the current version of the framework, the Euler forward integration methods is used.

Note that constraints on state variables, such as upper and lower limits, for example on lines 28–29 in listing 1, cannot be enforced during simulations. In particular, a poorly tuned controller might, for instance, violate temperature bounds. Thus, the state trajectories need to be checked for their validity after the simulations.

A sample solution of a simulation of the model in figure 19 is given in figure 23.

2.2.2 Optimizing Graph-Based Models

Instead of simulating the system response controllers acting on the degrees of freedom, the system can also be optimized. For this purpose, an objective function needs to be defined. Listing 3 shows an example of how this can be set up in the code.

Apart from the objective in lines 5–6, some additional constraints are defined on lines 2–3. These constraints require the solution to have the same energy contents (temperatures) at the end and in the beginning of the optimization horizon. Note that constraints on state variables like these, but also upper and lower state constraints, can be enforced if the control inputs are optimized.

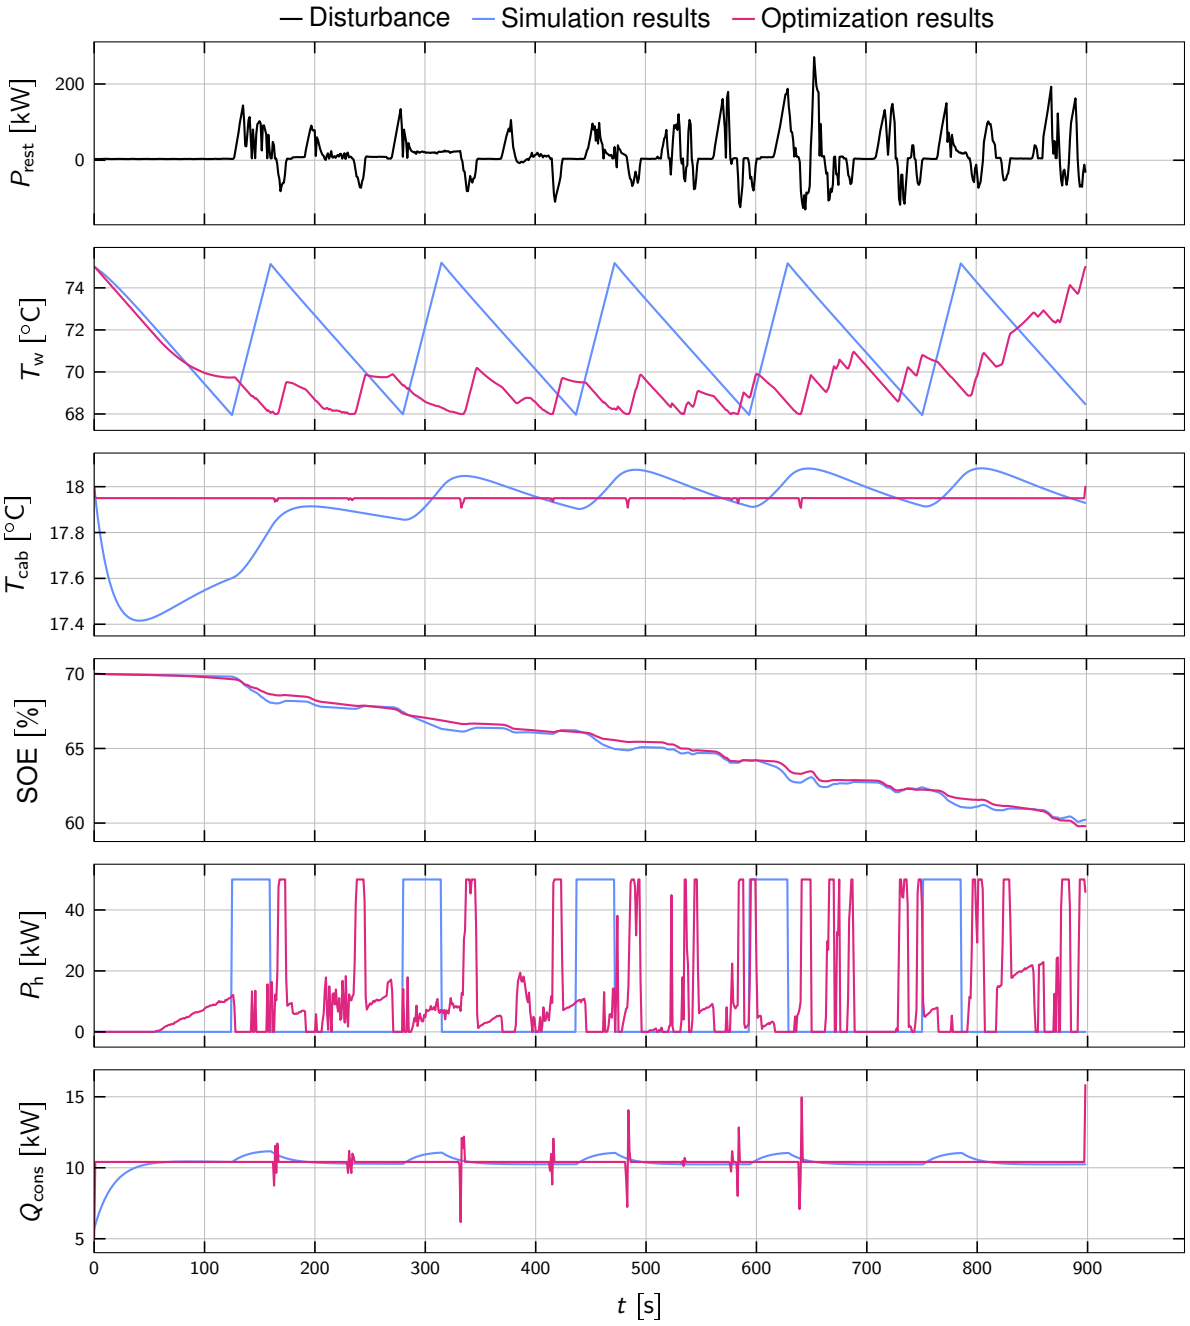

Figure 23: Sample results of a simulation (blue) and an optimization (magenta) of the model shown in figure 19.

Based on the structure of the graph-based model, the equations forming the OCP are automatically collected to form an OCP. For the example shown in listing 1 and figure 19, the following OCP results:

$$
\begin{aligned}
\min_{\substack{u \in \mathcal{U},\\ x \in \mathcal{X}}} \quad & -E_\mathrm{b}(t_\mathrm{f}) + \tfrac{1}{2\times 10^6} \cdot \textstyle\int_{t_0}^{t_\mathrm{f}} (E_\mathrm{cab}(t) - E_\mathrm{cab,ref})^2 \mathrm{d}t \\
\text{s.t.} \quad & \text{disturbance:} && P_\mathrm{rest} \\
& \text{DC link node:} && P_\mathrm{b} = P_\mathrm{h} + P_\mathrm{rest} \\
& \text{battery node:} && \tfrac{\mathrm{d}}{\mathrm{d}t} E_\mathrm{b} = -P_\mathrm{b} - P_\mathrm{b,loss} \\
& && P_\mathrm{b,loss} = \tfrac{(U_\mathrm{b} - \sqrt{-4 \cdot P_\mathrm{b} \cdot R_\mathrm{b} + U_\mathrm{b}^2})^2}{4 \cdot R_\mathrm{b}} \\
& \text{water node:} && \tfrac{\mathrm{d}}{\mathrm{d}t} E_\mathrm{w} = P_\mathrm{h} - \dot{Q}_\mathrm{loss} - \dot{Q}_\mathrm{cons} \\
& && E_\mathrm{w}(t_0) = E_\mathrm{w}(t_\mathrm{f}) \\
& \text{cabin node:} && \tfrac{\mathrm{d}}{\mathrm{d}t} E_\mathrm{cab} = \dot{Q}_\mathrm{cons} - \dot{Q}_\mathrm{cab,loss} \\
& && E_\mathrm{cab}(t_0) = E_\mathrm{cab}(t_\mathrm{f}) \\
& \text{edge constr.:} && \dot{Q}_\mathrm{loss} = k_\mathrm{w} \cdot (T_\mathrm{w} - T_\infty) \\
& && \dot{Q}_\mathrm{cons} = u_\mathrm{heat} \cdot (T_\mathrm{w} - T_\mathrm{cab}) \\
& && \dot{Q}_\mathrm{cab,loss} = k_\mathrm{cab} \cdot (T_\mathrm{cab} - T_\infty)
\end{aligned}
\tag{21}
$$

with the variables

$$u = \left\{ P_\mathrm{b}, P_\mathrm{b,loss}, P_\mathrm{h}, \dot{Q}_\mathrm{loss}, \dot{Q}_\mathrm{cons}, \dot{Q}_\mathrm{cab,loss}, u_\mathrm{heat} \right\} \in \mathcal{U}\,, \tag{22}$$

$$x = \{ E_\mathrm{b}, E_\mathrm{w}, E_\mathrm{cab} \} \in \mathcal{X}\,, \tag{23}$$

where the sets $\mathcal{U}$ and $\mathcal{X}$ represent the feasible sets of u and x based on the variable bounds, respectively.

The solution approach used to solve this OCP is the same that we have used in the publication [35]. It is explained in some detail in Appendix C in [35]. It can be summarized as follows:

1. Discretize the continuous-time signals with a constant sampling time to obtain a finite number of discrete variables.
2. Use forward Euler integration method to discretize the ordinary differential equations (ODEs) of all state variables.
3. Formulate an NLP using direct multiple shooting.
4. By using the CasADi Matlab interface [36], the variables and equations can be entered in a symbolic manner. The resulting problem is then automatically parsed and sent to Ipopt [37].

A sample solution of an optimization of the model in figure 19 is given in figure 23.

2.3 Outlook

We are confident that the developed toolchain targeted for graph-based models is a powerful tool for upcoming simulation and optimization studies. One element that might be missing in the current implementation is the treatment of integer (or: binary) decisions in the optimization of systems. While the numerical treatment of such decisions is generally tricky, we think that the developed framework offers enough freedom to include such decision variables in general.

3 Work Package 3: Simulations

Despite the title “Simulations”, this work package entails not only the numerical imitation of the evolution of a real-world system over time, but also some accompanying and related tasks. In fact, the contents of this work package can be understood as the counterpart to the fourth work package: While the fourth work package deals with how the *online* operation of a thermal system can be improved, this third work package treats questions that need to be answered “offline”.

In that regard, section 3.1 deals with the automatic identification and selection of representative simulation scenarios. In section 3.2, results of simulations using different vehicle types are presented. Section 3.3 shows a potential study, dealing with the task of battery lifetime extension by utilization of a thermal energy buffer. Finally, section 3.4 concludes this section with a brief outlook.

3.1 Scenario Identification based on Machine Learning

In the context of the “simulations” envisioned in this work package, a simulation “scenario” encapsulates all disturbances acting on the model over the course of the simulation. This includes, among others,

- the meteorological conditions, i.e., the ambient temperature, precipitation, and solar irradiation,
- the number of passengers in the bus,
- the door opening and closing times, and
- the velocity and altitude profile of the bus route.

In order to carry out simulation studies, suitable simulation scenarios must be selected. There is a conflict of objectives to be considered:

- On the one hand, there is the *reliability* of the statements gained from the simulations. In order to maximize this, one should simulate as many of the different scenarios which can occur during operation as possible, ideally even all scenarios that have ever occurred.
- However, this clearly contradicts the criterion of *efficiency*: In order to arrive at statements as quickly as possible, only as few scenarios as necessary should be simulated. However, it is not intuitively clear what this means in concrete terms, as can be illustrated by the following question: Is it sufficient to simulate one winter and one summer scenario each, or are the statements distorted by neglecting the intermediate seasons?

The general idea of the learning-based scenario selection is visualized in figure 24. By having a representative selection of simulation scenarios, results can be generated much more efficiently, compared to if all scenarios have to be simulated in order to get a result.

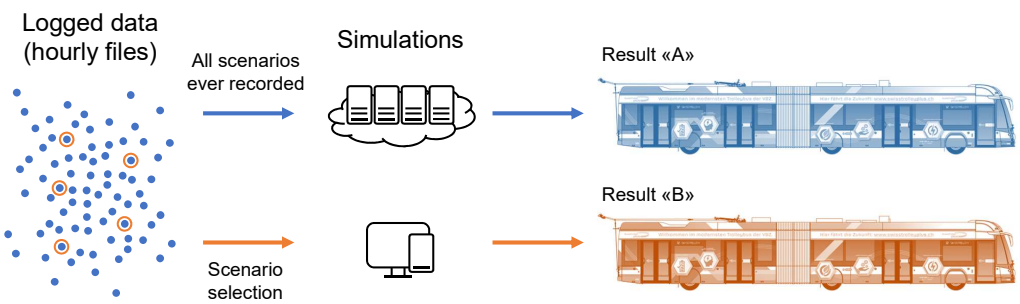

Figure 24: Concept of the learning-based scenario selection: The goal is to select a representative sample of all available scenarios, such that simulations with only the selected samples yield comparable results to simulations considering all samples.

The semester project of Ashutosh Pani [38], was the first attempt at dealing with this trade-off using classical machine learning approaches.

In a first step, hourly data from the following three sources are combined.

- Log database from the project SwissTrolley plus with more than 6000 h of usable data. This data contains, for example, the average indoor and outdoor temperatures, the average speed and the driven bus line.
- Meteorological data from the Federal Office for Meteorology and Climatology MeteoSwiss, via IDAweb[3]. This data includes outdoor temperature, precipitation and solar radiation.
- Data from the passenger counting system in the SwissTrolley plus, provided by VBZ.

The data set obtained in this way can be interpreted as a 16-dimensional data cloud, whereby the individual dimensions represent one of the above features each, for example, the average outside temperature, the average number of people on the bus or the average travel speed. Each point in this data cloud represents a one-hour scenario.

To prevent the “curse of dimensionality” during analysis of this data cloud, Stähelin refined the learning-based scenario selection toolchain in his semester project [39]. Instead of the original features, a “feature engineering” step was performed and the four main modes of heat transfer, i.e., heat input

- by solar energy $\dot{Q}_{\text{solar}}$,
- by passengers $\dot{Q}_{\text{passengers}}$,
- by door openings $\dot{Q}_{\text{doors}}$, and
- by conduction through the walls $\dot{Q}_{\text{walls}}$

are considered to span a four-dimensional feature space. The data is prepared for the next steps by balancing it over the year (i.e. the same number of data points for each month) by a combination of a heuristic and random oversampling. Since the four features all have the same physical dimension, a normalization step is not necessary.

Uniform manifold approximation and projection (UMAP), suggested by Leland et al. [40], can then be used to project this four-dimensional space onto two dimensions for visualization purposes. Figure 25 shows how the different features are located within this manifold.

Finally, a cluster analysis using the k-medoids algorithm has been conducted, which is used to identify structures in this data cloud and to group similar data points. Eventually, the centrally located data point (the medoid) of each cluster can be selected as representative for its cluster. Figure 26 shows the result of the clustering algorithm. Extreme scenarios are based on the extremum values of the single features.

For most clustering algorithms, the number of clusters k, and thus the number of representative simulation scenarios, must be specified as an input value. Thus, the scope of the found scenarios can be varied in the future depending on the application: If a simple, fast simulation study is to be carried out, it may be sufficient to select 5-10 scenarios; for a more comprehensive study, for example, 100-200 scenarios can be selected.

3.2 Steady-State Simulation Results

In this section, we present the results from the evaluation of the model briefly introduced in section 1.3. As mentioned, this model is explained in more detail in a conference publication [20], where also some of the results are shown.

