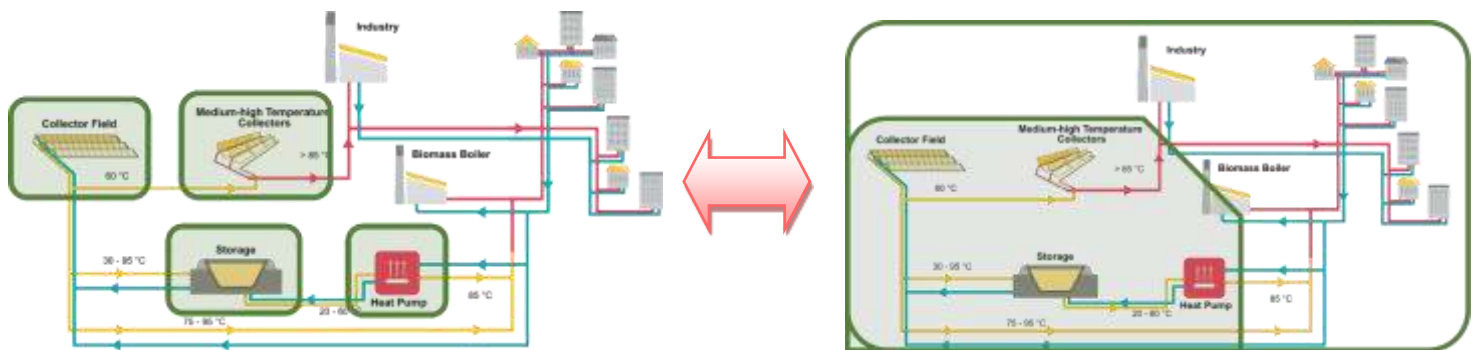


Control of Solar District Heating Systems

Comparison of state-of-the-art and advanced control strategies on the component and system level



Control of Solar District Heating Systems

**This is a report from SHC Task 68:
Efficient Solar District Heating Systems
and work performed in Subtask B:
Data preparation & utilization**

Lead Author: Klaus Lichtenegger (BEST – Bioenergy and Sustainable Technologies)

Supporting Authors and Reviewers: Markus Gölles, Valentin Kaisermayer, Daniel Muschick, Thomas Reiter-Nigitz, Julia Schönfelder, Sandra Staudt (all BEST), Lukas Emberger (SOLID S.E.S.), Viktor Unterberger (formerly BEST), Magdalena Berberich, Thilo Walser (both Solites)

In some sections, this report contains updated material from IEA SHC Task 55, in particular from Factsheets 55.A-D4.1 (**Gölles, Unterberger, Kaisermayer, Nigitz, & Muschick, 2021**) and 55.A-D4.2 (**Gölles, et al., 2021**)

Date: 13.04.2026

Report SHC.Task68.RB3, DOI: [10.18777/ieashc-task68-2026-0001](https://doi.org/10.18777/ieashc-task68-2026-0001)

The contents of this report do not necessarily reflect the viewpoints or policies of the International Energy Agency (IEA) or its member countries, the IEA Solar Heating and Cooling Technology Collaboration Programme (SHC TCP) members or the participating researchers.

Solar Heating and Cooling Technology Collaboration Programme (IEA SHC)

The Solar Heating and Cooling Technology Collaboration Programme began its work in 1977 as one of the first multilateral technology initiatives (“Implementing Agreements”) of the International Energy Agency.

Our mission is *To bring the latest solar heating and cooling research and information to the forefront of the global energy transition.*

IEA SHC members carry out cooperative research, development, demonstrations, and exchanges of information through Tasks (projects) on solar heating and cooling components and systems and their application to advance the deployment and research and development activities in the field of solar heating and cooling.

Our focus areas, with the associated Tasks in parenthesis, include:

- Solar Space Heating and Water Heating (Tasks 14, 19, 26, 44, 54, 69)
- Solar Cooling (Tasks 25, 38, 48, 53, 65)
- Solar Heat for Industrial and Agricultural Processes (Tasks 29, 33, 49, 62, 64, 72)
- Solar District Heating (Tasks 7, 45, 55, 68)
- Solar Buildings/Architecture/Urban Planning (Tasks 8, 11, 12, 13, 20, 22, 23, 28, 37, 40, 41, 47, 51, 52, 56, 59, 63, 66)
- Solar Thermal & PV (Tasks 16, 35, 60, 73)
- Daylighting/Lighting (Tasks 21, 31, 50, 61, 70)
- Materials/Components for Solar Heating and Cooling (Tasks 2, 3, 6, 10, 18, 27, 39)
- Standards, Certification, Test Methods and LCA/LCoH (Tasks 14, 24, 34, 43, 57, 71)
- Resource Assessment (Tasks 1, 4, 5, 9, 17, 36, 46)
- Storage of Solar Heat (Tasks 7, 32, 42, 58, 67)

In addition to our Task work, other activities of the IEA SHC include our:

- SHC Solar Academy
- *Solar Heat Worldwide*, annual statistics report
- SHC International Conference

Our members

Australia	France	SACREEE (UNIDO GN-SEC)
Austria	Germany	SICREEE (UNIDO GN-SEC)
Belgium	ICIMOD/REEECH (UNIDO GN-SEC)	Solar Heat Europe
Canada	International Solar Energy Society	Slovakia
CCREEE (UNIDO GN-SEC)	Italy	South Africa
China	Netherlands	Spain
Denmark	Norway	Switzerland
EACREEE (UNIDO GN-SEC)	Poland	Türkiye
ECREEE (UNIDO GN-SEC)	Portugal	United Kingdom
European Commission	RCREEE (UNIDO GN-SEC)	

For more information on the IEA SHC work, including many free publications, please visit www.iea-shc.org.

Contents

- 1 Executive Summary 2**
- 2 Introduction 3**
- 3 Component-Level Control 6**
 - 3.1 Control of Flat Plate Solar Collectors 7
 - 3.1.1 Control of the feed temperature in the primary circuit 7
 - 3.1.2 Static feed forward control with integrating controller 7
 - 3.1.3 Dynamic feed forward control..... 8
 - 3.2 Control of Concentrating Collectors: Tracking Methods 10
 - 3.3 Further Aspects of Collector Array Control..... 10
 - 3.4 Notes on Further Relevant Technologies 11
 - 3.4.1 Hybrid Collectors (PVT) 11
 - 3.4.2 Absorption Heat Pumps 11
 - 3.4.3 Further Heat Sources..... 11
 - 3.4.4 Thermal Storage Tanks..... 12
 - 3.5 Control of the Heat Exchanger and the Secondary Circuit..... 12
 - 3.6 Interaction with the Control of District Heating Networks..... 13
- 4 System-Level Control..... 14**
 - 4.1 Basic Aspects and Considerations 14
 - 4.1.1 Main Objectives of Superordinate Control..... 14
 - 4.1.2 Planning and Superordinate Control 15
 - 4.1.3 Forecasting 15
 - 4.1.4 Resilience, Safety and Security 15
 - 4.2 Rule-based Supervisory Control 16
 - 4.2.1 Classical State-based Rules 16
 - 4.2.2 State-based Fuzzy Rules 17
 - 4.2.3 Prediction-Based Rules for Supervisory Control 17
 - 4.3 Optimization-based Control Strategies..... 21
 - 4.3.1 Centralized Optimization-based Predictive Supervisory Control 21
 - 4.3.2 Decentralized and Agent-based Methods of Supervisory Control 23
 - 4.4 Machine-Learning and AI-based Control Strategies..... 24
 - 4.4.1 Supervised Learning 24
 - 4.4.2 Reinforcement Learning 24
 - 4.5 Supervisory Control: Conclusions 25
- 5 Conclusions and Outlook 26**
- 6 References 27**

1 Executive Summary

The integration of large-scale solar thermal plants into district heating systems has the potential to play a key role in the decarbonization of the heat supply. Solar energy is abundant, largely predictable at certain timescales, and well-suited to large-scale thermal applications. Thus, solar district heating solutions (SDH) can reliably provide significant solar fractions.

The fluctuating nature of solar radiation and the interaction of multiple components, however, require advanced control strategies for efficient and reliable operation. Control strategies can be broken down into two distinct but equally important levels:

- Component-level control, which must react within seconds and governs pumps, valves, collector arrays, heat exchangers and storage charging/discharging.
- System-level supervisory control, which optimizes overall operation using forecasts of heat demand, solar yield and electricity prices, typically updating schedules on 15-minute cycles.

This report covers both levels of control:

- Section 3 covers component control for flat-plate and concentrating collectors, hydraulic networks, thermal storage, heat exchangers and district heating interfaces.
- Section 4 compares supervisory control strategies, including rule-based approaches, predictive scheduling, optimization-based energy management systems (EMS), and emerging concepts from machine learning and artificial intelligence (AI), such as surrogate models and reinforcement learning.

Across all approaches, the objective remains consistent: maximize solar utilization while ensuring reliability, resilience and economic performance. Predictive strategies that combine physical modelling, operational forecasts and optimization are now considered the most promising single approach for SDH operation. AI-based methods, though rapidly evolving, still require careful handling due to transparency, robustness and data-availability challenges.

The report concludes with a synthesis of the strengths and limitations of the main control strategies and identifies key areas for future research, including hybrid physics-AI approaches and improved explainability of advanced supervisory control systems.

2 Introduction

Solar Energy: Among all renewable energy sources, solar energy certainly has the greatest potential. The global human energy consumption per year is approximately equal to the amount of energy received by the Earth from the sun within one hour. This energy ultimately powers wind and hydro energy as well as biomass, but it can also be used directly, either to generate electric energy (photovoltaics, PV) or to generate heat via solar thermal technology. While PV provides the more flexible form of energy, solar thermal solutions are often the better option for large-scale heat applications, like industrial process heat or heat supply for communities (Lichtenegger, et al., 2025).

Due to their pronounced time characteristics (day-night cycle and seasonal effects) and due to the fluctuating nature of solar energy, solar solutions are usually embedded in some complex energy system, which may consist of distribution channels, short- and long-term storage, complementary energy sources and the possibility of load shifts. This also includes district heating grids, which have the capability to integrate vastly different heat sources. With the need to phase out fossil energy, renewable solutions will have to replace all fossil heat sources in existing grids within a few decades. In addition, expanding heating grids and building new ones can be a viable alternative to other ways of replacing coal, oil and gas for domestic heating.

Control Challenges: While typical heating grids are certainly complex energy systems, this already applies to single solar thermal plants, which already consist of many elements that interact in a complex way. Thus, the control of large-scale solar thermal systems and heating grids (or even more general hybrid energy systems) in which they are embedded, faces several challenges for control. While the control aspect is often a bit neglected, a sound control strategy is of critical importance for the performance of complex energy systems.

The resulting control tasks are carried out in different layers: At a higher level, supervisory controllers, often loosely called energy management systems, decide on the operating mode of the different plants and components, and provide the reference signals for their controllers. These modes of operation or set points of the different plants and components are used by the local controllers at plant and component level, and by those responsible for the operation of the district heating grid. Thus, the control of large-scale solar thermal systems can be divided into two main categories:

1. Low-level control of single components, like valves and pumps (which may affect solar collectors, heat distribution and the use of storage tanks), but also more complex hydraulic networks. These controllers typically must be able to react very fast (within seconds) to changes in the system.
2. High-level supervisory control of the whole system (energy management system, EMS), which usually make use of predictions for energy yield and consumption and optimize the use of energy. A typical time step for updating the schedule is 15 minutes, though slightly shorter update intervals (e.g. 5 or 10 minutes) or longer ones (e.g. 30 or 60 minutes) are possible as well.

Structure of the Report: In this report, in Sec. 3, we will first present key aspects of low-level control, in particular for components specific to solar thermal plants or of vital importance for them, including the hydraulic control of complex systems. In the final technical section, in Sec. 4, we give an exposition of various approaches to the high-level task of creating schedules for complex energy systems. This is followed by conclusions and outlook.

A list of symbols used in this report is given in Table 1 on p. 4, a list of abbreviations in Table 2 on p. 5.

The present report is partially based on Deliverables of the previous IEA SHC Task on solar district heating, Task 55. It combines and updates material from Fact Sheets 55.B-D3.1 (Control of large-scale solar thermal plants) and 55.A-D4.1 (Supervisory control of large-scale solar thermal systems), with some additional material from 55.A-D4.2 (Control of DHC networks and Reduction of the operating temperatures in DH systems).

Table 1 – List of Symbols

Symbol	Meaning	SI unit
A_{coll}	Gross collector area	m^2
c_1	Heat loss coefficient	$W/(m^2 K)$
c_2	Heat loss coefficient	$W/(m^2 K^2)$
c_p	Average isobaric heat capacity	$J/(kg K)$
$c_{p,P}$	Heat capacity at the primary side of the heat exchanger	$J/(kg K)$
$c_{p,S}$	Heat capacity at the secondary side of the heat exchanger	$J/(kg K)$
ΔT	Temperature difference	K
ΔT_c	Temperature difference on the cold side	K
ΔT_h	Temperature difference on the hot side	K
η_0	Optical efficiency	-
H_{stor}	Heat from storage	J
I_g	Global solar radiation	W/m^2
$K(\theta)$	Incident angle modifier (IAM)	-
\dot{m}_p^*	Desired mass flow at the primary side of the heat exchanger	kg/s
\dot{m}_p	Mass flow at the primary side of the heat exchanger	kg/s
\dot{m}_S^*	Desired mass flow at the secondary side of the heat exchanger	kg/s
\dot{m}_S	Mass flow at the secondary side of the heat exchanger	kg/s
\dot{Q}	Heat flow	W
$\hat{Q}_{Cons}[k n]$	Effective heat consumption, based on predictions	W
$\hat{Q}_{Prod}[k n]$	Predicted heat production	W
θ	Angle of incidence of the global solar radiation	rad
t	Time	s
T_{amb}	Ambient temperature	K
T_{feed}^*	Desired feed temperature	K
\bar{T}_f	Arithmetic mean fluid temperature between inlet and outlet of the collector	K
$T_{p,in}$	Input temperature at the primary side of the heat exchanger	K
$T_{p,out}$	Output temperature at the primary side of the heat exchanger	K
T_{ret}	Return temperature	K
$T_{S,in}$	Input temperature at the secondary side of the heat exchanger	K
$T_{S,out}$	Output temperature at the secondary side of the heat exchanger	K
$T_{toBuffer}$	Temperature of fluid fed into the storage buffer	K