3.2.1 Procedure

The general idea is to solve the optimization problem stated in (5) for a large number of environmental conditions that are representative for a year-round operation of the bus. For the vehicle-specific information

[3] `https://gate.meteoswiss.ch/idaweb/`

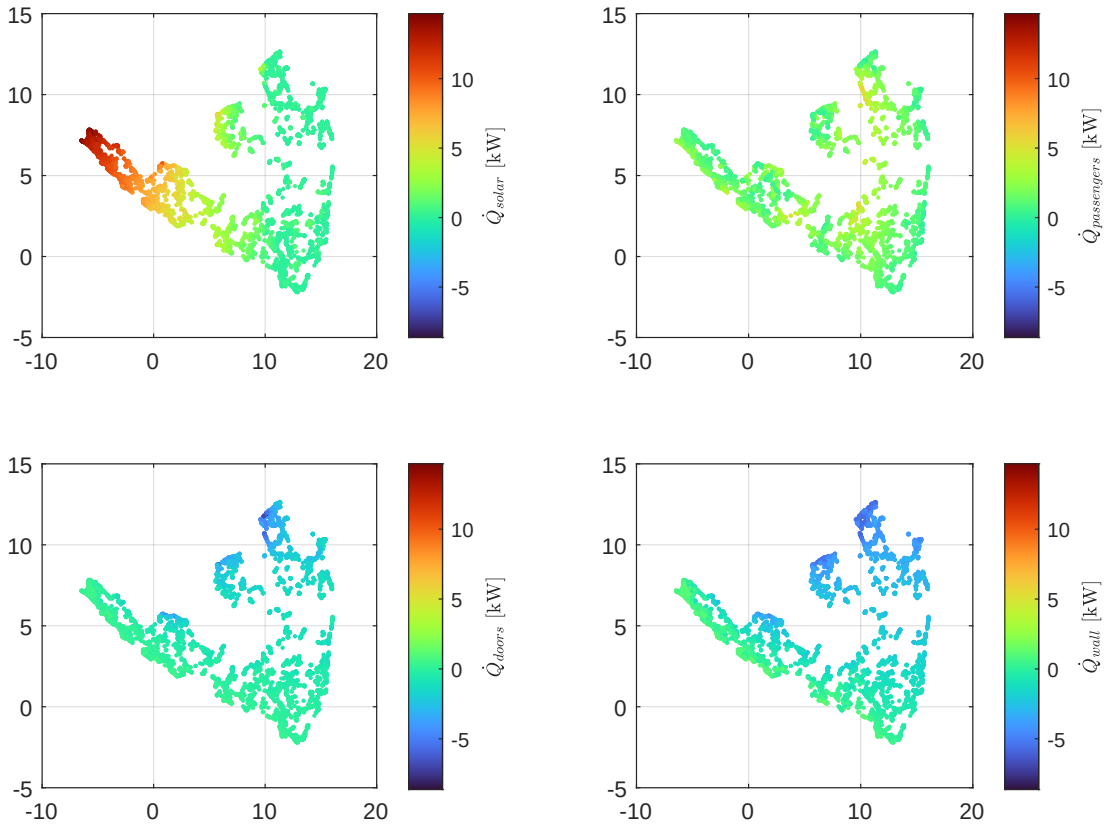

Figure 25: Low-dimensional representation of the four-dimensional feature space [39]. The x and y coordinates in this representation do not have any physical meaning. Thus, the color scales are used to locate the points within this two-dimensional representation.

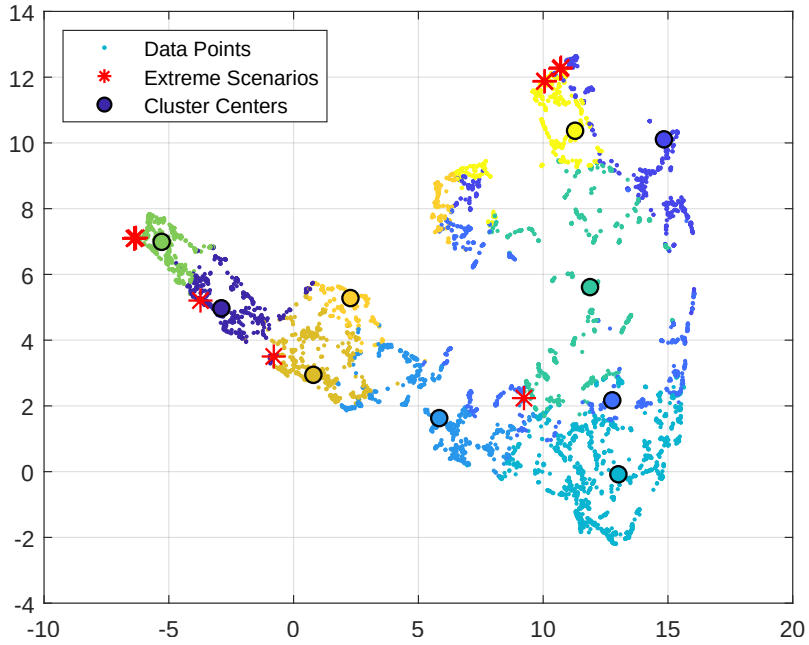

Figure 26: Result of the scenario selection tool [39]. The 10 cluster centers are selected as representative scenarios and extreme scenarios are included to represent outlier scenarios.

like the number of passengers and the status of the doors, we rely on data recorded on the SwissTrolley plus. We average this time-resolved data over operation periods of one hour. We extend this data with measurement data from MeteoSwiss, which we average over the same operation periods. Figure 27 provides an overview of the input data for the simulations presented in this section.

The scenarios obtained in this way are not uniformly distributed throughout the year. For instance, there are more than twice as many samples in May compared to June (see figure 27). Thus, we cannot directly average the results over the whole data set. Instead, we average the results over each month separately and then average over all twelve months.

3.2.2 Results

Using the simulation approach introduced above, we perform simulations with three different public transport vehicle types, i.e., electric buses, regional, and intercity trains. The parameters of all vehicles are given in tables 3 and 4.

Figure 28 visualizes the monthly average heat request and the corresponding power consumption for the electric bus model. No radiant heaters are used in this case to simplify the interpretation of the results. This overview shows that both the average as well as the peak heat demand is larger in winter compared to summer. In addition, cooling is needed in much fewer instances than heating. Together, these observations explains the fact that for our scenarios, the overall heating demand is about 8 times larger than the cooling demand. Furthermore, figure 28 shows that in spring, summer, and fall, often no active heating or cooling is necessary to meet the comfort requirements. In theses cases, the heat demand is zero, and thus, the HVAC system is considered “passive”. Of course, the narrower the allowed comfort range, the less often the HVAC system can be passive.

This trade-off between power consumption and thermal comfort is visualized in figure 29 for all three vehicle types. To generate the data for this figure, the allowed PMV is varied between $\Psi = 0$ (no variation, maximum comfort) and $\Psi \in [-2.0, 2.0]$ (least comfort). For each PMV window, all scenarios are evaluated for each vehicle concept. In this visualization, the most basic HVAC system is assumed, i.e., a purely

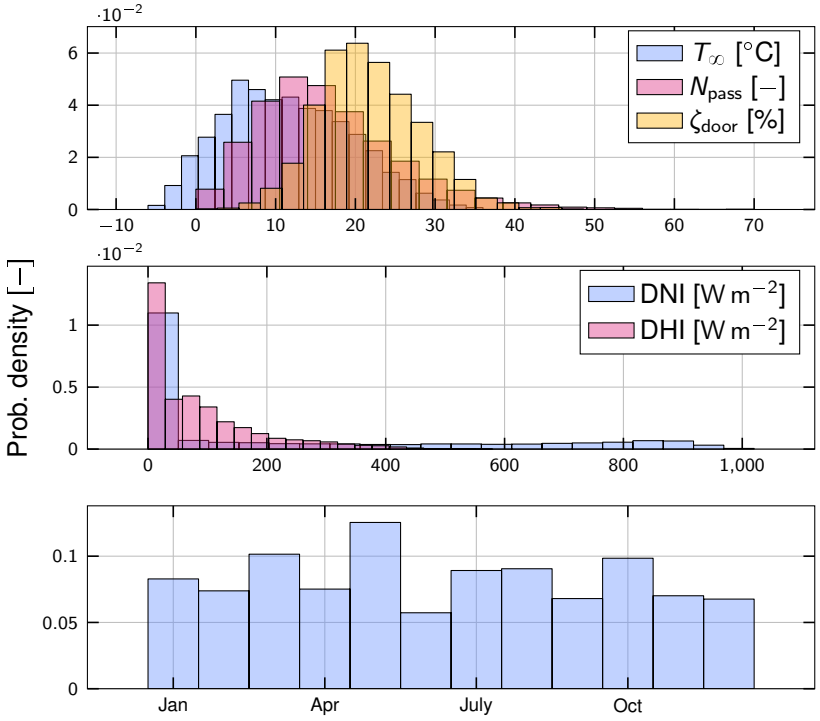

Figure 27: Overview of the 7500 scenarios used for the determination of the mean annual power consumption of the HVAC system [20]. The average door open fraction is represented by ζ_{door}. Solar irradiation is given in terms of direct normal irradiance (DNI) and diffuse horizontal irradiance (DHI).

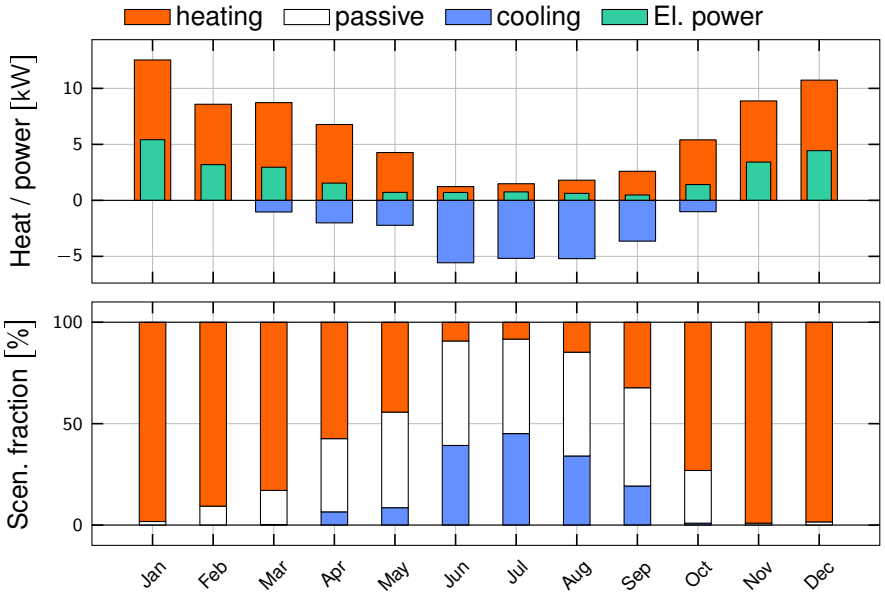

Figure 28: Visualization of the monthly heat and power consumption values required to achieve a PMV in the range $\Psi \in [-1, 1]$ (top) [20]. The bottom graph shows the fraction of the scenarios where heating or cooling is needed.

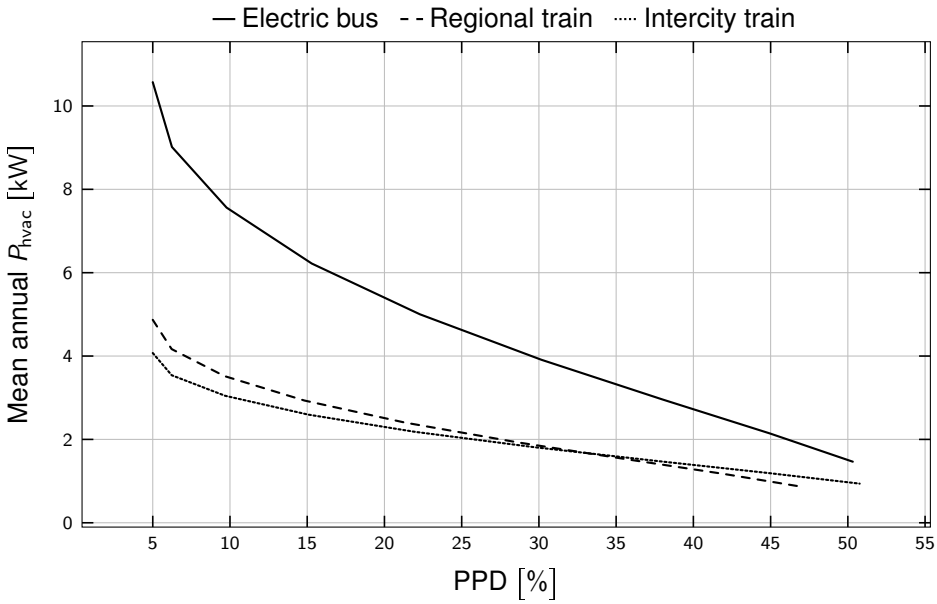

Figure 29: Pareto front for different vehicle types showing the trade-off between the conflicting objectives of power demand and thermal discomfort. Thermal discomfort is quantified using PPD. The HVAC systems used in these results consists of only PTC heaters.

electrical PTC heater with an efficiency of 100%, i.e., $\eta_{cop} = 1$. As expected from the definition of the PPD in [5], a value below 5% can never be reached in practice. The lower the PPD, the higher the power demand. Furthermore, the bus uses the most power, which is to be expected, since it has the lowest insulation and the largest doors which open more frequently than for the other vehicles (see table 4).

Figure 30 shows the achievable reductions in the mean power consumption by performing certain changes in the HVAC system. The following changes are evaluated:

Heat pump: Usage of a heat pump instead of the PTC heating elements. The COP of the heat pump is given in figure 10.

Air curtain: Installation of an air curtain to reduce door losses. Such a device has been found promising in a recent study by Sidler et al. [41]. Panthak et al. conducted a simulation in a mini-bus, where an air curtain was able to reduce door losses by 59% [42], which is the assumption for the simulations conducted in this section.

Insulation improvement: The second option to lower the thermal loads is to improve the insulation of the vehicle hull. For this case, a reduction of the mean thermal conductivity by 20% was assumed.

Radiant heaters: Through the usage of radiant heaters, the mean radiant temperature felt by passengers can be increased in (1). Therefore, the air temperature can be lowered to achieve the same level of comfort. However, depending on the environmental conditions, the power consumption might be lower if the infrared heaters are not used. This is particularly true in summer conditions. We assume that for each scenario, the radiant heaters are only used if they lower the overall power consumption. For the radiant heater surface temperature, we assume $70\,^{\circ}\mathrm{C}$. The layout of the radiant heaters in the cabin ceiling is shown in figure 31, along with the resulting mean radiant temperature of 30 passengers. This radiant heater layout corresponds to a coverage of the ceiling area of 16.5%. For the other vehicles, a radiant heater layout with the same coverage was assumed.

From figure 30, we see that the installation of a heat pump yields the largest benefits for all vehicle types. Throughout the the comfort range, the power consumption can be reduced by 40% to 60%. For electric

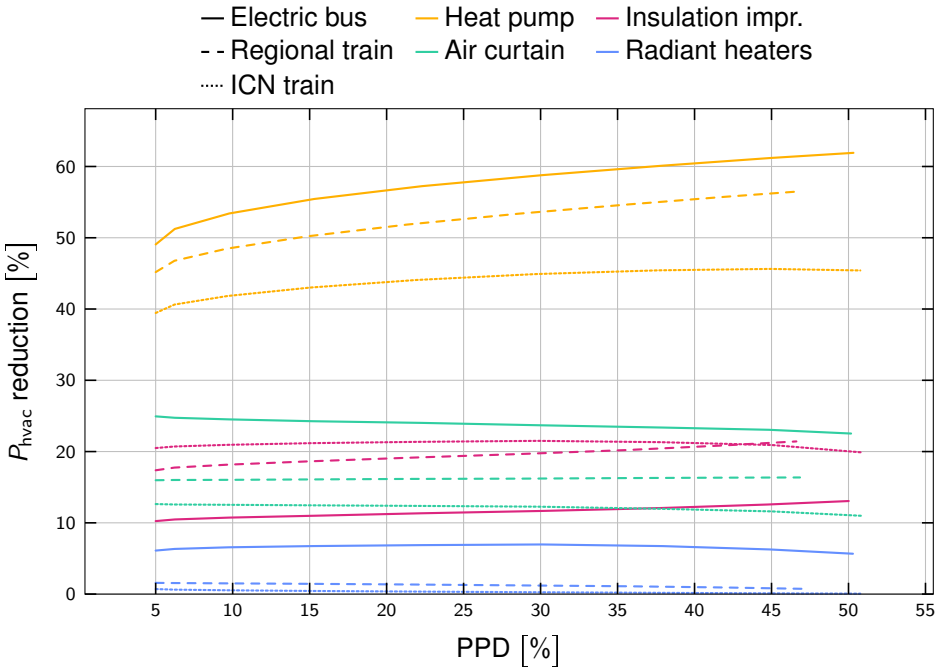

Figure 30: Potential reductions of the mean HVAC consumption by changes in the HVAC system for all three vehicle concepts. A HVAC system with PTC heating elements serves as the baseline for these improvement calculations.

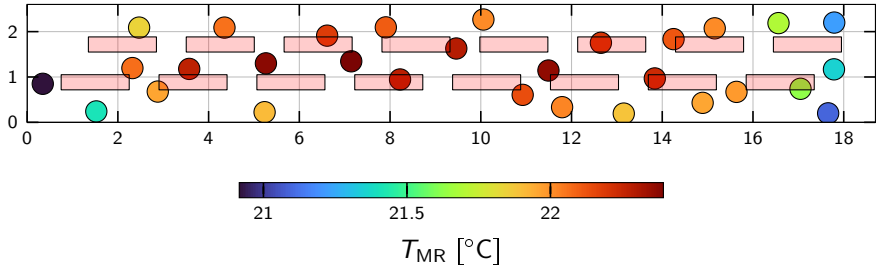

Figure 31: Visualization of the mean radiant temperature throughout the bus cabin (black outline) subject to radiant heater surfaces in the ceiling (light red) [20].

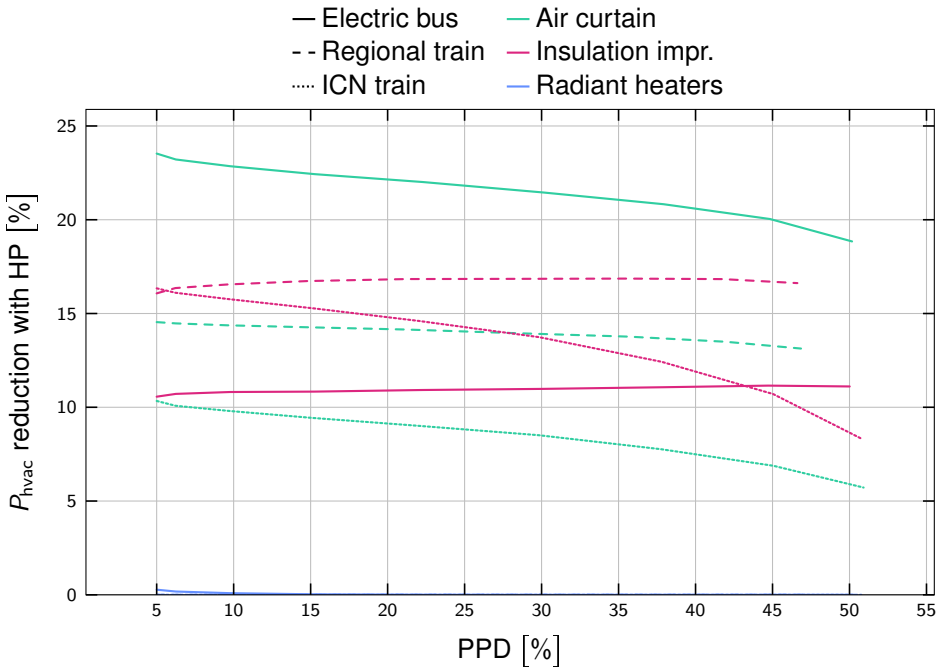

Figure 32: Potential reductions of the mean HVAC consumption by changes in the HVAC system for all three vehicle concepts. A HVAC system with a heat pump serves as the baseline for these improvement calculations.

buses, the installation of air curtains yields the second largest benefit, as it has large doors that open frequently. For the two train wagons, a 20% improvement of the hull insulation yields larger benefits than an air curtain. Finally, the radiant heaters yield the largest benefits for a city bus. The reason for this observation is that buses have large door losses on the one hand, and low insulation values on the other. Both effects benefit infrared heaters, as the cabin air temperature can be reduced. For the electric bus, improvements of 6% to 7% are expected with infrared heaters, while for the train wagons, the improvement potential is only in the order of 1%.

As heat pumps are becoming increasingly common in public transport vehicles, figure 32 evaluates further improvements of the HVAC system with respect to a heat pump-based system. Similarly to can be observed in figure 30, buses profit heavily from air curtains, while trains benefit more from improved insulation, with the expected improvements in a similar range as with PTC heaters. Radiant heaters, on the other hand, do not offer any significant benefits when coupled with a heat pump. While the usage of radiant heaters allows to reduce the cabin air temperature to achieve the same thermal comfort, which results in lower heat losses to the environment, this comes at the expense of the power necessary to sustain the radiant heater temperature, which can usually not be compensated by a reduction in the heat pump power due to its high COP. While we have tested many radiant heater configurations (size and temperature) to supplement the heat pump, we have found no configuration where a reduction in power consumption of more than 0.3% can be achieved for any of the vehicles. Considering the added complexity of radiant heater installation and the investment costs, we consider radiant heaters not a viable extension to a heat pump. This observation is also confirmed in simulations with lower and higher ambient annual mean temperatures.

A sensitivity study on these results showed that the annual energy consumption is about ten times more sensitive towards the heating COP, compared to the cooling COP, due to the lower cooling requirements shown in figure 28. This observation further encourages the use of alternative refrigerants, which offer

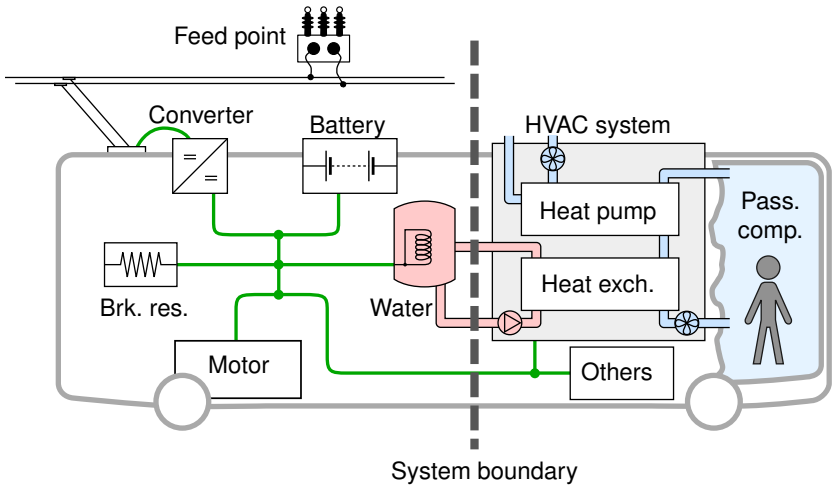

Figure 33: Schematic overview of the battery-assisted trolley bus under consideration [35]. The braking resistor can be used to dissipate excess recuperation energy. The air in the passenger compartment is heated by a heat pump, which extracts heat from the ambient air, and by the electrically heated water (shown in red) via an additional heat exchanger. All other auxiliaries (infotainment system, lights, etc.) are summarized in “Others”. Air paths are shown in blue, electric connections are shown in green. The vertical dashed line symbolizes the system boundary: Everything to the right of this line is considered a disturbance to the system. The power flows are shown in figure 34.

advantages in heating efficiency at a disadvantage in cooling efficiency. CO_2 (R744) which is used in the project Swiss eBus Plus, is such a refrigerant.

3.3 Combined Control and Sizing Optimization of the Water Heater System

For a scientific publication, we performed a detailed analysis on the potential benefits of a combined consideration of traction and thermal energy management [35]. Specifically, we analyzed the potential of including a thermal reservoir in the traction energy management by considering the trade-off between energy efficiency and battery degradation. The system considered for this purpose is visualized in figure 33. As we can see schematically in figure 33 and in more detail in figure 2, the hot water is used to support the heat pump in providing heat to the cabin (e.g., if the capacity of the heat pump is not sufficient for the demand). Additionally, the hot water is used to provide heat to the driver’s compartment. As visualized in figure 33, for the analysis presented here, we consider the amount of heat consumed by the HVAC system a disturbance. Thus, the only thing we can influence is “when” we heat the water, but not “how much” we heat it.

The battery’s SOH is a dimensionless quantity represented by $\Psi(t)$, which characterizes the current state of the aging process of the battery in terms of capacity degradation. It is typically defined as the fraction of the remaining charge capacity $Q_{\mathrm{b}}(t)$ to the initial charge capacity $Q_{\mathrm{b,nom}}$,

$$\Psi(t) = \frac{Q_{\mathrm{b}}(t)}{Q_{\mathrm{b,nom}}}\,. \tag{24}$$

The main factors affecting the aging process of lithium-ion batteries are the C-rate, the state of charge (SOC), the depth of discharge (DOD), and the cell temperature. We base the analyses on the battery degradation models of the open-source project OpenSesame [43] by the BFH. This model describes the

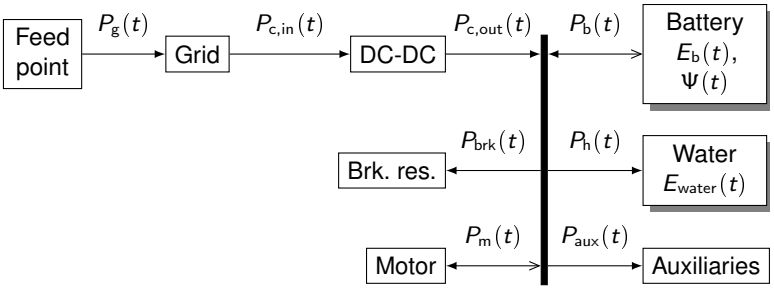

Figure 34: Overview of the electric power flows present in the system introduced in figure 33 [35]. Filled arrows indicate the positive direction of power flows. Open arrows indicate that power flows in the reverse direction are possible. The vertical black bar represents the DC link where the power balance between consumers and suppliers is established. Blocks with a drop shadow represent dynamic components, the state variables of which are given in the block descriptions.

degradation of battery cells based on all aforementioned influence factors, where a rainflow counting algorithm is used to decompose a given battery usage cycle into an equivalent set of subcycles. The degradation is then calculated as a superposition of cycle and calendar aging.

An energy management system of a battery-assisted trolley bus as shown in figure 33 can affect battery aging mainly in three ways, which are listed in the following. Both of these effects are visible when analyzing the optimal power split values as shown in figure 35.

1. Generally, the heating is preferred in recuperation phases, i.e., when the requested power is negative. This can be seen at Ⓐ This can reduce the battery charging power.
2. The grid can be used more often to prevent battery operation points that are particularly damaging (e.g., high C-rates). For example, in a particular range of requested power, the battery is not used by the optimal control strategy in order to prevent degradation, as indicated by Ⓐ. Consequently, the power request is entirely covered by the grid, as shown by the the converter power values lying on the identity line.
3. Strong recuperation C-rates can be prevented by dissipating some energy in the braking resistor. The "strong" battery degradation limit enforces the use of the braking resistor to dissipate recuperation power in order to mitigate the battery aging process, as indicated by Ⓑ.

The Pareto fronts between the energy consumption and the battery lifetime for various driving missions are shown in figure 36. For convenience, the battery lifetime is given in a driving distance before the battery needs to be replaced,

$$d_b = (1 - \Psi_{EOL}) \cdot \frac{d}{\Delta\Psi}, \tag{25}$$

where d and $\Delta\Psi$ represent the driven distance and the battery degradation in a certain time horizon, respectively. $\Psi_{EOL} = 80\%$ represents the SOH at the battery's end-of-life (EOL). In this representation, the goal is to push a solution towards the bottom right corner (i.e., long battery life and low energy consumption).

It is clear that, on all driving missions, the improvement from the dotted to the dashed solution in figure 36 is significant. This corresponds to the improvements achievable when including the thermal buffer in the energy management. For example, due to an optimal heating control strategy, a realistic target lifetime of $d_b = 1.2 \times 10^6$ km can be reached on Routes 46 and 33 with a reduction in energy consumption of 2.5% and 6.7%, compared to a bang-bang heating strategy, respectively. On the other hand, the improvement between the dashed and the solid line is only marginal. This corresponds to the benefits achievable by optimally sizing the thermal buffer. We thus conclude that the sizing of the buffer in the ISOTHERM prototype vehicle is not too far from the optimum.

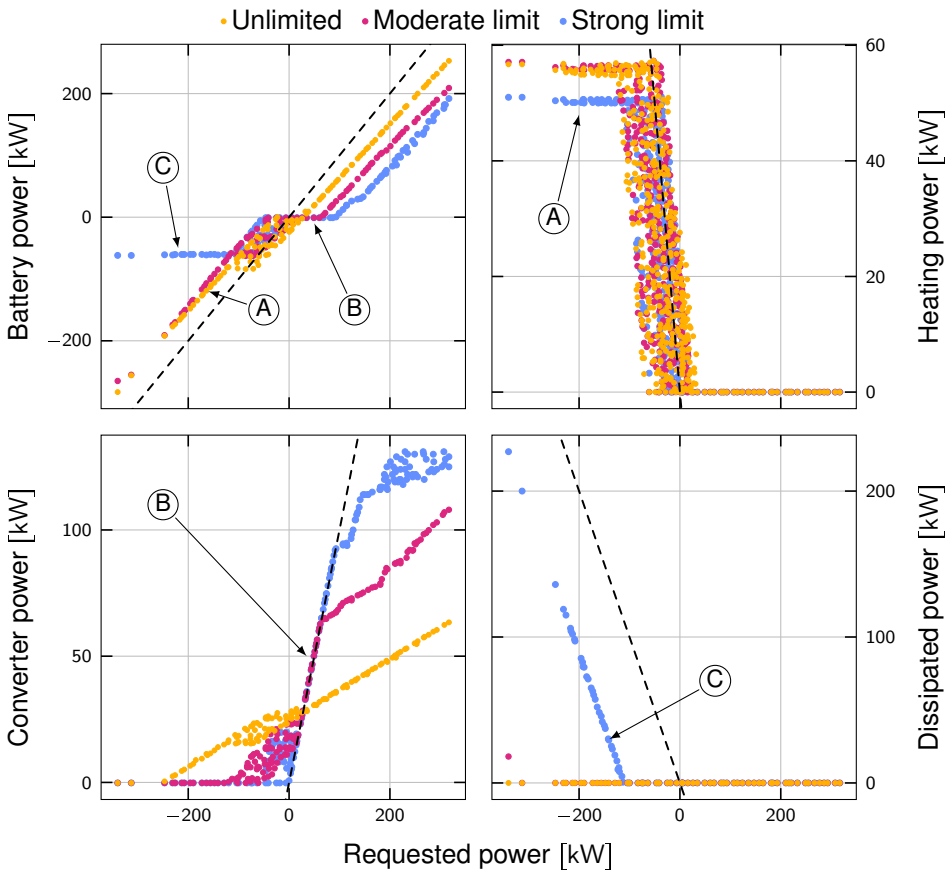

Figure 35: Power distribution at the DC-link (see figure 34) obtained by the OCP for Route 46 subject to various battery degradation limits [35]. The dashed black lines represent 45° lines.

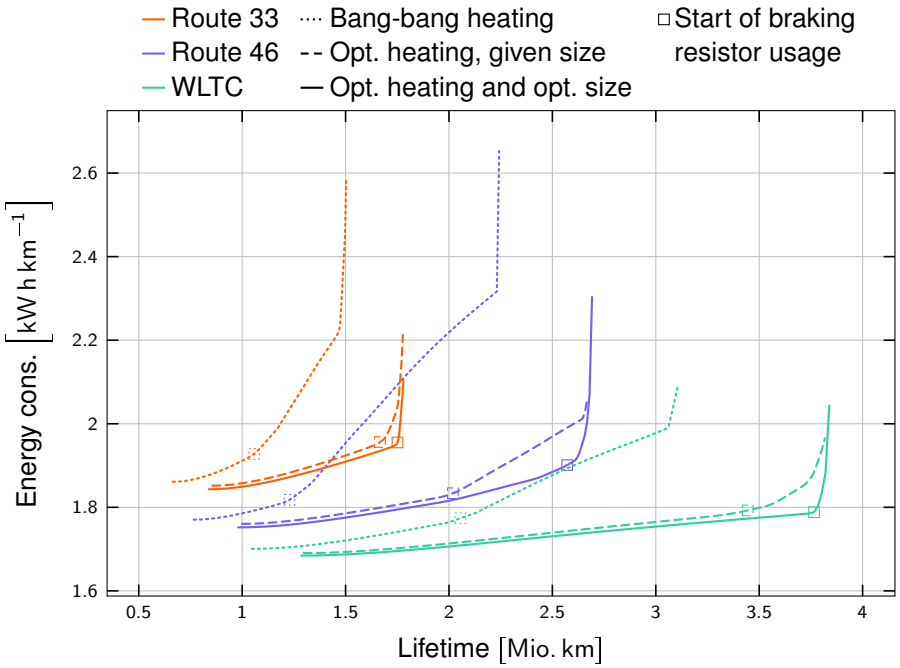

Figure 36: Resulting Pareto fronts for three different driving missions [35]. The dotted lines represent solutions of the baseline implementation with bang-bang heating, the dashed lines represent the solution with optimal heating for the specific vehicle design, and the solid lines represent the optimal sizing and control solution.

Specifically, the results of the design optimization reveal that high buffer power values are important. A value of at least 50 kW should be targeted, which is what is currently installed. This is an important finding, since it is well above the minimum value to satisfy thermal requirements. On the other hand, high buffer energy capacity is not important – the installed heating water capacity is already enough.

We conclude this section by stating that coupling thermal and traction is very important and not “optional”, as originally suggested in the funding application.

3.4 Outlook

The evaluation of the energetic benefit of the thermal coupling between the cooling circuit and the VCC responsible for HVAC of the passenger compartment is still an open task. Due to the change of refrigerant in the project Swiss eBus Plus to the more future-oriented CO_2 (R744), as mentioned in section 1.1.4.4, we think that it makes more sense to analyze this question in detail with the new system.

4 Work Package 4: Online Control

The goal of this work package is to develop and test algorithms that can improve the performance of the bus during real-world operation. Section 4.1 explains a concept for the optimized control of the water heater, which has been implemented and successfully tested on a prototype bus. Section 4.2 introduces a concept for the optimized control of a hybrid propulsion system (such as the one of a battery-assisted trolley bus), where the battery degradation is balanced against energy efficiency, such that a target battery lifetime can be achieved. As this last work package is the final work package originally proposed in the description of the project ISOTHERM, some of its envisioned tasks are not completed yet. These are listed in the outlook in section 4.3, along with some initial ideas for further proceeding in the project Swiss eBus Plus.

Not that work concerning disturbance characterization, estimation, and prediction is explained in the first work package, in section 1.2, although this work is also closely related to online control.

4.1 Online Water Heater Control

This section deals with the development and implementation of an intelligent water heater controller. The theoretical development of the control strategy, including the mathematical foundations, are provided in section 4.1.1. Section 4.1.2 explains the practical implementation. Finally, section 4.1.3 shows results from real-world operation.