Table 2 – List of Abbreviations

Abbreviation	Meaning
ADMM	Alternating direction method of multipliers
AI	Artificial Intelligence
AR	Auto-regressive
ARIMA	Auto-regressive integrated moving average
ARIMAX	Auto-regressive integrated moving average with eXternal factors
ARMA	Auto-Regressive Moving average
CAPEX	Capital expenditure (investment costs)
CHP	Combined heat and power
CNN	Convolutional neural network
DH	District heating
DHC	District heating and cooling
DHG	District heating grid
EMS	Energy management system
HSt	Heat storage
IAM	Incident angle modifier
LLM	Large language model
LSTM	Long short-term memory
MA	Moving average
MILP	Mixed-integer linear programming
MINLP	Mixed-integer nonlinear programming
MIQCP	Mixed-integer quadratically constrained programming
MIQP	Mixed-integer quadratic programming
ML	Machine learning
MLP	Multi-layered perceptron
MPC	Model-predictive control
OPEX	Operational expenses (operation costs)
PGM	Probabilistic graphical model
PID	Proportional-integral-derivative
PLC	Programmable logic controller
PPO	Proximal Policy Optimization
PV	Photovoltaics
PVT	Photovoltaic-thermal collectors
PWA	Piece-wise-affine
RL	Reinforcement learning
RNN	Recurrent neural network
SCADA	Supervisory control and data acquisition
SDH	Solar district heating
ST	Solar thermal
SSSM	Sequential state space model
SVM	Support vector machines
XAI	eXplainable Artificial Intelligence

3 Component-Level Control

In the present section, we present key aspects of low-level control, in particular for components specific to solar thermal plants or of vital importance for them. The components we discuss in detail are highlighted in Figure 1.

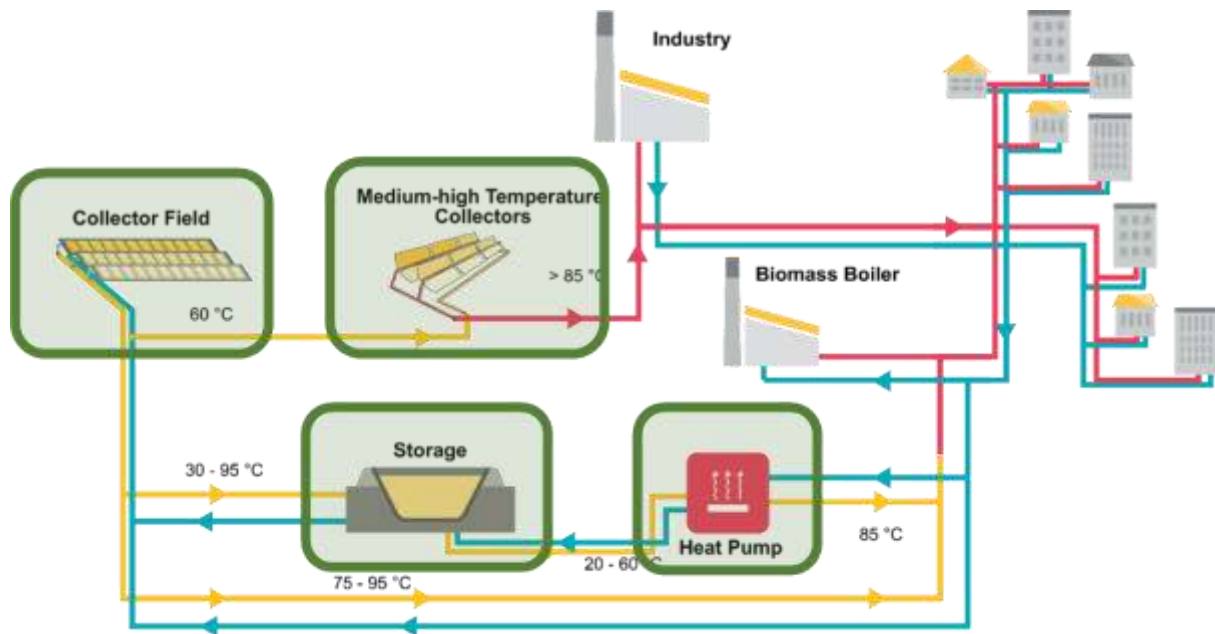


Figure 1: Solar district heating system with single components highlighted

Among all solar thermal plants, flat-plate collectors are by far the most used technology. In Europe about 72% of all solar thermal systems installed contain flat-plate collectors (Weiss & Spörk-Dür, 2024). This ratio is even higher for large-scale solar thermal plants at least partly feeding into DH networks. For this reason, the focus of this section will be on the control of large-scale solar thermal plants consisting of flat-plate collectors connected in series to so-called collector arrays, and with several collector arrays connected in parallel.

Since, other technologies such as parabolic troughs or Fresnel collectors can provide higher temperatures (Tamm & Berberich, 2024), their importance in particular for large district heating grids (which typically operate at forward-feed temperatures above 100 °C) is increasing and a brief description of possible additional or different control tasks for these plants is provided as well.

Independent of the exact type of integration, the availability of on-site consumers or buffer storages, etc., the actual solar thermal plant usually consists of a primary circuit (collector circuit) and a secondary circuit, connected by a plate heat exchanger. As illustrated in Figure 2, the control of these solar thermal plants has the following two tasks:

1. Control of the feed temperature (collector outlet temperature) in the primary circuit.
2. Control of the mass flow in the secondary circuit to optimize the heat transfer.

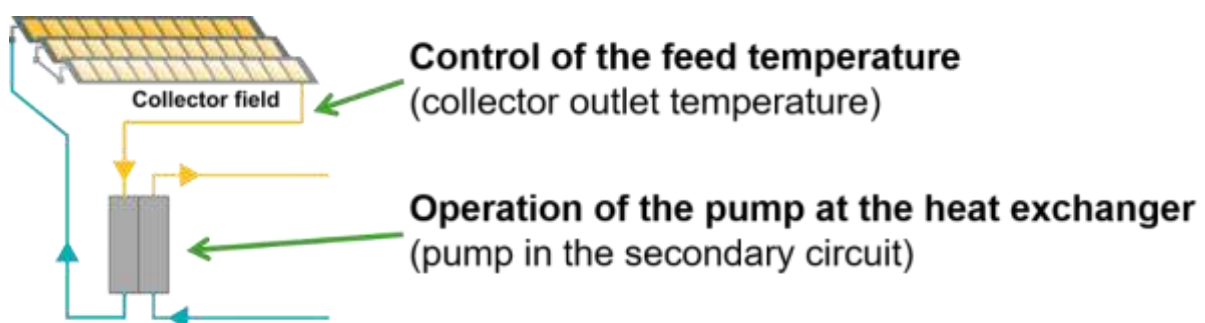


Figure 2: Basic overview of control tasks at large-scale solar thermal plants

3.1 Control of Flat Plate Solar Collectors

3.1.1 Control of the feed temperature in the primary circuit

The feed temperature (collector outlet temperature) in the primary circuit is controlled by adjusting the rotation speed of the pump. In rare cases, an additional 3-way valve can be used for very low mass flows, however, since this is rare and it does not make a methodological difference, it is not considered in the following.

In a simple approach, one uses only PID controllers. While PID controllers are a viable option for many control tasks, they also have severe limitations. For solar collector arrays, PID controllers tend to exhibit rather poor control performance since the control parameters have to be chosen such that reactions are relatively slow. Otherwise, one would obtain rather unstable behaviour because of the long and variable delay times from inlet to outlet – in particular in the case of large solar fields, low flow speed, and rapidly fluctuating irradiation (e.g. clouds on sunny day). Thus, more advanced strategies are required to obtain good control behavior.