4.1.1 Theoretical Development

The idea of optimized water heating control, which will be presented in the following, has been initially evaluated in the Master's thesis by Speckien [44]. Just as in section 3.3, we consider the amount of heat consumed from the hot water system a disturbance. Thus, we deal with the question "when" to heat the water, and not "how much" heat is consumed by the HVAC system.

The water heater on the bus consists of a simple resistive heating element. Additionally, it includes a shut-off relay which disconnects the voltage, and thus prevents heating when the water temperature exceeds a set temperature T_{max}. After the temperature drops below a second temperature setpoint T_{min}, the relay again connects the voltage and thus allows further heating. Using this hysteresis control, the water temperature is held within $[T_{\mathrm{min}}, T_{\mathrm{max}}] = [68\,^{\circ}\mathrm{C}, 75\,^{\circ}\mathrm{C}]$. Obviously, this controller does not consider the state of the rest of the vehicle, most notably the traction system.

In practice, the heating power can be regulated using pulse-width modulation (PWM) to any value between

$$0 \leqslant P_{\mathrm{h}} \leqslant \frac{U_{\mathrm{b}}^2}{R_{\mathrm{water}}}\,. \tag{26}$$

A more intelligent control algorithm would prefer heating the water when energy is available from braking energy recuperation. On the other hand, if the energy demand of the traction system is large, the heating system should be temporarily disabled. A visualization of such a strategy is given in figure 35. As we show in section 3.3, such a heating strategy could reduce the amount of energy cycled through the battery, thus reducing both losses as well as battery wear.

Figure 34 gives an overview of the power flows involved. With the goal of minimizing overall energy consumption, the water heating control task can be included in the energy management problem, i.e.,

$$\begin{aligned} \min_{P_{\mathrm{c,out}},\, P_{\mathrm{h}}} \quad & \int_{t_0}^{t_f} P_{\mathrm{g}}\,\mathrm{d}t \\ \text{s.t.} \quad & x_{\mathrm{SOE}}(t_0) = x_{\mathrm{SOE}}(t_f) \\ & T_{\mathrm{water}}(t_0) = T_{\mathrm{water}}(t_f) \\ & \text{system dynamics and constraints}\,. \end{aligned} \tag{27}$$

In various simulations experiments, Speckien was able to show empirically that the influence of the water heating strategy $P_\mathrm{h}(t)$ on the charging strategy $P_\mathrm{c,out}(t)$ is very small [44]. This is most likely due to the fact that

- the magnitude of the water heating power P_h is typically a lot smaller than the rest of the consumed power P_rest, and
- the amount of energy stored in the battery is significantly larger than the amount of (thermal) energy stored in the hot water circuit.

The optimal control problem (27) can thus be solved in a cascaded fashion. It is thus divided into the "standard" energy management problem for a battery-assisted trolleybus, introduced in [45], i.e.,

$$\begin{aligned} P^*_\mathrm{c,out}(t) = \arg\min_{P_\mathrm{c,out}} \quad & \int_{t_0}^{t_f} P_\mathrm{g}\,\mathrm{d}t \\ \text{s.t.} \quad & x_\mathrm{SOE}(t_0) = x_\mathrm{SOE}(t_f) \\ & \text{battery dynamics and constraints}\,. \end{aligned} \tag{28}$$

and, additionally, the water heating control problem,

$$\begin{aligned} \min_{P_\mathrm{h}} \quad & \int_{t_0}^{t_f} P_\mathrm{loss,bat}\,\mathrm{d}t \\ \text{s.t.} \quad & T_\mathrm{water}(t_0) = T_\mathrm{water}(t_f) \\ & P_\mathrm{c,out}(t) = P^*_\mathrm{c,out}(t) \\ & \text{water temperature dynamics and constraints}\,. \end{aligned} \tag{29}$$

We introduce the following assumptions about the optimal control problem (29):

- The hot water system is modeled using only one lumped temperature T_water, representing the whole circuit.
- The dynamic equation of the hot water temperature only depends on the heating power P_h, the "consumed" heat $\dot{Q}_\mathrm{cons}$, and some losses $\dot{Q}_\mathrm{loss}$,

  $$\frac{\mathrm{dT_{water}(t)}}{\mathrm{d}t} = \frac{1}{mc}\cdot\left(P_\mathrm{h} - \dot{Q}_\mathrm{cons} - \dot{Q}_\mathrm{loss}\right)\,, \tag{30}$$

  however, for simplicity, we assume that the losses are not dependent on the water temperature. Of course, this is not completely true, but keeping in mind that the temperature varies only in a small window $[68\,^\circ\mathrm{C}, 75\,^\circ\mathrm{C}]$, this simplification does not induce major errors.
- The consumed heat $\dot{Q}_\mathrm{cons}$ is perfectly controlled by the roof and front box, respectively. Specifically, with correct setting of valves and blower speed, $\dot{Q}_\mathrm{cons}$ can be selected independently of the current water temperature T_water.

Applying PMP and using the simplifications introduced above, we can derive the optimal heating power as

$$P^*_\mathrm{h}(t) = P_\mathrm{c,out}(t) - P_\mathrm{rest}(t) + \frac{U^2_\mathrm{b,OC}}{4R_\mathrm{b}}\left(1 - \left(\frac{1}{1-\lambda}\right)^2\right)\,, \tag{31}$$

with the co-state λ being a constant [44]. An example of such a trajectory is shown in figure 37.

Determining the constant value of the co-state λ requires perfect knowledge of the future and is thus not a strategy which can be applied in real life. In the energy management of hybrid electric vehicles (HEVs), a similar problem arises, which is often solved by introducing a PI controller which continuously adapts the value of λ to track a state trajectory [46]. This approach is called Adaptive equivalent consumption minimization strategy (ECMS).

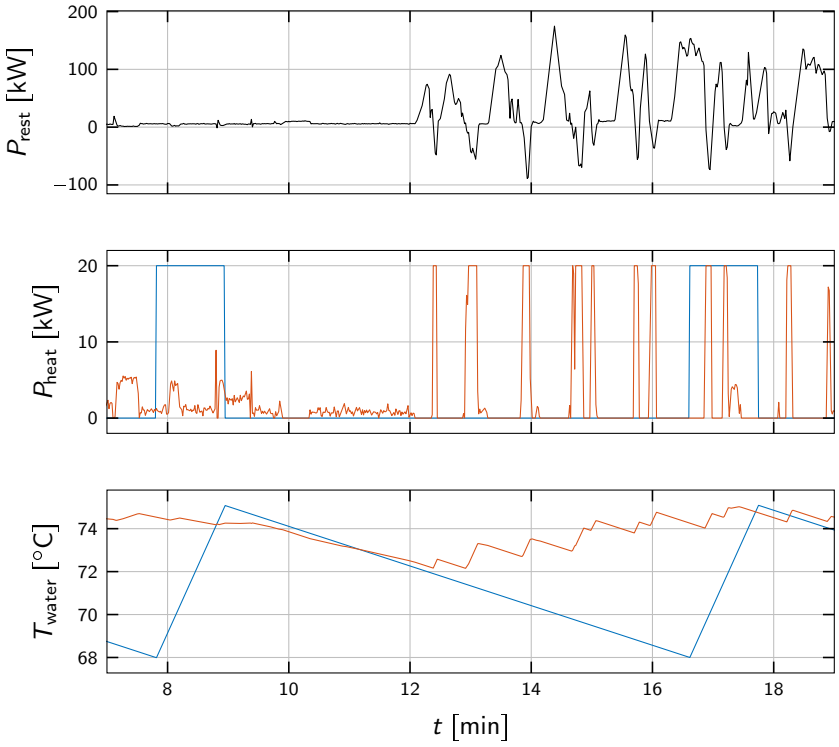

Figure 37: Simulated comparison of the optimal control (“offline-optimal” solution with perfect information) of the water heater (orange) with a hysteresis control (blue) [44]. In both cases, the water temperature (bottom) remains within clearly defined limits. The optimized heat output (center) responds intelligently to the current power requirement of the traction system (top): If energy is recuperated in the traction system (negative power), part of it is used for heating.

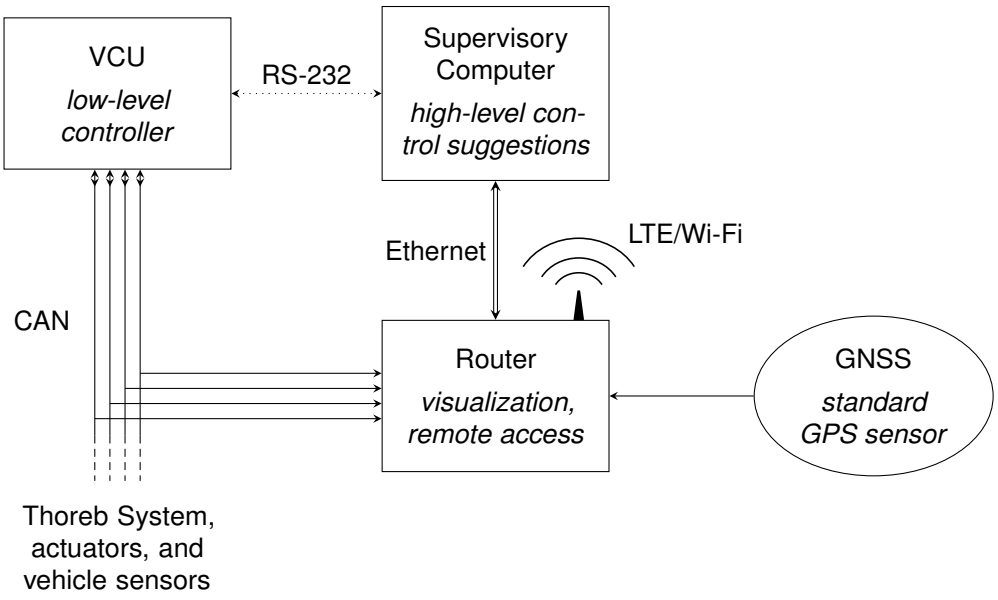

Figure 38: Hardware components on the vehicle and the corresponding communication channels. The high-level control implemented in Java runs on the “Supervisory Computer” and sends suggestions over the RS-232 interface to the low-level VCU, which in turn sends set points and commands over multiple controller area network (CAN) buses. The Router provides GNSS data for the Supervisory Computer via the Ethernet connection and allows remote access over mobile internet.

Following this idea Speckien suggests introducing a PI controller constantly adapting the value of λ to keep the temperature T_{water} between the limits T_{min} and T_{max} [44]. In simulation studies, Speckien showed that this approach can yield the following benefits, compared to the current hysteresis controller.

- 0.3–0.4% less consumed grid energy
- 10–11% reduction in battery wear

4.1.2 Practical Implementation

The implementation is explained in the following. First, we explain the implementation on the high-level side in section 4.1.2.1, which has been implemented in Java by ETH researchers. This high-level sends suggestions to the vehicle control unit (VCU), which is implemented by engineers of Carrosserie HESS AG. Some safety and fallback mechanisms implemented on that level are explained in section 4.1.2.2. Figure 38 provides an overview of the involved computing hardware.

4.1.2.1 High-Level Implementation

The code for the high-level implementation is written in Java. For this purpose, the modular software framework that was originally developed in the SwissTrolley plus project and is briefly documented in, e.g., [45] was extended with the necessary communication channels. Then, a new “module” was implemented, the program structure of which is given in the following list:

1. Handle override commands:
   Check if any “override commands” have been received and adapt the internal state accordingly. These commands can be sent at any time over secure shell (SSH) and transmission control protocol (TCP), for example if the controller does not perform as expected. This additional opportunity to intervene is important in the early stages of development.

2. Determine state of heater relay:
   Check if the relay is open, i.e., the controller has no control authority.
3. Update the equivalence factor:
   The current temperature "error" ΔT_{curr} from the target value is calculated. This error is integrated into ΔT_{int} if and only if the controller has control authority, i.e., all of the conditions below are fulfilled:
   - The suggestions are reportedly being used by the VCU.
   - The relay in the water heater is closed.

   The equivalence factor λ is then updated using the initial value λ_0 and a discrete-time implementation of a PI controller:
   $$\lambda = \lambda_0 + k_{\text{P}} \cdot \Delta T_{\text{curr}} + k_{\text{I}} \cdot \Delta T_{\text{int}} \tag{32}$$
4. Calculate optimal heating power and heating fraction:
   Using the analytical solution to the local Hamiltonian minimization problem (31), the optimal heating power can be calculated analytically. By first calculating the maximum possible heating power, calculate the optimal PWM fraction.
5. Publish suggestion:
   The calculated values are used to send an HVAC suggestion message to the VCU. The message contains the following fields.

   `optimalHeatingFraction`: The optimal heating fraction f^*_{PWM} according to (26), (31),
   $$f^*_{\text{PWM}} = \frac{P^*_{\text{HW}}}{P_{\text{HW,max}}} = \frac{P^*_{\text{HW}} \cdot R_{\text{HW}}}{U^2_{\text{bat}}}\,, \tag{33}$$
   where the software limits this value to $0 \leqslant f^*_{\text{PWM}} \leqslant 1$.

   Experience shows, that during recuperation phases, we often use the full heating power available. As the maximum heating power depends on the DC voltage according to (26), it is more robust and simpler to provide the suggestion in terms of a PWM fraction instead of directly providing the desired heating power P_{HW}.

   `heatingSensitivity`: This denotes the sensitivity γ of the optimal heating fraction with respect to the requested power, i.e.,
   $$\gamma = \frac{\partial f_{\text{PWM}}}{\partial P_{\text{rest}}}\,. \tag{34}$$
   This value must be calculated with special care if the optimal PWM fraction is either 1 or 0: In this case, γ is set to 0.

   `powerRequestBaseline`: This is the value of the power request P^*_{rest}, which was used during Hamiltonian minimization.

   `hvacSuggestionStatus`: A single 8-bit integer denoting the validity of the suggestion. The suggestion should only be considered by the VCU, if this is set to the value 1.

4.1.2.2 Low-Level Implementation

The following paragraph first explains how the suggestions introduced above are "translated" to set points in the VCU. Further below, we introduce the safety and fall-back implementation on the VCU.

The suggested phase change material (PCM) setpoint f^*_{PWM}, the sensitivity γ, and the power request baseline P^*_{rest} are used on the VCU to calculate the PWM set point as follows:
$$f_{\text{PWM}} = f^*_{\text{PWM}} + \gamma \left(P_{\text{rest}} - P^*_{\text{rest}}\right)\,. \tag{35}$$
Using this implementation, the VCU can adapt the optimal heating fraction to sudden changes in power request P_{rest}.

Due to the temperature-controlled relay, it is impossible to overheat the hot water circuit. The safety measures on the VCU-side thus only have to ensure that the temperature stays high enough, i.e., that a certain lower temperature limit is respected. The implemented approach is the following.

- Start a timer as soon as the temperature falls below 68 °C. Reset this timer, whenever the temperature is above that limit.
- If the timer reaches 30 s, i.e., the temperature has been too low for at least 30 s, start overriding the suggestions.
- Stop overriding the suggestions, once the temperature reaches 71.5 °C.

We were able to test the above-mentioned safety mechanisms using a special software version on the high-level side which was heating too much or too little on purpose. In these tests, we could confirm that both the safety heating shutdown at too high temperatures and the one at too low temperatures work as intended.

4.1.3 Results

The reliability tests and the tests of the fallback mechanisms could be finished by about June 2022. Since then, as of December 2022, the water heater controller has been active for more than 1100 h of operation without any issues. Figure 39 shows a sample recording of the control algorithm in action on the prototype trolley bus. It shows that the water temperature is kept relatively constant although the heating power is frequently adapted based on the power request by the remaining consumers (P_{rest}).

An comparison of battery currents between the operation with and without the adaptive water heating is shown in figure 40. This comparison indicates that high charging and discharging currents in the battery can be reduced by the adaptive heating strategy. However, there may be other factors that could influence the observed behavior, such as

- ambient temperature,
- driving style,
- passenger volume.

To equalize such random factors, it is important to observe the operation for longer periods than just one day.

Therefore, we analyze the operation of the bus over several months, the results of which are shown in figure 41. For this analysis, we consider only full trips (i.e., starting in the depot until returning to the depot) where the bus was operated on Route 83 to ensure that different route deployments have no influence on the results. The achievable lifetime distance is calculated based on the Open-SESAME tool, developed at BFH [43], and the formula (25). Of course, although approximately 500 h and 300 h of operation are considered for 2021 and 2022, respectively, the variance in the data still does not permit clear conclusions yet. For instance, figure 41 shows that the mean energy consumption was varying significantly more in 2022, compared to 2021. The reason for this observation might be the higher variation in ambient temperature, which is also shown in figure 41. Despite these uncertainties in the data, we are confident that the observed lifetime expectancy improvement of an average of almost 12% in 2022, compared to the data of 2021, is not coincidence, but indeed caused by the intelligent heating control. This achieved value even slightly exceeds the value estimated originally by Speckien, as mentioned in section 4.1.1 and [44].

4.2 SOH control

Based on the findings presented in section 3.3 (see figure 36), we know that there is a strong conflict of objectives between the achievable battery lifetime and the energy efficiency. An EMS that best deals with both objectives in order to establish a good trade-off is a desirable development. Such an EMS should balance the longevity of the battery with the energy efficiency such that the EOL of the battery

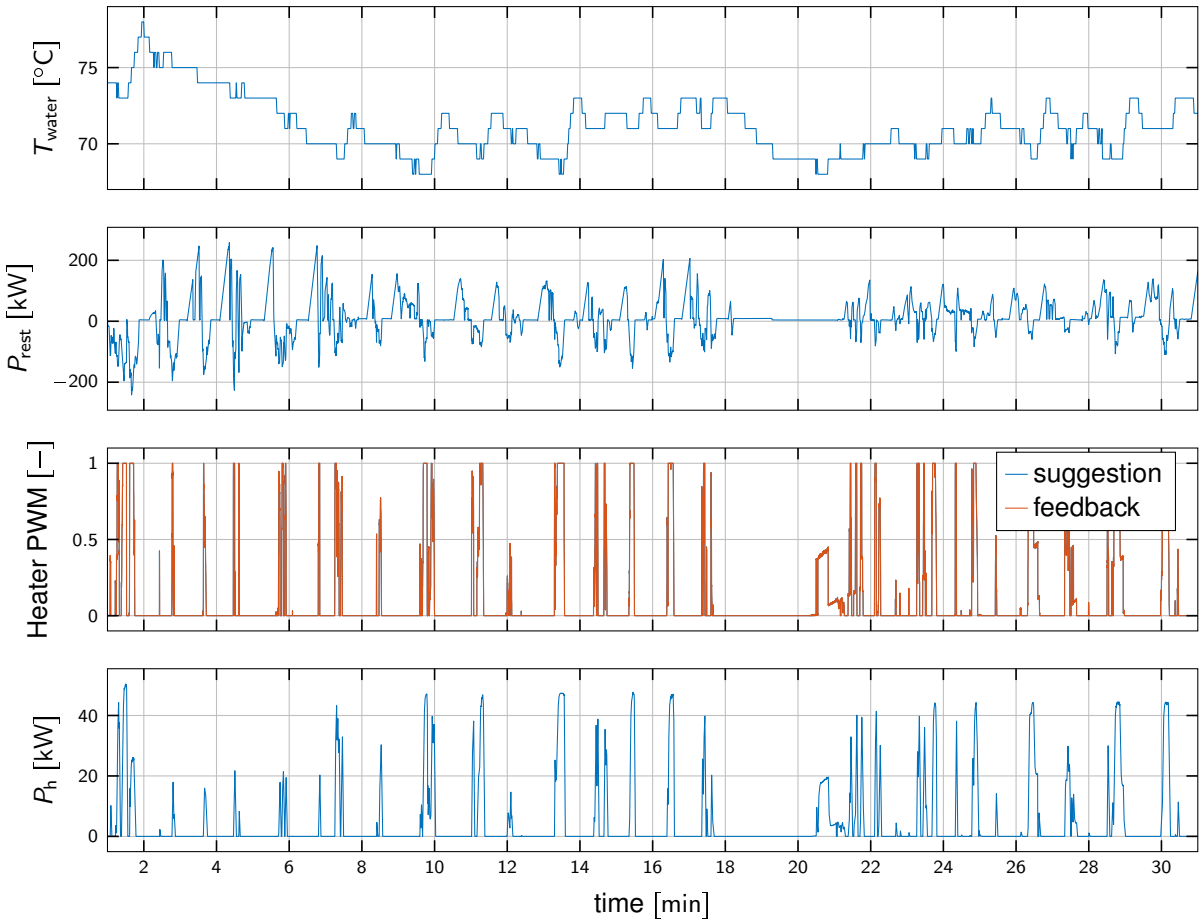

Figure 39: Sample trajectory recorded on the prototype trolley bus on the morning of 2022-10-01. Based on the power request (second graph), the heating PWM fraction is varied (third graph). The “feedback” value from the VCU perfectly follows the suggestion by our control algorithm. The bottom graph shows the measured heating power values.

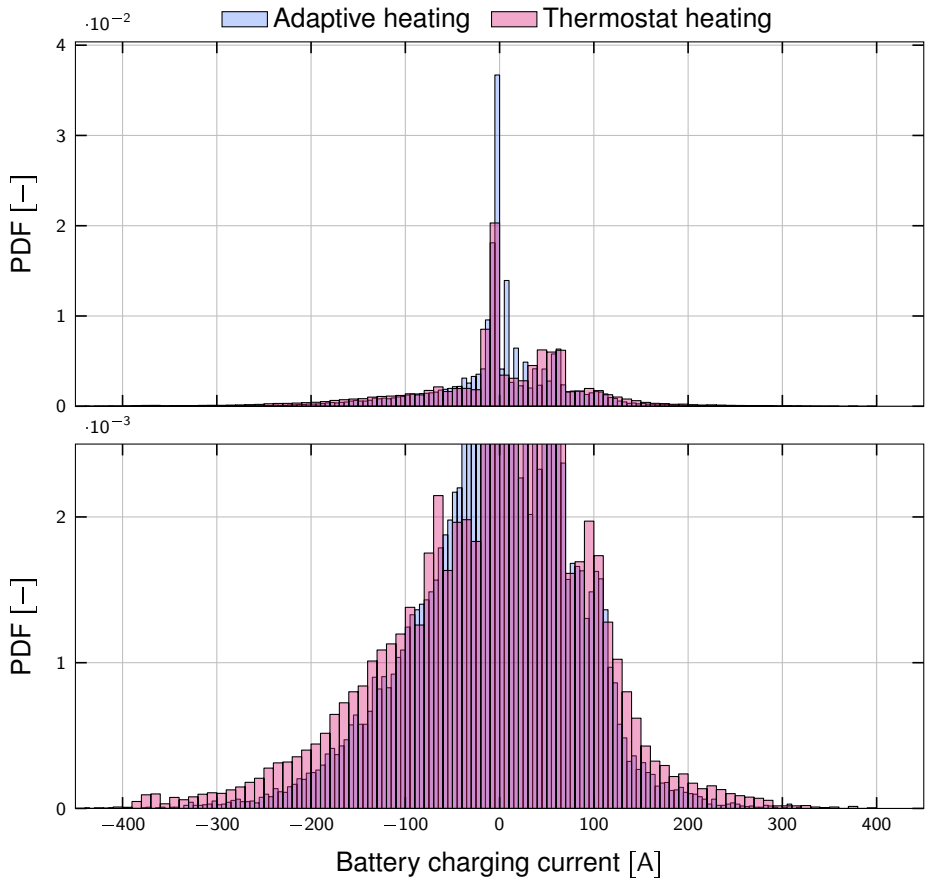

Figure 40: Histogram of the battery charging current. The same data is shown in both graphs, the bottom graph is cropped to allow clearer view of the higher currents. The graphs show data for a case with adaptive heating and one without, both for very similar operating conditions: Both scenarios have been recorded on Route 83, and both have an overall mean power consumption of 20.4 kW. The case with adaptive heating has been recorded during 13 h of operation on 2022-10-08, the case with thermostat heating has been recorded during 14 h of operation on 2021-09-29.

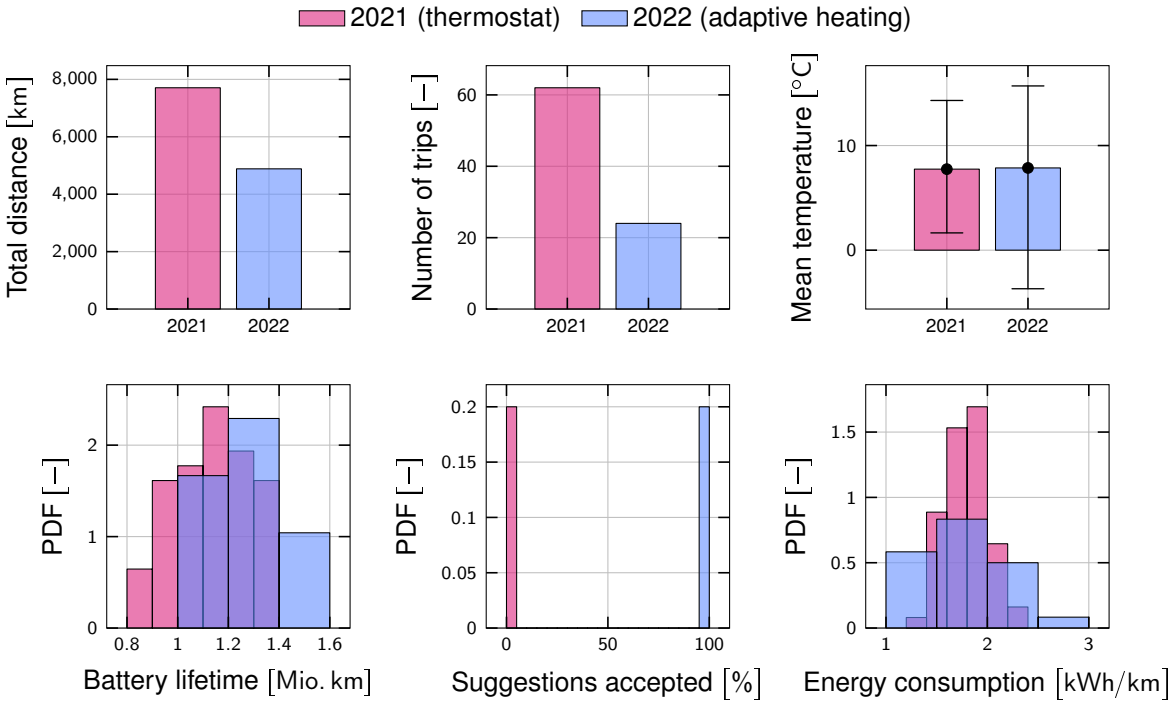

Figure 41: Comparison between trip-averages of fall 2021 (thermostat heating) and fall 2022 (adaptive heating). For this visualization, we consider only trips where the bus was operated on Route 83 and that happen within 01.10. and 13.12. of the respective years.

is reached at the exact time targeted by the battery replacement strategy. In a paper submitted to the journal “eTransportation”, we propose a control framework that is able to utilize SOH measurements to adapt the EMS such that a battery degradation reference can be tracked in a closed loop [47]. A graphical summary of the contributions of this publication is given in figure 42. The three contributions introduced in this figure are further explained in sections 4.2.1 to 4.2.3 below.

4.2.1 Suggested Control Approach

The goal of the controller is to minimize the vehicle energy consumption while ensuring that a certain battery replacement strategy can be followed (e.g., no battery replacements during the vehicle lifetime). To achieve this goal, we propose the three-level control structure as briefly illustrated in the left part of figure 42. Each of the three layers performs a different task and is therefore executed at a different frequency.

In the first level, the power split is determined based on a classical ECMS approach. The second level extends this approach with a second feedback loop, where the battery degradation is estimated based on an internal model and is then compared to a reference trajectory. The novelty of our approach is the third level, where periodic SOH measurements are used to correct for any model mismatch and compensate for non-modeled disturbances. The details of the three levels is explained in the accompanying publication [47].

In a case study, we can show why the feedback measurements suggested are very important in practice, see figure 43. Therefore, we use two scenarios, which are explained in the following. Scenario 1 includes the following disturbances:

- Change in the usage pattern of the bus: Initially, the bus is only operated on the Routes 33, 46, and 72, which feature grid coverage rates between 90% and 100%. After the first $0.3\,\mathrm{Mio.\,km}$ of operation, for the next $0.1\,\mathrm{Mio.\,km}$, the bus is operated only on Route 83 which features a low grid coverage ratio of around 50%, after which it is switched back to the initial route selection.
- The degradation estimation in the health tracking controller based on a simple approximation and it is disturbed by some model mismatch.

Scenario 2 includes the following disturbances:

- The degradation estimation used for health tracking does not properly model calendar aging during standstill of the bus (i.e., overnight). An incorrect estimation of calendar aging during standstill might be a realistic scenario, since the VCU might not correctly record the influence factors such as the cell temperature during standstill. A constant calendar aging of -2.5×10^{-6} is assumed between successive missions.
- At half the lifetime distance, an unexpected health drop of 1% is introduced. This is not detected by the degradation estimation in the health tracking. Such a drop in SOH could be caused by a prolonged standstill of the bus for several weeks, for instance, or by an improper handling of the battery.

The results of a lifetime simulation on the first scenario are visualized in the left graph of figure 43, which shows one simulation where the model mismatch rejection is active (blue) and one where it is deactivated, i.e., no measurements of the SOH are considered (magenta). Due to the systematic overestimation of the degradation, the controller without SOH measurements is consistently over-restricting battery usage. However, the real SOH trajectory slowly deviates due to the model mismatch. In this case, a final SOH of 85.5% is reached, which results in a significant increase in the overall energy consumption by 8.8%. On the other hand, if the measurements are taken into account, a final SOH of 80.3% is achieved despite the disturbances introduced in this scenario.

The results of a lifetime simulation on the second scenario are shown in the right graph of figure 43. The effect of the approximation of overnight calendar aging leads to a systematic underestimation of the overall battery aging. Once a new SOH measurement is available after the unexpected health drop at

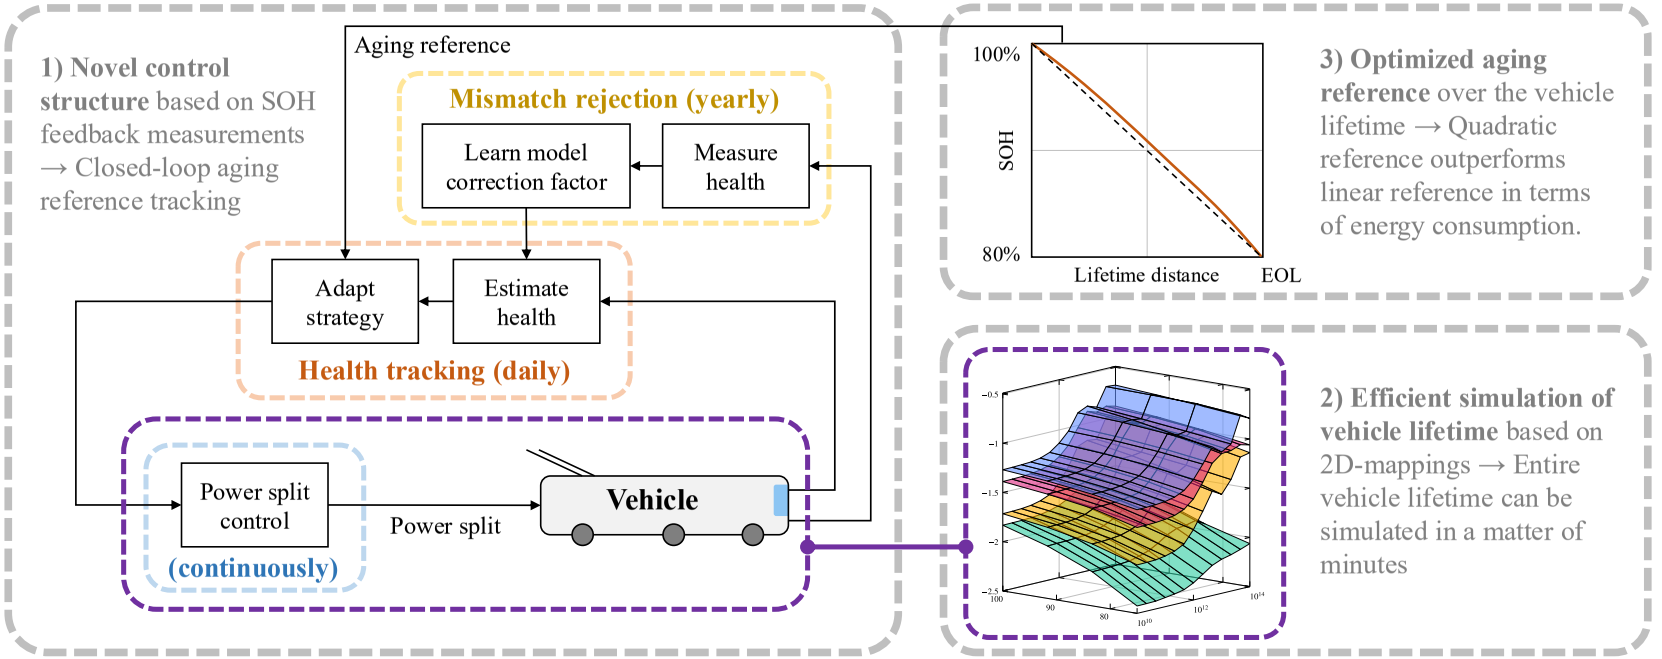

Figure 42: Graphical summary of the contributions of the SOH control publication [47]. A rough overview of the control structure is given in the left part of the figure. The right part of the figure visualizes the other contributions of the publication, i.e., an efficient simulation algorithm and an optimized aging reference trajectory.