3.1.2 Static feed forward control with integrating controller

Among most technology providers and operators, a common strategy is the use of a static feed forward control in combination with an integrating controller, see e.g. (Bava & Furbo, Impact of different improvement measures on the thermal performance of a solar collector field for district heating, 2018; Lemos, Neves-Silva, & Igreja, 2014; Bava, Modeling of solar collector fields for solar heating plants in district heating systems, 2017). The static feed forward control calculates the mass flow \dot{m}_p^* needed to achieve the desired feed temperature T_{feed}^* in steady state for the return temperature T_{ret} , the global radiation I_g and ambient temperature T_{amb} .

In the quadratic efficiency model, which is used in ISO 13061 (International Organization for Standardization, 2022), the corresponding heat flow \dot{Q} is given by

$$\dot{Q}(t) = A_{\text{coll}}K(\theta)\eta_0I_g(t) - A_{\text{coll}}c_1\Delta T(t) - A_{\text{coll}}c_2\Delta T(t)^2 \quad (1)$$

with the temperature difference

$$\Delta T(t) = \bar{T}_{\text{fl}}(t) - T_{\text{amb}}(t), \quad (2)$$

where A_{coll} denotes the gross collector area, I_g the global radiation received by the collector surface, \bar{T}_{fl} the arithmetic mean fluid temperature between the inlet and the outlet of the collector, i.e. the return temperature T_{ret} and the desired feed temperature T_{feed}^* , and T_{amb} denotes the ambient temperature. The coefficients represent the optical efficiency η_0 and the heat loss coefficients c_1 and c_2 . The function $K(\theta)$ represents the incident angle modifier (IAM) which describes the dependency of the optical efficiency η_0 on the angle of incidence θ of the global solar radiation I_g , which varies from collector to collector and is typically estimated through experiments, (Duffie & Beckman, 2013), and given in the data sheet.

By using the average isobaric heat capacity c_p , the mass flow \dot{m}_p^* needed to achieve the desired feed temperature T_{feed}^* in steady state can thus be calculated as

$$\dot{m}_p^* = \frac{A_{\text{coll}} \left[K(\theta)\eta_0I_g - c_1 \left(\frac{T_{\text{feed}}^* + T_{\text{ret}}}{2} \right) - T_{\text{amb}} \right] - c_2 \left(\frac{T_{\text{feed}}^* + T_{\text{ret}}}{2} - T_{\text{amb}} \right)^2}{c_p (T_{\text{feed}}^* - T_{\text{ret}})} \quad (3)$$

Some plants do not directly use the current measurement of the global solar radiation I_g , but filter the signal with a low pass filter in order to remove high-frequency components (noise).

To achieve stationary accuracy, the feed forward control is additionally combined with an integrating controller, however, in many cases the output signal of the controller is not added to the control signal calculated by the integrating controller but multiplied with this signal. For long return and supply pipes, special care must be taken about the positions of the temperature sensors, and possibly the effect of heat losses in the pipes. The desired mass flow \dot{m}_p^* , which must not be below a certain lower limit, is finally maintained by a lower-level flow controller, which must be parametrized so that it reacts reasonably faster than the temperature controller.

A schematic overview of this control strategy, with the controller's output signal being added to the control signal calculated with the feed forward control, is shown in Figure 3.

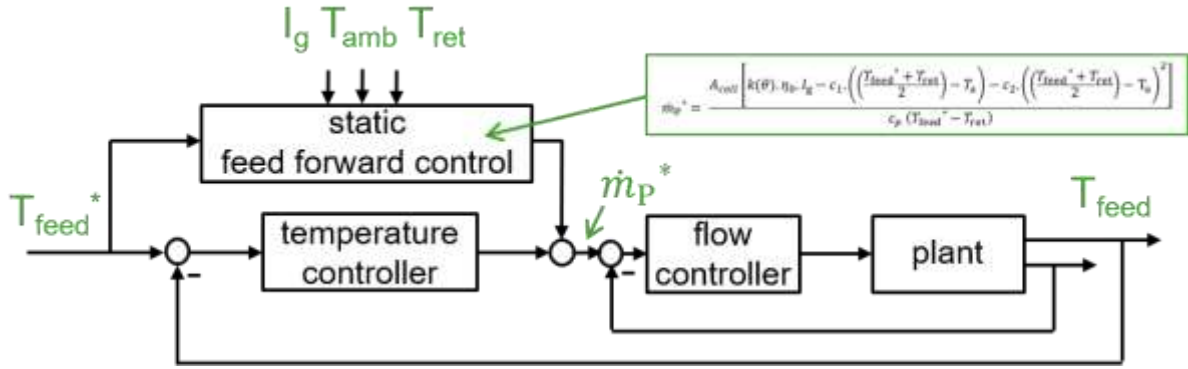


Figure 3: Static feed forward control with integrating controller

3.1.3 Dynamic feed forward control

A control based on static relations, as described in Sec. 3.1.2 has obvious limitations. Thus, the dynamic behaviour of the system should be taken into account as well. While control strategies based on the simulation of heat transport with partial differential equations (PDEs) are possible in principle, this is a rather cumbersome approach.

A more lightweight way of representing the dynamics with a small number of ordinary differential equations (ODEs) has been proposed by K. Lichtenegger in 2015 and been further extended and validated by V. Unterberger, who has given a comprehensive presentation of this approach, followed by some simulation results, in Sec. 3.3.2.5 of (Unterberger, Modelling and control of large-scale solar thermal systems, 2021). Here, we only give a brief presentation of the main ideas.

Starting point is the observation, that from the heat flow equation (1), the time-dependent temperature of a moving fluid element can be described by ordinary differential equation

$$\frac{dT}{dt} = b_0 + b_1 T + b_2 T^2 \quad (4)$$

with coefficients b_i that contain the ambient temperature, collector coefficients and properties of the fluid. This ODE is separable and thus, in principle, can be solved by two integrations:

$$\int_{T_0}^{T(t)} \frac{dT}{b_0 + b_1 T + b_2 T^2} = \int_0^t d\tau \quad (5)$$

While the integral on the right-hand side is trivial, the one on the left-hand side is trickier. It can be analytically evaluated by fractional decomposition, which yields to terms that integrate to logarithms. Recombining these logarithms and some algebraic transformations yields an analytic expression for $T(t)$, which for $t = t_{tot}$, the total throughput time, yields the outlet temperature

$$T_{out}(t) = T_- + \frac{(T_+ - T_-)(T_0 - T_-)}{T_0 - T_- + (T_+ - T_0) \exp\{-a_2 \sqrt{p^2 - 4\tilde{q}} t_{tot}(t)\}} \quad (6)$$

with $T_0 = T_{in}(t - t_{tot}(t))$, $T_{\pm} = -\frac{p}{2} \pm \frac{1}{2} \sqrt{p^2 - 4\tilde{q}}$ and system-specific coefficients a_2 , p and \tilde{q} . The throughput time $t_{tot}(t)$ is a time-dependent quantity as well, which in principle could be calculated from time integration of the mass flow $\dot{m}(t)$. Since in practice this is difficult, it makes sense to promote it to a state variable, which is described by a separate ODE, for which a reasonable form is given by

$$\frac{dt_{tot}}{dt} = 1 - t_{tot}(t) \frac{\dot{m}(t)}{\rho A_{CS} L} \quad (7)$$

with fluid- and plant-specific constants ρ (fluid density), A_{CS} (cross-section area) and L (total length). To incorporate further dynamic effects in a way suitable for controller design, it has been suggested to include two further state variables, namely an effective solar radiation $I_{eff}(t)$, which yields an approximate description of the solar radiation that has acted on the fluid element so far, and a quantity $I_M(t)$, which describes the inertia effect of the metal parts. Those two variables are governed by two relatively simple ODEs:

$$\frac{dI_{\text{eff}}}{dt} = \frac{1}{\tau_I} (I_M(t) - I_{\text{eff}}(t)) \quad (8)$$

$$\frac{dI_M}{dt} = \frac{1}{\tau_M} (I'_g(t) - I_M(t)) \quad (9)$$

with suitable, system-specific time constants τ_I and τ_M .