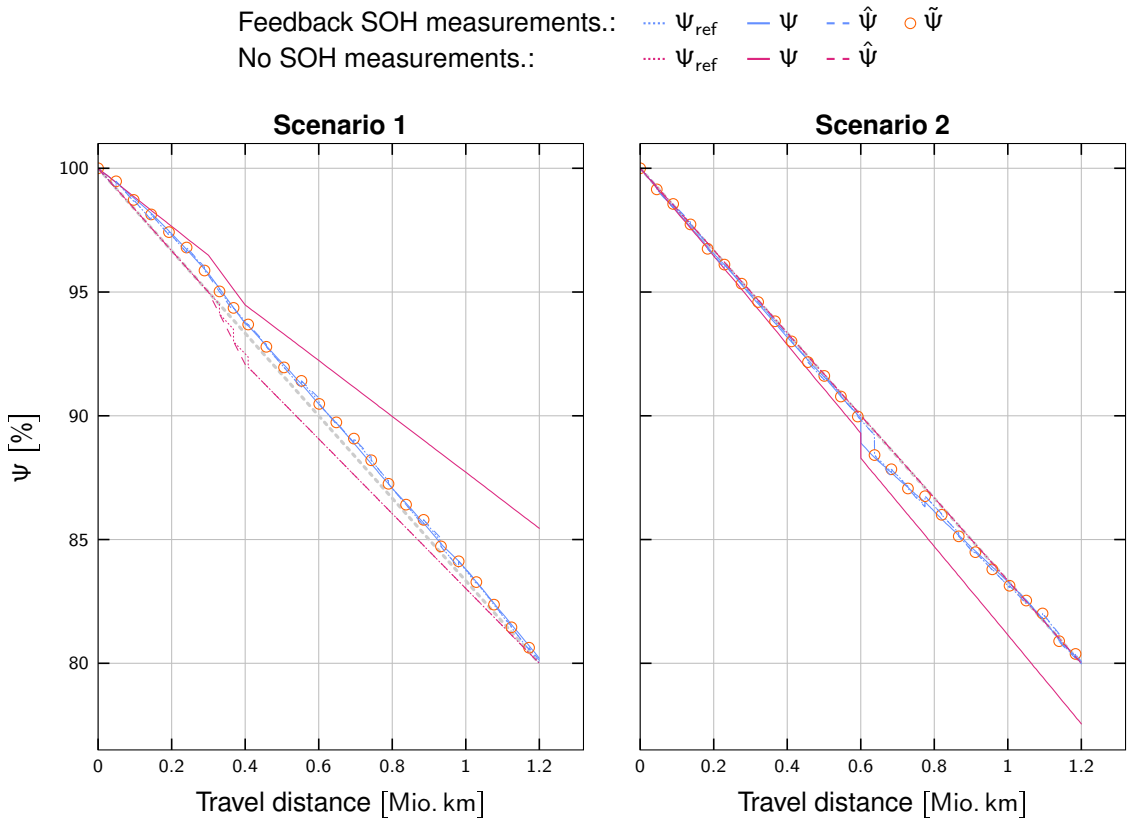

Figure 43: Results of two case studies comparing the suggested control strategy that includes feedback SOH measurements to a classical health-aware EMS that does not include feedback measurements.

0.6 Mio. km, this event is detected by the model mismatch rejection. If no SOH measurement is available, the controller can neither compensate for the model mismatch nor for the health drop, which results in a final SOH of 77.5%. In other words, the EOL of the battery ($\Psi_{\mathrm{EOL}} = 80\%$) is reached already at $d = 1.06$ Mio. km instead of the planned 1.2 Mio. km. Thus, a battery replacement might be necessary. The controller that takes SOH measurements into account, on the other hand, achieves a final SOH of 80.0% with an energy consumption increase of only 3.3% compared to the controller without model mismatch rejection.

Of course, a key "component" of the approach presented here is the measurement of the SOH, i.e., the measurement of the remaining battery capacity. Thanks to the impetus that this concept has been able to provide, several interested parties are now collaborating on the practical implementation of periodic SOH measurements in the field, i.e., VBZ, Basler Verkehrs-Betriebe (BVB), Industrielle Werke Basel (IWB), BFH, Carrosserie HESS AG, and ETH. Furthermore, Carrosserie HESS AG is planning to file a patent for a part of this approach.

4.2.2 Efficient Simulation Algorithm

To be able to evaluate the performance of a controller like the one presented above, it is necessary to simulate the entire vehicle lifetime. Such a simulation is a time-consuming process, because the time scales involved vary over several orders of magnitude: On the one hand, the SOC and the power split are determined every second, whereas the SOH changes only slowly over the period of years. Thus, we show an approach that is able to simulate the entire vehicle lifetime in a matter of minutes, while fully considering the effects of battery degradation on the fast dynamics of the SOC and the power split controller [47].

By comparing the results generated with the suggested approach to the ones generated by simulating the entire vehicle lifetime, we can quantitatively show that the approximations necessary have only minor impacts on the simulation accuracy. The deviations for the degradation and the energy consumption are only about 1.0%. In terms of computational efficiency, we show a speed-up of the simulation by a factor of about 70 000 from 50 days to around 15 min.

4.2.3 Optimization of the Aging Reference

Most published research dealing with battery aging in HEVs or BEVs considers a linear battery degradation reference. However, for two reasons, a nonlinear reference, where the battery aging is slower in the beginning and faster towards the end of the lifetime, is expected to perform better. First, the efficiency of a healthy battery is high, since its internal resistance is low. Thus, the operation with a higher SOH should be preferred. Second, battery degradation increases if the SOH is lower, since the DOD increases. Therefore, the lower the SOH, the more difficult it is to follow a linear aging profile.

We thus use the simulation algorithm introduced in section 4.2.2 above to numerically optimize the degradation reference trajectory [47]. The resulting optimized piecewise linear degradation reference trajectory with 15 sample points between $d = 0$ m and $d = d_{\mathrm{EOL}}$, i.e., 15 degrees of freedom, is shown in figure 44. We can see in this visualization that the optimized reference trajectory can be well approximated with a quadratic function. However, we for our particular application and battery cell type, the assumption of a linear aging profile does not induce any major penalties in achievable energy efficiency. In fact, with an optimized quadratic reference trajectory, the lifetime energy savings achievable are about 0.05%, amounting to approximately 1 MW h over the bus lifetime of 15 to 20 years.

4.3 Outlook

Generally, some of original ideas that have been proposed in the project ISOTHERM description remain open tasks for the subsequent project Swiss eBus Plus. Specifically, the following two points will be tackled in this follow-up project:

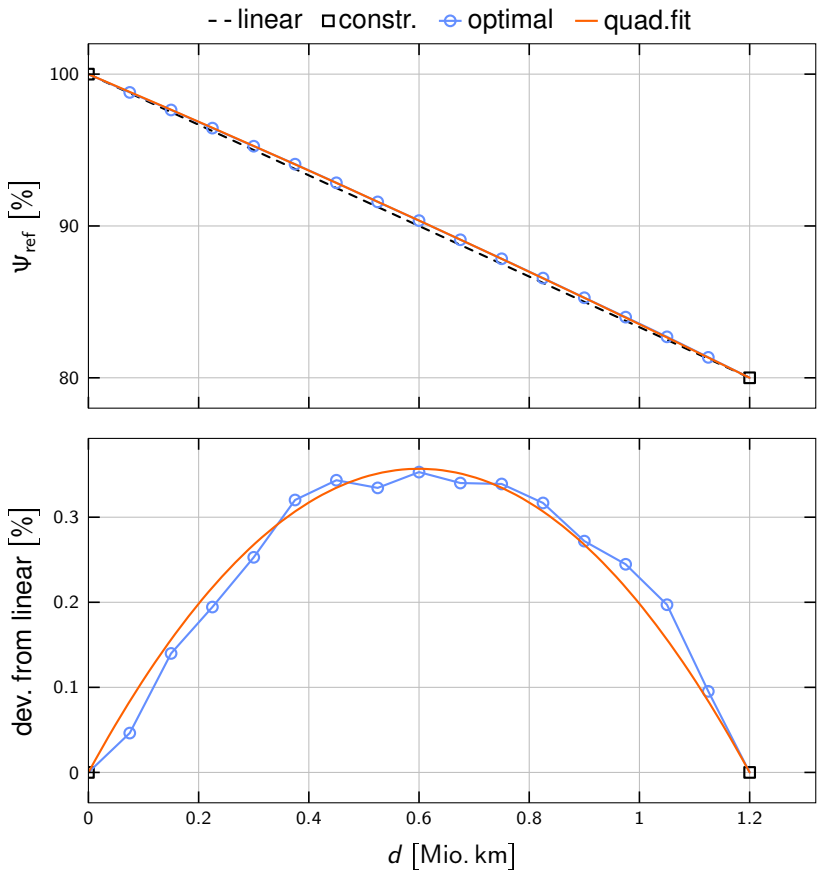

Figure 44: Visualization of the optimized piecewise linear battery degradation reference and a quadratic fit thereof [47]. The lower diagram shows the deviation from the linear reference trajectory for clarity.

- Optimized control of the VCC for summer (i.e., HP) and winter (i.e., AC) operation.
- Development of an integral thermal management that can provide suggestions for the joint operation of the thermal systems and thus address the trade-off between thermal comfort and energy consumption.
- Implementation and testing of (at least part of) the developed algorithms in the field.

The following sections 4.3.1 and 4.3.2 investigate two specific initial ideas for control algorithms in some more detail.

4.3.1 Potential for Optimization of Dual-Source HPs

In this paragraph, we highlight a potential for improvement in the operation of a dual-source heat pump. For this purpose, we compare a simple operating strategy to an improved strategy that is based from results from the scientific literature. Both strategies rely on multiple two-point hysteresis controllers. The baseline strategy is defined as follows:

- A two-point hysteresis controller is used to decide if cabin heating is requested. Its limits are set to $[18\,°\mathrm{C}, 21\,°\mathrm{C}]$. The compressor is turned “on” and the ambient evaporator is used, whenever cabin heating is requested.
- Additionally, a second two-point hysteresis controller is used to decide whether battery cooling is requested. Its limits are set to $[15\,°\mathrm{C}, 20\,°\mathrm{C}]$. The battery cooling request itself does not turn the compressor “on”. However, if the compressor is “on” and battery cooling is requested, the coolant evaporator is used in addition to the ambient evaporator.

For a dual-source HP system similar to the one one the bus under consideration (one air-source evaporator and one coolant evaporator), Ahn et al. show that an “alternating” strategy, where each source is used in isolation, can perform better than a dual-source strategy [48]. The reason for this observation is that the waste heat can be “buffered” in the coolant during air source only operation. After some time, the operation can then be switched to waste heat only, where the high temperature can be used more efficiently. Thus, the improved strategy is defined exactly as the baseline introduced above, however, the use of the ambient evaporator is prevented whenever the coolant evaporator is used.

To compare the two approaches, we conduct a simulation using the modular framework introduced in section 1.4.2. An overview of the model is given in figure 45. The cabin is heated via the dual-source VCC and loses heat to the ambient, which is assumed to be at a constant temperature of $0\,°\mathrm{C}$. For the shell, we assume an overall U-value of $570\,\mathrm{W\,K^{-1}}$, which is consistent with the models introduced above without door losses. The coolant receives a constant flow of waste heat from the traction components of $1\,\mathrm{kW}$. Using these assumptions, we run the simulation for a long period of over $4\,\mathrm{h}$ to estimate its mean performance.

Figure 46 shows the energy balance achieved using both controllers. Indeed, the mean COP can be improved by about 1.2%, i.e., on average, 1.2% more heat is produced per Watt of mechanical input power. Figure 47 shows that this improvement is achieved without negatively affecting the temperature in the passenger compartment or the quality of the battery cooling. Further improvements are expected, for example, if the coolant temperature is allowed to increase further. However, this might come at the expense of a higher battery degradation.

Note that the results presented in this subsection depend on the models of the VCC that have been presented in section 1.1.4.3. As mentioned in section 1.1.4.3, these models have not been properly validated and thus, the results have to be taken with a grain of salt. In particular, Ahn et al. obtained a much higher potential: Using the alternating approach, they showed that 10% more heating capacity and a 5% higher COP can be achieved, compared to the continuous operation [48].

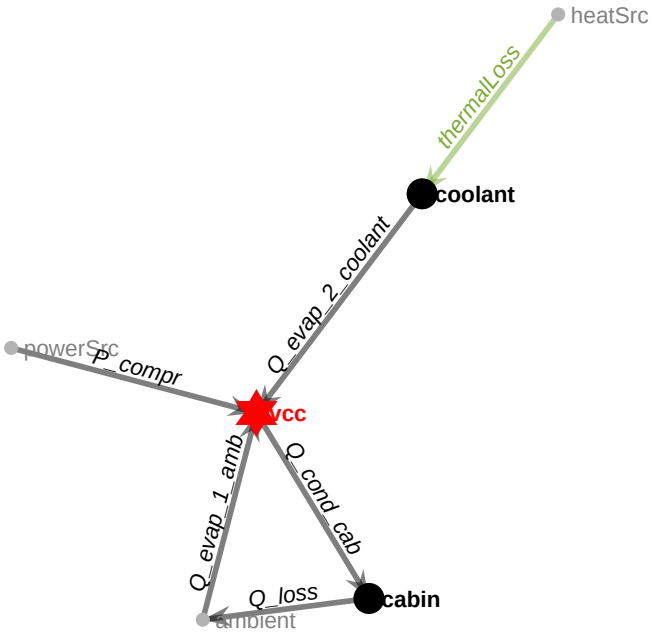

Figure 45: Automatically generated visualization of the dual-source simulation model.

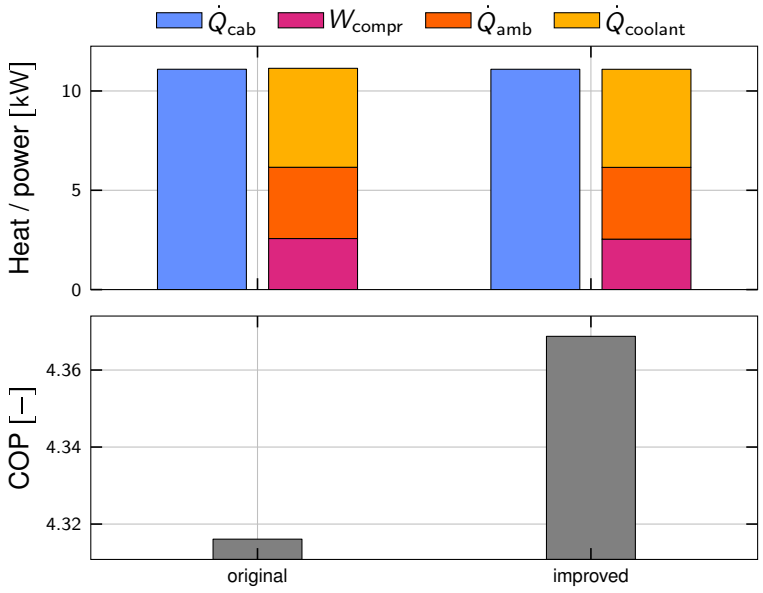

Figure 46: The top graph shows the mean power balance achieved over a long operation using both dual-source controllers. The bottom graph shows the resulting COP, obtained by dividing the useful heat $\dot{Q}_{cab}$ by the compressor mechanical power W_{compr}.

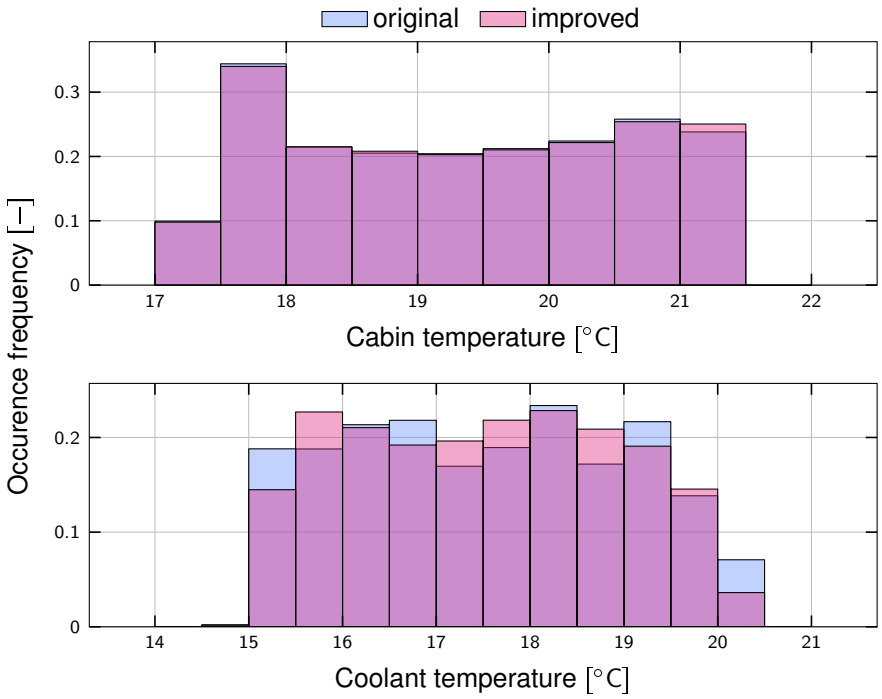

Figure 47: Comparison of the thermal performance of the two controllers.