In total, one has a system of three ordinary differential equations (7), (8), (9) and an algebraic equation (6), with coefficients that can be determined from the fluid and system properties. Using them requires knowledge of the time-dependent mass flow $\dot{m}(t)$, the global radiation $I'_g(t)$, the inlet temperature $T_{\text{in}}(t)$ and the ambient temperature $T_{\text{amb}}(t)$.

Such a formulation is suitable for the development of a controller, as it is described in Sec. 5.2.1.2 of (Unterberger, Modelling and control of large-scale solar thermal systems, 2021). The structure of the controller is shown in Figure 4, a comparison of simulation results in Figure 5. These results strongly indicate that the more advanced controller has superior performance, yielding output temperatures $T_{\text{out}}(t)$ closer to the desired set point values $T_{\text{out,SP}}(t)$.

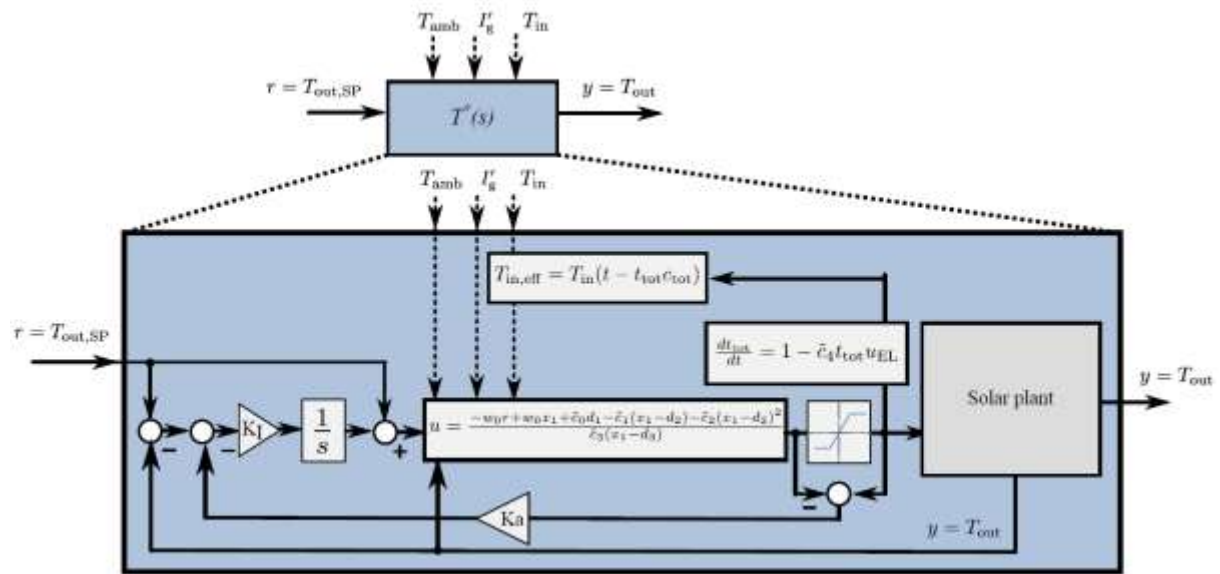


Figure 4: Control structure of an advanced model-based control for controlling the solar outlet temperature, Figure 5.12 from (Unterberger, Modelling and control of large-scale solar thermal systems, 2021)




Figure 5: Comparison of the results for a controller based on static equations (left) and a more advanced one based on dynamic equations (right), Figures 5.16 and 5.17 from (Unterberger, Modelling and control of large-scale solar thermal systems, 2021)

3.2 Control of Concentrating Collectors: Tracking Methods

Concentrating collector technologies, e.g. Linear Fresnel or parabolic trough, become increasingly interesting for mid- and high-temperature applications, see (Tamm & Berberich, 2024). These collector technologies allow to efficiently produce heat at high temperatures, also in regions with moderate climates, see (Krüger, Fischer, Hirsch, & Labairu, 2020). Even though this technology is continuously gaining relevance, no standard strategy for its control has yet been established. It is rather the case that the control strategies applied strongly depend on the used technology, its integration and the final utilization of the heat. However, the control strategies previously described for flat-plate collectors serve as a good starting point. An overview on the control for these systems typically used for process heat or electrical power generation is given in (Camacho, Berenguel, Alvarado, & Limon, 2021).

Control of tracking devices

Although tracking devices are used primarily for concentrating collectors, they also gain relevance for flat-plate collectors due to their decreasing prices as a result of their increased use in the PV sector, see e.g. (Nsengiyumva, Chen, Hu, & Chen, 2020). For this purpose, typically one-axis tracking devices are used to rotate entire collector arrays. This leads on the one hand to a higher solar share and provides, on the other hand, the possibility to stop the heat production of the plant by turning the collectors away from the sun. The second aspect allows to decrease the size of the heat storage, eliminates the problem of overheating and is especially of interest when the heat is provided to an industrial process, where the heat demand can be highly fluctuating, e.g. on weekdays vs. on weekends. The control of these devices is rather simple with the goal to follow the path of the sun as it is already known in advance. This is typically accomplished by PID controllers. An issue for tracking can be regions where strong wind disturbs the tracking, and controllers have to react to such disturbances. Besides the tracking aspects, the main control goal in these applications is still the control of the outlet temperature which is done in the same way as described in the previous section.

An alternative to traditional tracking methods are AI-based approaches, mainly Computer Vision (CV) methods based on Convolutional Neural Networks (CNN), as available almost out-of-the-box as open-source solutions from <https://opencv.org/>. Some success of this approach has been reported for example in (Carballo, Bonilla, Berenguel, Fernández-Reche, & García, 2019). Vision transformers offer an increasingly interesting alternative. Also, other Machine Learning methods like Support Vector regression have been used successfully (Mallikarathne, Abeysignhe, Rathnayake, & Perera, 2023). By now, several manufacturers claim to use AI-based methods, usually without giving details.

3.3 Further Aspects of Collector Array Control

Decoupled control of the flow in parallel collector arrays. In case of large subfields that are differently oriented and/or differently sized, it can make sense to install motor driven balancing valves which allow a dynamic balancing of the fields in order to achieve the same feed temperature, see e.g. (Unterberger, Innerhofer, Lichtenegger, & Gölles, 2018). In doing so, special care of the hydraulic correlations needs to be taken, e.g. the influence of a single valve on the remaining arrays and consequently the overall mass. A possible strategy to do this is outlined in (Unterberger, Muschick, & Gölles, Model-based control strategies for an efficient integration of solar thermal plants into district heating grids, 2017). Unfortunately, motor-driven balancing valves for large subfields are rarely installed due to the considerable cost of these appliances.

Predictive model-based methods for short times. If a constant outlet temperature is of high importance, it can be beneficial to use a model-based approach and predictions for the near future (the next few minutes). That way, fluctuations in solar gains can be compensated by changes in pump power, which change the transfer time. Obtaining constant temperatures (with only very small variations) is typically not an issue as soon as an intermediate storage tank is used, since fluctuations of the outlet temperature are averaged out by the storage.

Constant temperatures are rarely required for district heating with centralized feed-in. (For decentralized feed-in into the supply line, the situation may be different, with constant temperatures being of very high importance, especially in dense urban areas, where space for storages is limited, see (Jentsch, 2020).) Since such methods depend on high-resolution irradiation forecasts (obtained e.g. by real-time image processing for a camera pointed towards the sky), which are rarely available, we will not elaborate on this topic.

Combination of different collector types. As sketched in Figure 1 and discussed e.g. in (Jensen & Sifnaios, 2022), it can be beneficial to combine different types of collectors: While flat-plate collectors provide some pre-heating at lower temperatures, for which their efficiency is high, the final temperature lift is achieved in concentrating collectors that typically have a lower decrease of efficiency for higher temperatures. The control problem for such a two-stage approach is typically not that different from the one for a single-stage plant, since there are no additional actor-variables present; the whole fluid flow just sequentially passes through both stages. Predictive model-based methods for short times (as discussed above), however, may provide a larger yield gain for such setups, provided the increased complexity of the two-stage system is modelled in an adequate way.

Start-up and Shut-Down Process. In addition to the regular operation, the start-up and the shut-down of solar plants must be handled by the controller as well. Usually, this is done by simple step chain controls starting and stopping the actuators as soon as certain threshold levels for irradiation are achieved.

3.4 Notes on Further Relevant Technologies

3.4.1 Hybrid Collectors (PVT)

Hybrid collectors (Photovoltaic-Thermal collectors, PVT) are an interesting development and enable the double use of solar energy by extracting both electric and thermal energy. To some extent, however, there are conflicting objectives, since the efficiency of PV cells drops considerably at higher temperatures, typically 0.4-0.5% per Kelvin above the optimum temperature range (Affolter, et al., 2005). Thus, obtaining heat at high temperatures, as it is usually desired in solar thermal applications (in particular for industrial applications and district heating), conflicts with maximizing electric energy output. While for local appliances, the combination of PVT with compression heat pumps (see Sec. 3.4.3) can work well by upgrading the low-temperature heat from PVT to usable temperature levels, such methods seem not (yet) to be well-suited for traditional district heating. For low-exergy and anergy grid, the situation is different, though, and the optimal operation of PVT in a complex, cross-sectorial energy system constitutes an interesting (and in the future presumably important) problem, cf. Sec. 4.1.

3.4.2 Absorption Heat Pumps

Solar thermal systems are well-suited for the combination with absorption (or alternatively adsorption) heat pumps. This is not only the basic principle of solar cooling but also has applications for solar district heating. Standard methods for control of absorption heat pumps tend to be simple (linear independent controllers for the individual control loops) and thus yield suboptimal results. An advanced method based on multiple-input-multiple-output (MIMO) linearization has been developed in (Staudt, 2023).

3.4.3 Further Heat Sources

While solar thermal has great potential for sustainable heating, it has to be a supplement and complemented by additional technologies in order to provide heat at times when solar energy is not available. While thermal storage (see Sec. 3.4.4) can achieve some decoupling, typically additional heat sources are still required. The two most prominent sources are:

Compression Heat Pumps. Compression heat pumps are versatile heat sources that can make use of ambient and waste heat. Operation of heat pumps, however, requires some care. Various control strategies (for compression and absorption heat pumps, see Sec. 3.4.2) are discussed in (Goyal, Staedter, & Garimella, 2019).

Combustion. The most common heat source is still combustion of fuels (usually coal, hydrocarbons or biomass). While fossil fuels have to be phased out, combustion of renewable gases (biogas, synthetic methane from green hydrogen), liquid and solid biomass is still a viable sustainable source of heat, best to be used in combined heat and power (CHP) plants. While gas burners have very good response properties, combustion of solid biomass is a tough problem that greatly benefits from advanced control, see e.g. (Gölles, Reiter, Brunner, Dourdoumas, & Obernberger, 2014) and (Schörghuber, Gölles, Reichhartinger, & Horn, 2020).

Other heat sources include direct electric heating (usually to be avoided due to exergy aspects), industrial high-temperature waste heat, geothermal energy and waste incineration (a special type of combustion). While component control for some heat sources is already nontrivial, even greater challenges arise at the system level (see Sec. 4).

3.4.4 Thermal Storage

Sensible thermal storage, from buffer storage tanks for a few hours up to seasonal storage, are crucial for reaching considerable solar fractions. On the component level, control of charging and discharging of thermal storages is rather straightforward (and usually handled by the controller of the main heating system). The main challenges, as in the case of auxiliary heating, arise at system level (see Sec. 4), ranging from reliable state estimates to operation planning for one or two days and to robust strategies for seasonal operation, taking into account the necessity of auxiliary heating to make up for significantly decreasing temperatures during long-term operation cycles.

Other types of thermal storage such as phase change materials (PCM), sorption and thermochemical storage technologies pose greater challenges for control, in particular since the state is harder to estimate, but they are still rarely used for solar district heating.

3.5 Control of the Heat Exchanger and the Secondary Circuit

The usual operating strategy of a heat exchanger aims at maximizing the efficiency, i.e. to minimize the sum of the terminal temperature differences on the cold side ΔT_c and on the hot side ΔT_h , see Figure 6.




Figure 6: Heat exchanger connecting the primary (collector) circuit and the secondary circuit

This is achieved by keeping the heat capacity flows the same at the primary (P) and the secondary (S) side of the heat exchanger, i.e. the products of the respective mass flows \dot{m}_i and heat capacities $c_{p,i}$,

$$\dot{m}_P c_{p,P} = \dot{m}_S c_{p,S} \quad (3)$$

A constant value for the feed temperature in the secondary circuit (i.e. the outlet temperature of the heat exchanger on the secondary side) $T_{S,out}$ is typically achieved by applying a simple cascading control structure, where the outer slower controller adjusts the set point for the feed temperature in the primary (solar) circuit T_{feed}^* , i.e. $T_{P,in}$, so that the feed temperature in the secondary circuit $T_{S,out}$ achieves the desired value. This is depicted in Figure 7, which is an extended version of Figure 3.

In some cases, the feed temperature in the secondary circuit is not directly controlled. In that case, only the feed temperature in the primary circuit is subject to a strategy control, with the set point being a certain offset above the feed temperature required in the secondary circuit.

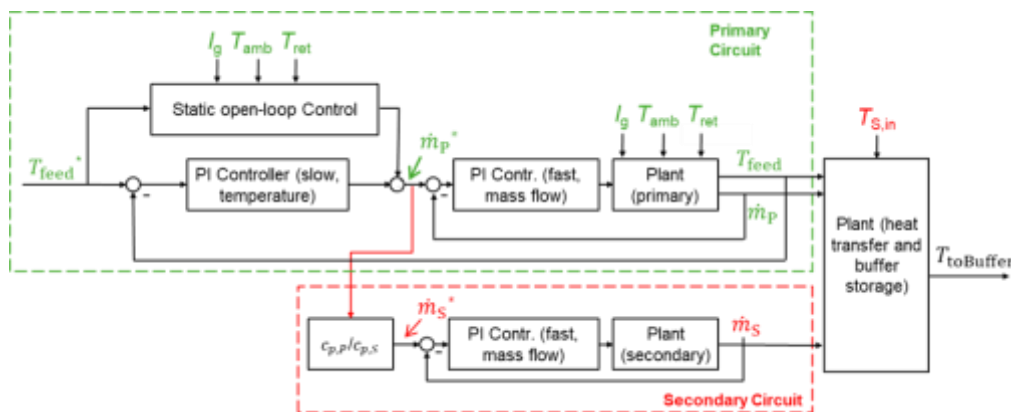


Figure 7: Static feed forward control with integrating controller, combining primary and secondary circuit.