4.3.2 Potential for Optimization of Compressor Control

Similarly to the control of the water heater, shown in sections 3.3 and 4.1, we expect large potential benefits when combining the control of other consumers linked to thermal reservoirs with the operation of the traction system. The operation of the VCC's compressor is an obvious candidate, as

- the water heater is not needed in summer conditions,
- the HVAC system of future HESS buses is planned without a PTC-based water heater.

Therefore, other consumers need to be considered to harness the potential to reduce battery degradation. Compared to the water heater, the VCC poses the following new challenges:

- VCCs typically only allow binary operation (i.e., "on" or "off").
- The VCC cannot be turned on or off arbitrarily. In particular, the compressor longevity can suffer from frequent switching. Additionally, the VCC's performance is usually poor in the first instances after switching it on. Thus, the VCC should be running for at least a couple of tens of seconds.

The latter constraint, in particular, poses a challenge to control algorithms. A good and robust solution to this task will likely require short-term prediction of the driving profile, or at least involve assumptions about the future driving profile.

5 Conclusion

5.1 Contributions

In the following paragraphs, we attempt to summarize the findings of the last three years that we find most relevant. Where applicable, we link to the corresponding sections of the text where the results are explained in greater detail.

- Temperature stratification is present in reality, but much smaller than in climate chamber measurements.
  A detailed investigation of the temperature stratification evolving in the passenger cabin suggests that the effect of stratification is less pronounced during real-life operation than what has previously been observed in a climate chamber. New measurements and simulations show that the temperature stratification at moderate winter temperatures is in the range of 3 °C to 4 °C, which is in the acceptable range according to, e.g., the ASHRAE Standard 113-1990 [9] or other studies [10]. Section 1.1.3 provides more detail on these investigations.
- The developed modular model framework can offer various simulation and optimization capabilities.
  A framework based on the paradigm of “graph-base modeling” has been proposed. The framework ensures the modularity of the developed models by requiring that the model components only directly influence their “local neighborhood” withing the graph. The developed framework is able to automatically perform both simulations with causal controllers, such as PI controllers, as well as dynamic optimizations. The framework forms a solid basis for further studies to be conducted in the project Swiss eBus Plus. Further information on the modeling framework is given in section 1.4.2, while section 2.2 explains the simulation and optimization procedures.
- Heat pumps offer significant benefits.
  For central European climate conditions, replacing electric PTC heaters with heat pumps yields significant benefits in annual energy consumption: Depending on the vehicle type and the comfort requirements, the mean annual energy consumption can be reduced by 40–60%. This observation is mainly explained by the fact that heating is much more influential than cooling in terms of annual energy consumption. Section 3.2 and the publication [20] provide more insights into these results.
- Radiant heaters not relevant.
  Radiant heaters lead to an increased mean radiant temperature perceived by the passengers. Thus, the same level of thermal comfort can be achieved with a lower air temperature, which reduces thermal losses. For an electric bus heated with electric PTC heaters only, the radiant heaters can reduce the annual energy consumption by 6–7%. However, for other vehicle types like regional or intercity trains, or for vehicles equipped with a heat pump, the energy saving potential of radiant heaters diminishes. Considering the fact that heat pumps are becoming more and more common in the automotive sector, we think that radiant heaters probably have no future in the vehicle categories considered. Section 3.2 and the publication [20] provide more insights into these results.
- Depending on the vehicle type, air curtains or insulation improvements offer larger potential.
  Considering air curtains to reduce losses through open doors is a very promising concept. Assuming a reduction of door losses by 59% [42], the mean annual energy consumption can be decreased by roughly 25% for articulated electric city buses. For trains, the expected improvement is only in the range of 10–15%. On the other hand, a reduction of thermal conductivity through the vehicle shell can reduce annual energy consumption by roughly 20% in trains, while for buses the potential is only at around 10–13%. Section 3.2 provides more insights into these results.
- Thermal buffers usage significantly influences battery degradation.
  An intelligent control of thermal buffers can have a great effect on battery degradation. To harness this potential, only a small thermal energy capacity is needed, but a large converter of at least 50 kW is required (i.e., conversion from electric to thermal energy). Particularly, the trade-off between battery lifetime and energy requirements can be improved. For example, due to an optimal heating

control strategy, a realistic target lifetime can be reached on some routes with a reduction in energy consumption of up to 6.7%, compared to a bang-bang heating strategy. More details on these results are given in section 3.3 and the publication [35]. These results highlight the importance of including battery degradation in the energy and thermal management.

- Intelligent water heater control reduces battery degradation.
  Based on PMP, an intelligent adaptive water heating strategy has been developed for the battery-assisted trolley bus. It has been implemented and tested on a prototype bus of VBZ. Since around June 2022, the controller has been active for over 1000 h without any issues. Extensive data analysis suggests that it is effectively able to reduce battery degradation by about 12%. The controller is explained in more detail in section 4.1.
- Periodic SOH measurements can allow a closed-loop tracking of the SOH.
  To ensure that a battery can be operated for a certain period of time without using an overly conservative strategy or having to replace the battery prematurely, regular SOH measurements are required. Assuming that such measurements are available, for instance, once per year, a three-level control structure has been developed. In simulation studies, this controller is able to perfectly meet the battery lifetime requirements even in the presence of unmodeled disturbances. It is introduced in more detail in section 4.2.1 and the publication [47].
- Linear battery aging reference is not optimal, but “good enough”.
  In extensive simulation studies, we were able to show quantitatively that an ideal battery SOH trajectory declines slower in the beginning of the vehicle lifetime (i.e., slower aging) and shows accelerated aging towards the end. In fact, a quadratic trajectory is able to approximate the optimal trajectory very well. However, for the lithium-titanate-oxide (LTO) cells considered, the lifetime energy savings achievable are only about 0.05%. We thus consider the standard linear aging trajectory to be “good enough”. These results are explained in depth in section 4.2.3 and the publication [47].

5.2 Limitations and Outlook

The following subject areas were planned for in the description of the project ISOTHERM. Due to the initiation of the follow-up project Swiss eBus Plus and the associated premature termination of the project ISOTHERM, they are now covered by the project Swiss eBus Plus. This switch was also necessary as the project Swiss eBus Plus addresses a new vehicle with a different thermal setup.

- Intelligent VCC control:
  Being the central component in the thermal systems in an electric city bus, it is clear that the VCC operation should be analyzed and optimized in detail. Particularly in summer, cooling can only be provided by means of the VCC in any thermal configuration. Furthermore, VCCs are increasingly common also for heating purposes (i.e., heat pumps). However, the operation of VCCs presents new challenges for controllers due to the high complexity of the VCC on the one hand and the restricted on-off operation on the other.

  The optimization of VCC control should be done in simulations first. In later stages of the project, control approaches should also be demonstrated in real-life operation.

  Finally, during the work on this topic within the project ISOTHERM, limitations of the hardware implemented on the bus became apparent: Specifically, on the bus considered in the project ISOTHERM, a direct manipulation of the VCC operation is not possible. This further motivated the setup of a follow-up project on a new vehicle with different thermal components.
- Integral thermal management addressing the comfort-energy trade-off:
  On the basis of the intelligent VCC control, the integral thermal management envisioned in the project ISOTHERM description can be developed. This overarching controller should address the trade-off between thermal comfort and energy consumption in an optimized way, offering simple

intuitive tuning parameters to influence the operation. Again, this development should be performed first in simulation, before testing it in real-life.

In addition to the two tasks mentioned above, we see the following interesting streams of further research and development:

- Continuous adaptation of compressor frequency:
  As mentioned previously, the compressors in VCCs typically only allow “on/off” operation. While a continuously variable compressor frequency is possible in theory, it requires changes in the electrical setup of the vehicle. It would be interesting to determine whether such a development would be desirable from an energetic standpoint.
- Periodic SOH measurements:
  As mentioned in section 4.2, a group of interested parties, formed by VBZ, BVB, IWB, BFH, Carrosserie HESS AG, and ETH, are now collaborating on the practical implementation of periodic SOH measurements in the field. In addition to the practical implementation of the algorithm presented in section 4.2.1, many interesting areas of research might present themselves based on the data generated by such measurements, such as bus scheduling, predictive maintenance, etc.

A List of Publications

The following sections serve as an overview and reference for all published material in relation to the project ISOTHERM. Appendix A.1 contains a list of scientific papers that have been (or will be) peer-reviewed. Appendix A.2 lists the appearances in national conferences. Finally, appendix A.3 provides a complete list of all student projects conducted.

A.1 Scientific Publications

- A. Ritter, F. Widmer, B. Vetterli, and C. H. Onder, “Optimization-based online estimation of vehicle mass and road grade: Theoretical analysis and experimental validation”, in Mechatronics, Volume 80, 2021, `https://doi.org/10.1016/j.mechatronics.2021.102663`.

  Estimation algorithm for road grad and GVM based on measurements of traction torque, covered in section 1.2.2. See also ref. [18].

- F. Widmer, A. Ritter, P. Duhr, and C. H. Onder, “Battery lifetime extension through optimal design and control of traction and heating systems in hybrid drivetrains”, in eTransportation, Volume 14, 2022, `https://doi.org/10.1016/j.etran.2022.100196`.

  Inclusion of a thermal energy buffer into the battery health-aware EMS of a HEV, covered in section 3.3. See also ref. [35].

- F. Widmer, A. Ritter, J. Ritzmann, D. Gerber, and C. H. Onder, “Battery Health Target Tracking for HEVs: Closed-Loop Control Approach, Simulation Framework, and Reference Trajectory Optimization”, submitted to eTransportation, 2022.

  Control approach for closed-loop reference tracking of the SOH in a hybrid drivetrain, covered in section 4.2. See also ref. [47].

- F. Widmer, A. Ritter, M. Achermann, F. Büeler, J. Bagajo, and C. H. Onder, “Highly Efficient Year-Round Energy and Comfort Optimization of HVAC Systems in Electric City Buses”, submitted to the IFAC World Congress, 2023.

  Steady-state simulations of the HVAC system, covered in sections 1.3 and 3.2. See also ref. [20].

A.2 Presence at National Conferences

- SFOE Mobility Research and Innovation in Switzerland Workshop, Lucerne, 2021:
  Presented a poster entitled “Combined Heating and Traction Energy Management to Prolong Battery Life”.
- Center for Sustainable Future Mobility (CSFM) Kick-off Symposium, Zurich, 2022:
  Speed talk in the session “Research on Sustainable Future Mobility” entitled “Managing Battery Lifetime in an Electric City Bus”.

A.3 Student Project Reports

The following abbreviations are used for the different types of theses conducted at ETH. The number of European Credit Transfer and Accumulation System (ECTS) credits gained indicate the amount of work expected to be put into each thesis, where 1 ECTS credit corresponds to about 30 h of work.

SOM Studies on mechatronics – A literature research project performed by some students before the Bachelor's thesis. 5 ECTS credits.

BT Bachelor's thesis – The final project of the Bachelor's degree. 14 ECTS credits.

SP Semester project – A project performed in the course of the Master's degree. 8 ECTS credits.

MT Master's thesis – The final project of the Master's degree. 30 ECTS credits.

A.3.1 2018

SOM: Signal Post Processing from Different Sources – Pol Eyschen [49]

A project providing a comprehensive overview of the various communication systems used on the former SwissTrolley plus prototype. Following this overview, multiple approaches to merge log data from various sources are explored and implemented.

A.3.2 2019

BT: Software in the Loop Tests – Pol Eyschen [50]

Conceptualization and implementation of a software in the loop (SIL) framework for controller testing. The framework is able to execute control algorithms as if they were operated "online" on a bus by emulating the communication with the bus and simulating its response to the control action. This framework allows extensive testing of control algorithms and their implementation before deploying them on the physical prototype. To speed up such tests, they can be executed in an "accelerated" manner.

SP: HVAC System Modeling – Felix Kieser [51]

Re-evaluation of measurement data acquired in climate chamber experiments [6]. Thereby, the focus was put on air stratification and air velocity measurements during door openings. This project represents an initial attempt to bring the air velocity measurements and the performance measurements together via a door loss model based on air exchange.

A.3.3 2020

MT: Automated Generation and Solving of OCPs from Object-Oriented Models – Pietro Pelizzari [31]

Development of an automatic software pipeline that can translate object-oriented model descriptions from Modelica together with a cost function into a nonlinear optimization problem and solve it directly. Development of the first complete vehicle model that represents all relevant thermal reservoirs in the simplest possible way. Test of the developed pipeline against this model and analysis and presentation of these first results. See sections 1.4.1 and 2.1.

SOM: Vapor Compression Cycle Modeling Approaches – Jürg Furrer [11]

Literature survey to identify and categorize the different ways to model the air conditioning system (or heat pump). See section 1.1.4.

BT: Cooling Circuit Modeling – Jürg Furrer [52]

Creation of a model of the traction cooling circuit of the SwissTrolley plus. The individual components are modeled based on the underlying physics and, where possible, the parameters are identified using measured data. In addition to the static model from Carrosserie HESS AG, which were also analyzed, a

dynamic model was created that can calculate the temperature trajectories of the cooling circuit. The comparison with measured data shows that in the future, especially a better loss model of the inverter, as well as of the VCC are necessary.

SP: Scenarios in HVAC Simulations – Ashutosh Pani [38]

Merger of different data sources (hourly bus log data, meteorological data from MeteoSwiss, passenger data from VBZ) into one database. Analysis, completion and simplification (dimension reduction) of this database using machine learning algorithms. Cluster analysis of this data to find similarities and differences. Finally, selection of representative data samples for the entirety of the database. See section 3.1.

MT: Trolley Bus Thermal Management – Jan Speckien[44]

Development of an online control strategy for the water heater. The basic idea is that less energy has to be cycled through the battery if heating is preferred during recuperation (and thus, conversely, heating is no longer required during acceleration). The result is an algorithm based on optimality theory that can be applied online thanks to some simplifications. In simulation presented in this thesis, this algorithm manages to reduce the total energy consumption minimally (0.3%), but above all to reduce the battery damage strongly (10%). See section 4.1.

MT: Vehicle Mass and Road Grade Estimation Using Most Likely Explanation – Basil Vetterli [53]

Further development of an estimation algorithm for simultaneous vehicle mass and road gradient estimation. The algorithm uses signals from the drive system and a model of the vehicle's longitudinal dynamics to set up and solve an optimization problem. In one study each with artificially generated and real measurement data, the new algorithm is shown to perform significantly better than previous algorithms from the literature. This thesis forms the basis of an article that has later been published in the journal "Mechatronics" [18], see also section 1.2.2.

SOM: HVAC Components in Public Transport – Joshua Bagajo [54]

The objective of this literature review was to search technical literature to obtain a list of all measures used to reduce the energy consumption of HVAC systems. The measures are then explained individually and categorized for ease of reference. In addition to the various HVAC components that are typically installed, measures for thermal load reduction and thermal load shifting are discussed in this thesis.

SOM: Thermal Comfort for Passengers – Lukas Huber [2]

In this literature study, the thermal comfort sensation models proposed in the literature are collected and categorized. The models found are categorized into two-block models on the one hand, in which the physiology of the human body is separated from its sensation (psychology). The other large category is formed by the "integral" one-block models, in which no explicit distinction is made between physiology and psychology. See section 1.1.2.

SP: Learning-Based Representative Scenario Selection for HVAC Systems in Electric Buses – Manuel Stähelin [39]

Further development of the automatic simulation scenario selection procedure [38]. The goal is to process the collected, hourly stored log data from the SwissTrolley plus using machine learning methods. The idea is to automatically group "similar" log files (with e.g. similar outside temperature and similar passenger

volume) together. From each such group (“cluster”), it is then presumably sufficient for a simulation study to simulate only one representative scenario instead of the entire group. See section 3.1.

A.3.4 2021

MT: Holistic Energy Optimization for Electric Buses with Focus on the Development of a Control-Oriented Dynamic VCC Model – Manuel Stähelin [13]

Development of a holistic model of the complete bus. Compared to previous work, the focus in this project is put on the development of a dynamic VCC model. For this purpose, initially, a high-fidelity physics based model of the VCC is developed. This serves as a “ground truth” to validate a second, much simpler model. Being based on only two dynamic states, this simpler model can be successfully included in dynamic optimization problems of the entire bus operation. See section 1.1.4.3.

MT: Object-Oriented Modeling Approach applied on a Modern Trolley Bus – Alan Imperatori [32]

In this work, the principle of object-oriented, physical modeling was applied to the bus system for the first time. The well-known modeling language “Modelica” was used for this purpose. Since the work thus breaks new ground for the project, the report deals heavily with the concepts, as well as advantages and disadvantages of object-oriented modeling. Finally, a model for the traction cooling circuit is developed and shown. See section 1.4.1

SP: Optimal Controller Parameters for Robust Online Control of a Trolley Bus – Luca Sandel [55]

In-depth analysis of the control algorithm for regulating the battery charge level, which has been developed originally for the SwissTrolley plus. The analysis shows that the controller should be analyzed in two phases: There is an initial “settling phase” before the controller enters a kind of limit cycle. The transient phase can be shortened by choosing appropriate initialization parameters. The trade-off between robustness and energy consumption can be visualized very well in this work.