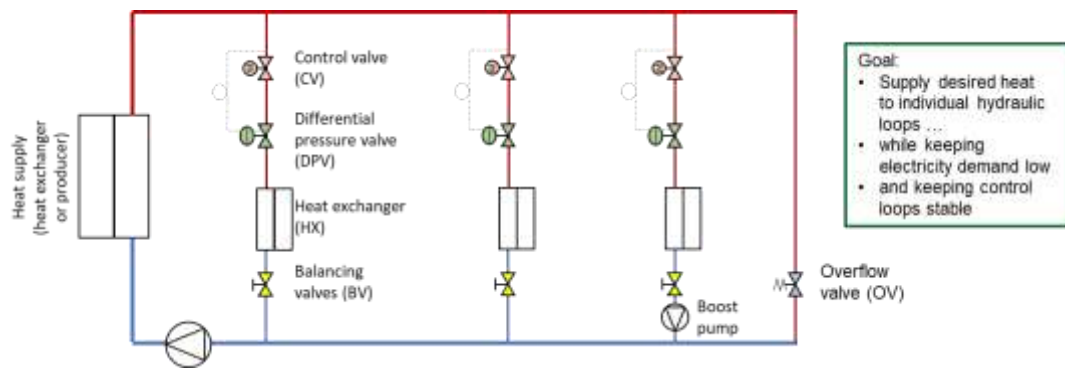


Figure 8: General structure of hydraulic distribution networks, with different types of valves.

3.6 Interaction with the Control of District Heating grids

Solar thermal plants can be quite complex hydraulic networks, and the complexity is even higher in many heat distribution networks, as sketched in Figure 8. The interplay of pumps and different valves have to be considered for this, and concepts like *valve authority* are important – which are beyond the scope of the present report. Still, in the following paragraphs, we will sketch a rough picture of the control challenges of integrating large-scale solar thermal into district heating grids.

The control of heat distribution networks must, on the one hand, adjust the feed temperatures of the different producers to provide sufficient heat while keeping losses low, and, on the other hand, adjust the pressure and mass flow conditions accordingly to the load situation. Standard methods for this are well-established for traditional heating grids, i.e. mostly the 2nd generation of the systematic introduced in (Lund, et al., 2014), but they often struggle with complex networks with several producers, as characteristic for 3rd generation and later ones. As soon as there is not a single dominant producer, pressure control becomes difficult, and there is the risk of pumps “competing against each other”, raising the overall pressure level and energy consumption without actual necessity.

In principle, later generations of heating grids with lower feed temperatures are better suited for integration of solar thermal plants, even those not operating at higher temperatures. There are several ways for the hydraulic integration, as discussed in some detail in Ch. 8 of (Dénarié, et al., 2025).

As outlined in (Gölles, et al., 2021), the distribution network, i.e. the hydraulic network connecting the thermal storage with the consumers, highly depends on the applications and can be very different from plant to plant. In some cases, only a simple hydraulic circuit connects the storage with on-site consumers while in other cases the hydraulic circuits can be very complex connecting several consumers in parallel or serial where each of the consumers have different needs regarding mass flows, temperatures and consequently heat flows. For this reason, it is not possible to propose a common strategy for these tasks. However, for example in (Unterberger V. , et al., 2020) a method aiming for a decoupled control of mass flow and temperature is described, which could get adapted to different specific configurations.

4 System-Level Control

In the present section, we discuss various approaches for the high-level (or system-level) supervisory control of the whole system (sometimes also called *energy management system*, EMS¹), i.e. the task of creating schedules for complex energy systems, typically by taking into account predictions for energy yield and consumption (instead of just making decisions based on the present state). For SDH, this system view is sketched in Figure 9.




Figure 9: Solar district heating system with the core solar plant subsystem highlighted

4.1 Basic Aspects and Considerations

4.1.1 Main Objectives of Superordinate Control

Superordinate control should operate the single elements (by turning them on or off, or by communicating set points to the low-level controls) in a way that ensures a desired mode of operation for the whole system. But what characterizes the desired mode of operation? Typically, there are several objectives that should be met, but are sometimes in conflict with each other:

1. Be able to provide the heat (or, in general, the energy) needed by the consumers at any time,
2. Minimize the cost of operation (OPEX), mostly fuel and electricity cost,
3. Minimize harmful emissions (toxic or climate-relevant),
4. Minimize the stress on the appliances to extend the lifespan (and effectively reduce CAPEX),
5. Resilience, being able to cope with unexpected issues (failure of components, prognosis errors, ...)

Typically, one finds (1.) as the primary and (2.) as the secondary objective, with (3.)-(5.) being far less prominently considered. Fortunately, for SDH, objectives (2.) and (3.) go along very well, since maximizing the use of solar energy usually both reduces the cost for other forms of energy and the associated emissions.

In rule-based approaches (Sec. 4.2), the objectives are usually implicitly incorporated in the rules, e.g. by setting thresholds in a way that guarantees sufficient heat supply, while keeping fuel costs and strain at bay. In optimization-based approaches (Sec. 4.3) and for Reinforcement Learning (Sec. 4.4.2), one has a more systematic way by defining a cost function, which is minimized (or a reward that is maximized). Violation of objective (1.) is typically either prohibited by hard constraints or heavily punished by corresponding terms in the cost function.

¹ Note, however, that, while superordinate control is indeed a viable method for improving energy management strategies, this use of the term only partially overlaps with the one in ISO 50001.

4.1.2 Planning and Superordinate Control

Superordinate control should already be considered when planning energy systems. On the one hand, the performance of system-level control influences the dimensioning of system, on the other hand, superordinate control requires some degrees of freedom to work effectively. These can be provided by storage devices, but also by “over-sizing” of components, as compared to standard dimensioning practice. Such “over-sizing” of some components, while raising the required CAPEX, can subsequently lead to reduced OPEX, since the use of fluctuating energy can be improved.

4.1.3 Forecasting

Most advanced high-level control strategies are *predictive*, i.e. they make use of forecasts for solar gains, heat demand, electricity prices etc. Typically, forecasts for directly relevant quantities are created for a horizon that is typically chosen to be 24 to 48 hours. Longer prediction horizons are possible, but typically, the temporal resolution is reduced for periods of times further in the future. Usually, a weather forecast is assumed to be available, since such forecasts are available at low cost (or, with limited functionality, even for free) from various providers such as Ubimet, Weather Underground, OpenMeteo, Meteoblue and national meteorological institutions.