BT: System Efficiency of a Combination of Infrared Heater and Heat Pump under Various Conditions on an Electric Bus – Joshua Bagajo [21]

With a simple model, which can only represent steady-state situations, the sensitivities of infrared heating could be investigated in this work. Based on the comfort model according to EN-ISO 7730, it is shown that in the given model with one passenger placed directly underneath an infrared heating surface, an energy saving potential exists by using infrared heaters. This is realized because through their use the air temperature can be slightly lowered to achieve the same comfort index. The work shows that this benefit strongly depends on the achievable COP of the heat pump: At high COP, the advantage of infrared heating disappears. See section 1.3.

SP: Forecasting Passenger Load on Municipal Bus Routes [19]

The task of the work was to develop a “predictor” for the number of passengers on the bus based on machine learning approaches. Special care was taken to ensure that only information available in real time on the bus (on the vehicle control unit) was used for the prediction – i.e., no information from the dispatch center. Data from the passenger counting systems of different buses were used as training data. Surprisingly, we were able to show that weather has only a negligible effect on passenger counts. The information about time and date (incl. weekday, vacations, etc.) as well as the location (GPS), on the other hand, are significantly more important. See section 1.2.3.

SP: Infrared Heating in an Electric Trolley Bus – Fabian Büeler [22]

With a simple model, which can represent steady-state situations, the influence of the infrared heating could be examined in this work. In contrast to earlier analyses [21], situations with several heaters and several passengers were investigated on the one hand, and the analysis was extended to a whole "typical" year on the other. The result of the investigation is that although the use of infrared heaters can reduce the interior temperature and thus heat losses, this gain is more than compensated by the energy demand of the infrared heaters if the bus is heated with a heat pump. If the bus is heated with conventional electric heaters, there is a savings potential of approx. 6%, which is, however, significantly smaller than the savings potential through the use of a heat pump. See section 1.3.

MT: Closed-Loop Battery Lifetime Target Tracking Using a Battery Health-Aware Energy and Thermal Management – David Gerber [56]

Development of an energy management framework that takes battery aging into account when determining the current power distribution. In addition, the structure includes periodic measurements of the battery remaining capacity to determine the SOH. These measurements are used by the control structure to guarantee a certain battery lifetime. This is to prevent the battery from aging too quickly and needing to be replaced, on the one hand, and to prevent the battery from being overused, resulting in unnecessarily high energy consumption, on the other. A mathematical simplification of the problem allows to simulate the whole lifetime of a bus within a few minutes. This thesis forms the basis of an article that has later been submitted to the journal eTransportation [47]. See section 4.2.

A.3.5 2022

SOM: Review on Optimization Building Design for Energy Efficiency – Mathias Achermann [57]

In this literature study, the focus is deliberately placed on the building sector. The motivation behind this is that thermal systems in the building sector have already been the target of scientific investigations much more frequently than is the case with vehicles. In particular, design optimization is analyzed. The report provides a list of the models used of the building and the climatic conditions, as well as the thermal comfort. Also the different mathematical and numerical tools are shown.

BT: Design Sensitivities in Thermal Comfort Systems. Optimizing the Design of a City Bus Based on Year-Round Performance – Mathias Achermann [23]

In this project, a thermal model of the bus including thermal comfort is presented. The model is based on the basic assumption of a steady state. This allows the model to be set up as a system of equations to calculate energy consumption in order to achieve a certain comfort criterion. The model extends models created in previous theses [21, 22] with the ability to represent cooling and ventilation scenarios. Therefore, the comfort criterion is set as an allowed band instead of a specific value. Based on this model, the annual energy consumption can be analyzed to determine how much a particular comfort criterion “costs”. In addition, sensitivity studies are conducted to determine which parameters have a particularly large impact on energy consumption. See section 1.3.

SOM: Dual-Source Heat Pumps in Automotive Applications. Overview of Current Research Topics – Maximilian Krause [58]

Dual-source heat pumps (heat pumps with at least two heat sources) are heat pumps that draw their heat not only from ambient air, but also from a second source (such as the coolant circuit in the case of the SwissTrolley plus bus concept). In this literature study, such dual-source concepts are first explained in

detail. Then, different configurations with different heat sources including the respective advantages and disadvantages are shown. Finally, concrete implementation examples are listed.

BT: High-Fidelity Heat Pump Modeling. Model Adaptation and Parameter Identification Based on Real-Life Data – Maximilian Krause [12]

In this work, a dynamic model of the heat pump installed on HESS buses is developed based on physical system knowledge. The Simulink extension “Simscape” is used for this purpose. Physical measurements (such as the dimensions of the heat exchangers), the data of the steady-state operating points of the manufacturer of the heat pump, and additionally generated time-resolved measurements of the pipe temperatures serve as the basis for the parameterization. With the model, both the steady-state operating points and the transient pipe temperatures during cyclic switching on and off can be simulated relatively well. As a validation, the model is run without further adjustments as an air conditioner instead of a heat pump, where it also reproduces the manufacturer’s steady-state operating points relatively well. This is an impressive demonstration of how good the generalization capability of the models is. See section 1.1.4.2.

BT: Stratified Thermal Model for a Bus Cabin – Flurin Solèr [7]

In this work, a static model of the cabin is developed based on various models from the literature as well as basic physical equations to realistically represent the temperature stratification during operation. The model is validated using data recorded in the SBB climate chamber in 2017 [6]. Furthermore, additional data recorded by means of new measurement equipment in real operation is shown. This data, as well as the results of the model, suggest that the temperature stratification is in the range of 3-4 °C at moderate winter temperatures, which is significantly lower than the values recorded in the climate chamber. In summer, the stratification is smaller. With the help of a literature review in regarding thermal comfort, this project concludes that the observed temperature stratifications probably have only a small effect on thermal comfort and can therefore be neglected for most investigations. A verification of this statement would require more elaborate measurements. See section 1.1.3.

References

[1] Andreas Ritter, Fabio Widmer, Christopher Onder, Alejandro Santis, Timan Schneider, Ludovic Lauber, Priscilla Caliandro, Andrea Vezzini, and Martin Widmer. SwissTrolley plus. Schlussbericht (F&E) SI/501321, Bundesamt für Energie BFE, 2019. URL: https://www.aramis.admin.ch/Grunddaten/?ProjectID=37064.

[2] Lukas Huber. Thermal comfort for passengers. Studies on mechatronics, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, December 2020. IDSC-CO-FW-12.

[3] Qiantao Zhao, Zhiwei Lian, and Dayi Lai. Thermal comfort models and their developments: A review. *Energy and Built Environment*, 2(1):21–33, 2021. doi:10.1016/j.enbenv.2020.05.007.

[4] Povl Ole Fanger. *Thermal Comfort: Analysis and Applications in Environmental Engineering*. Danish Technical Press, Copenhagen, Denmark, 1970.

[5] Ergonomics of the thermal environment – analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort crieteria. Norm, European Committee for Standardization, Brussels, Belgium, 2005.

[6] Franz Sidler, Stefan Ineichen, and Gerhard Zweifel. Messung Energieverlust durch Türöffnungen bei Linien-Bussen. Schlussbericht 098 (des ESöV 2050-Programms), Bundesamt für Verkehr BAV, 2019. URL: https://www.aramis.admin.ch/Grunddaten/?ProjectID=43762.

[7] Flurin Solèr. Stratified thermal model for a bus cabin. Bachelor's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, July 2022. IDSC-CO-FW-20.

[8] Sven Strebel. Untersuchung Heizenergiebedarf eines Cobra-Trams im Winterbetrieb. Master's thesis, HSR – Hochschule für Technik, Rapperswil, Switzerland, June 2017.

[9] Method of testing for room air diffusion. Norm, ASHRAE, 1990.

[10] H. Zhang, C. Huizenga, E. Arens, and T. Yu. Modeling thermal comfort in stratified environments. In *Proceedings, Indoor Air 2005: 10th International Conference on Indoor Air Quality and Climate*, 2005. URL: https://escholarship.org/uc/item/8q58k4hs.

[11] Jürg Furrer. Vapor compression cycle modeling approaches. Studies on mechatronics, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, March 2020. IDSC-CO-FW-07.

[12] Maximilian Krause. High-fidelity heat pump modeling. Bachelor's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, July 2022. IDSC-CO-FW-19.

[13] Manuel Stähelin. Holistic energy optimization for electric buses with focus on the development of a control-oriented dynamic VCC model. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, September 2021. IDSC-CO-FW-14.

[14] Q. Zhang, S. Stockar, and M. Canova. Energy-optimal control of an automotive air conditioning system for ancillary load reduction. *IEEE Transactions on Control Systems Technology*, 24(1):67–80, 2016. doi:10.1109/TCST.2015.2418322.

[15] Ian H. Bell, Jorrit Wronski, Sylvain Quoilin, and Vincent Lemort. Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library CoolProp. *Industrial & Engineering Chemistry Research*, 53(6):2498–2508, 2014. doi:10.1021/ie4033999.

[16] Yulong Song, Ce Cui, Xiang Yin, and Feng Cao. Advanced development and application of transcritical co2 refrigeration and heat pump technology—a review. *Energy Reports*, 8:7840–7869, 2022. doi:10.1016/j.egyr.2022.05.233.

[17] K. N. Patil, S. N. Garg, and S. C. Kaushik. Luminous efficacy model validation and computation of solar illuminance for different climates of india. *Journal of Renewable and Sustainable Energy*, 5(6):063120, 2013. doi:10.1063/1.4841195.

[18] Andreas Ritter, Fabio Widmer, Basil Vetterli, and Christopher H. Onder. Optimization-based online estimation of vehicle mass and road grade: Theoretical analysis and experimental validation. *Mechatronics*, 80:102663, 2021. doi:10.1016/j.mechatronics.2021.102663.

[19] Lau Stokholm. Forecasting passenger load on municipal bus routes. Semester project, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, December 2021. IDSC-CO-FW-18.

[20] Fabio Widmer, Andreas Ritter, Mathias Achermann, Joshua Bagajo, Fabian Büeler, and Christopher H. Onder. Highly efficient year-round energy and comfort optimization of HVAC systems in electric city buses. Submitted to the 2023 IFAC World Congress, 2022.

[21] Joshua Bagajo. System efficiency of a combination of infrared heater and heat pump under various conditions on an electric bus. Bachelor's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, July 2021. IDSC-CO-FW-13.

[22] Fabian Büeler. Infrared heating in an electric trolley bus. Semester project, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, February 2022. IDSC-CO-FW-17.

[23] Mathias Achermann. Design sensitivities in thermal comfort systems. Bachelor's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, July 2022. IDSC-CO-FW-21.

[24] George Havenith, Dusan Fiala, Krzysztof Błazejczyk, Mark Richards, Peter Bröde, Ingvar Holmér, Hannu Rintamaki, Yael Benshabat, and Gerd Jendritzky. The UTCI-clothing model. *International Journal of Biometeorology*, 56:461–470, 2012. doi:10.1007/s00484-011-0451-4.

[25] Building components and building elements – thermal resistance and thermal transmittance – calculation methods. Norm, European Committee for Standardization, Brussels, Belgium, 2017.

[26] Miguel Centeno Brito, Teresa Santos, Filipe Moura, David Pera, and Jorge Rocha. Urban solar potential for vehicle integrated photovoltaics. *Transportation Research Part D: Transport and Environment*, 94:102810, 2021. doi:10.1016/j.trd.2021.102810.

[27] Christoph Isenschmid. Potentialanalyse für Verbesserung der Wärmedämmung bei der Fahrzeug-Aussenhülle. Schlussbericht 010 (des ESöV 2050-Programms), Bundesamt für Verkehr BAV, 2018. URL: www.bav.admin.ch/energie2050.

[28] Franz Sidler. Simulation Energiebilanz Heizung Lüftung Klima: ICN (SBB). Schlussbericht 056 (des ESöV 2050-Programms), Bundesamt für Verkehr BAV, 2018. URL: www.bav.admin.ch/energie2050.

[29] The Modelica Association. https://www.modelica.org/. Accessed: 2022-10-18.

[30] Johan Åkesson. Optimica – an extension of Modelica supporting dynamic optimization. In *6th International Modelica Conference*, pages 57–66, March 2008.

[31] Pietro Pelizzari. Automated generation and solving of OCPs from object-oriented models. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, May 2020. IDSC-CO-FW-05.

[32] Alan Imperatori. Object-oriented modeling approach applied on a modern trolley bus. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, January 2021. IDSC-CO-FW-09.

[33] Justin P. Koeln, Matthew A. Williams, and Andrew G. Alleyne. Hierarchical control of multi-domain power flow in mobile systems: Part I — framework development and demonstration. In *ASME 2015 Dynamic Systems and Control Conference*, October 2015. V001T08A006. doi:10.1115/DSCC2015-9908.

[34] Donald J. Docimo and Andrew G. Alleyne. Electro-thermal graph-based modeling for hierarchical control with application to an electric vehicle. In *2018 IEEE Conference on Control Technology and Applications (CCTA)*, pages 812–819, 2018. doi:10.1109/CCTA.2018.8511390.

[35] Fabio Widmer, Andreas Ritter, Pol Duhr, and Christopher H. Onder. Battery lifetime extension through optimal design and control of traction and heating systems in hybrid drivetrains. *eTransportation*, 14:100196, 2022. doi:10.1016/j.etran.2022.100196.

[36] Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019. doi:10.1007/s12532-018-0139-4.

[37] Andreas Wächter and Lorenz Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106:25–57, 03 2006. doi:10.1007/s10107-004-0559-y.

[38] Ashutosh Pani. Scenarios in HVAC simulations. Semester project, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, May 2020. IDSC-CO-FW-06.

[39] Manuel Stähelin. Learning-based representative scenario selection for HVAC systems in electric buses. Semester project, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, December 2020. IDSC-CO-FW-10.

[40] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction, 2020. arXiv:1802.03426.

[41] Franz Sidler and Frank Gubser. Leistungsmessung an Bussen mit Türluftschleiern. Schlussbericht 159 (des ESöV 2050-Programms), Bundesamt für Verkehr BAV, 2021. URL: www.bav.admin.ch/energie2050.

[42] Aditya Pathak, Matthias Binder, Fengqi Chang, Aybike Ongel, and Markus Lienkamp. Analysis of the influence of air curtain on reducing the heat infiltration and costs in urban electric buses. *International Journal of Automotive Technology*, 21:147–157, 2020. doi:10.1007/s12239-020-0015-x.

[43] Marco Beyeler, Shubham Sharad Bhoir, Stefan Broennimann, and Yoann Moullet. Open-SESAME-Battery. V1.1, https://gitlab.ti.bfh.ch/oss/esrec/open-sesame, 2022.

[44] Jan Speckien. Trolley bus thermal management. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, June 2020. IDSC-CO-FW-04.

[45] Fabio Widmer. Trip-wise adaptive ECMS for HEVs in public transportation. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, May 2017. IDSC-CO-AR-12.

[46] Lino Guzzella and Antonio Sciarretta. *Vehicle Propulsion Systems*. Introduction to Modeling and Optimization. Springer, third edition, 2013.

[47] Fabio Widmer, Andreas Ritter, Johannes Ritzmann, David Gerber, and Christopher H. Onder. Battery health target tracking for HEVs: Closed-loop control approach, simulation framework, and reference trajectory optimization. Submitted to eTransportation, 2022.

[48] Jae Hwan Ahn, Hoon Kang, Ho Seong Lee, Hae Won Jung, Changhyun Baek, and Yongchan Kim. Heating performance characteristics of a dual source heat pump using air and waste heat in electric vehicles. *Applied Energy*, 119:1 – 9, 2014. doi:10.1016/j.apenergy.2013.12.065.

[49] Pol Eyschen. Signal post processing from different sources. Studies on mechatronics, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, December 2018. IDSC-CO-FW-01.

[50] Pol Eyschen. Software in the loop tests. Bachelor's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, May 2019. IDSC-CO-FW-03.

[51] Felix Kieser. HVAC system modeling. Semester project, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, June 2019. IDSC-CO-FW-02.

[52] Jürg Furrer. Cooling circuit modeling. Bachelor's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, June 2020. IDSC-CO-FW-08.

[53] Basil Vetterli. Vehicle mass and road grade estimation using most likely estimation. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, September 2020. IDSC-CO-AR-29.

[54] Joshua Bagajo. HVAC components in public transport. Studies on mechatronics, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, December 2020. IDSC-CO-FW-11.

[55] Luca Sandel. Optimal controller parameters for robust online control of a trolley bus. Semester project, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, July 2021. IDSC-CO-FW-15.

[56] David Gerber. Closed-loop battery lifetime target tracking using a battery health-aware energy and thermal management. Master's thesis, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, February 2022. IDSC-CO-FW-16.

[57] Mathias Achermann. Review on optimization building design for energy efficiency. Studies on mechatronics, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, April 2022. IDSC-CO-FW-22.

[58] Maximilian Krause. Dual-source heat pumps in automotive applications. Studies on mechatronics, ETH Zurich, Institute for Dynamic Systems and Control, Zurich, Switzerland, May 2022. IDSC-CO-FW-23.