Several methods can be used for the yield forecast:

- There are very simple approaches such as the persistence method (assuming everything to be the same as the day before at the same time) or statistical averages over / linear extrapolation from corresponding values of the last few years, assuming that the data is available. Such methods are not very accurate, but quite robust and can serve as a baseline.
- Classical time series analysis offers various Auto-Regressive (AR) and Moving Average (MA) methods, like ARMA, ARIMA and ARIMAX, by combining AR with MA, possibly including Integration (for non-stationary time series) and eXternal factors, see e.g. (Hamilton, 1994).
- Statistical learning / shallow machine learning (ML) methods include, among others, multi-linear and quasi-linear regression, regression tree learning, support vector regression, probabilistic graphical models (PGMs) and multi-layered perceptrons (MLPs).
- Deep learning methods include Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), 1-dimensional Convolutional Neural Networks (CNNs), Transformers and Sequential State Space Models (SSSMs), see e.g. (Murphy, Probabilistic Machine Learning: An introduction, 2022), (Murphy, Probabilistic Machine Learning: Advanced Topics, 2023) or (Bishop & Bishop, 2023).
- Physics-based models may be used to estimate gains, demand, and costs based on the weather forecast and physical models of the components. These approaches may also be combined with machine learning resulting in physics-informed machine learning models.
- An adaptive method which is powerful, but light-weight enough for implementation on programmable logic controllers (PLCs), has been introduced in (Nigitz & Gölles, 2019) and (Unterberger, Lichtenegger, & Gölles, Predictive Rule-Based Control Strategy for Optimizing the Operation of Solar District Heating Plants, 2024); it is also discussed in some detail in appendix A of (Gölles & Unterberger, Control of large-scale solar thermal plants, 2021).

The actual choice depends on the data available (with deep-learning methods being notoriously “data-hungry”), computational effort for training, computational effort for making predictions (inference) and the ability to incorporate new data on the fly (adaptive learning, online learning).

4.1.4 Resilience, Safety and Security

Maintaining resilient, safe and secure operation is an important issue.

Resilience can be achieved in planning by building redundancies (possibly by establishing an $N - 1$ approach, where any single component can fail without breakdown of the system), in operation in particular by stochastic optimization, doing Monte-Carlo simulations for different scenarios.

Aspects of **safety** include plausibility checks for sensor signals, initiation of emergency measures (usually handled on the level of low-level control for time-critical issues) and fallback solutions for low-level control, in case connection to the superordinate control is lost or certain plausibility checks fail.

For **security**, in addition to classical aspects, *cybersecurity* aspects are rapidly gaining importance. High-level control typically requires some internet access (for obtaining weather forecasts, possibly also electricity tariffs), and increasingly, data collection and data processing task are handled with cloud solutions.

4.2 Rule-based Supervisory Control

The state of the art is to derive the respective operating mode of the different plants and components from measured variables, and if available, other information, e.g. forecasts, by means of simple rules. This includes, for example, the adjustment of the set point for the feed temperature levels according to the weather conditions, typically based on the ambient temperature, or to decide on a specific operating strategy for the operation of on-site thermal storage(s). The operating strategies applied vary considerably among different applications and especially among different technology providers. Still, a distinction can be made between two basic approaches:

1. Strategies only considering the current state of the system
2. Strategies additionally using forecasts

Examples for these 2 approaches are examined in more detail in the following based on a typical, widely spread configuration consisting of a decentral collector field, an on-site buffer storage, on-site consumers and a bi-directional connection to a DH network. A simplified scheme of this configuration is given in Figure 10, however, the possibility to feed into the DH network is often realized directly, i.e. not via the buffer storage, which is displayed only in a simplified way in this scheme.




Figure 10: Scheme of a typical large-scale solar thermal system with a decentral, bi-directional connection to a DH network

4.2.1 Classical State-based Rules

The decision on the mode of operation finally comes down to two (possibly three) main decisions, first whether the collector field should be operated or not, second whether heat should get fed into or obtained from the DH network, and possibly third at which position of the buffer storage the heat from the collector field should get fed in, in case the buffer storages have different inlets. Indirectly these decisions also determine the temperature levels the different plants and components are operated at, however, these levels are typically control parameters or externally imposed, e.g. by the consumer or the DH network. The heat consumption takes place anyway, so in simple, state-of-the-art configurations (without demand side management) no decision must be taken to this end.

The decision on starting or stopping the operation of the collector field is taken based on the current levels of measured temperatures among the collector field and the state of the buffer storage. In both cases, deciding to start or stop implies a certain sequential control for gradually switching on/off the different actuators as certain temperature levels are achieved, finally ending in normal operation mode or no operation of the collector field.

The decision on the interaction with the DH network, i.e. whether to feed into the DH network, to obtain heat from the DH network, or neither of these, is taken based on the current state of the buffer storage. This is typically done very simply, by starting to feed in respectively to draw when certain thresholds of the state of the buffer storage are reached.

The decision on the inlet position to be used to feed into the buffer storage is also taken based on the current state of the buffer storage. This is typically simply done by switching between the inlets when certain thresholds of the state of the buffer storage and the current temperature level provided by the collector field are achieved.

In all these decisions the evaluation of the current state of the buffer storage is very important. The current state is in most cases only determined by the current values of vertically distributed temperature sensors. The next step of complexity is to use the available temperature levels to estimate the temperature levels in between the measurement positions and to use the resulting temperature distribution to estimate the current amount of heat stored and possibly also the current exergy level. The best estimation of the current state of the buffer can be achieved by additionally simulating the vertical temperature distribution online using a suitable mathematical model, and to continuously correcting the current simulated state by the currently available measurement values. However, such complex and sophisticated approaches are typically only used when also more complex control strategies are applied.

The threshold levels used for the decisions described above typically vary with the seasons of the year, indirectly considering to some extent the conditions of the near future, and that the required temperatures for feeding into the DH network, typically fixed by contracts, are always given. Often a certain correlation to the ambient temperature is fixed, but it could also be the case that the set point is defined by the supervisory controller of the entire DH network.

Even though these control strategies are based on comparatively simple rules, it could be reasonable to optimize these strategies by determining the optimal values for the different threshold levels through numerical simulation studies and possibly even by numerical optimization algorithms. This is not state-of-the-art yet among the most solar thermal systems connected to DH networks but is getting more common and particularly used by application-oriented R&D projects.

4.2.2 State-based Fuzzy Rules

Standard state-based rules, as described above, are formulated as “hard” (often binary) decisions. An alternative (and for some applications superior) approach is based on fuzzy logic. This way, expert rules, gained from experience (like “If the tank is already rather full, the local demand is low, the solar yield is on average level and almost stable, feed a moderate amount of heat into the grid”), can be codified as fuzzy sets. Ideally, the whole state space of the system is covered by overlapping patches of such rules. In this case, methods like the Mamdani-Assilian controller (Kruse, et al., 2013) can be used to obtain a consistent control strategy. For simulations of multi-energy systems that also contain solar thermal plants, such approaches seem to be promising, see e.g. (Ma, Li, Jin, Cao, & Zhao, 2022)

4.2.3 Prediction-Based Rules for Supervisory Control

As discussed in Sec. 4.1.3, forecasts have the potential to significantly improve the performance of system-level control. While such forecasts are most often incorporated in optimization-based control strategies (see Sec. 4.3), it is also possible to improve rule-based approaches by using such forecasts.

An example for a rather simple rule-based supervisory control strategy that can be implemented on a standard programmable logic controller (PLC) has been described in (Unterberger, Modelling and control of large-scale solar thermal systems, 2021) and (Unterberger, Lichtenegger, & Gölles, Predictive Rule-Based Control Strategy for Optimizing the Operation of Solar District Heating Plants, 2024). The setup, composed of a large-scale solar thermal field, a storage, on-site consumers and a direct connection to a local district heating grid (DHG) owned and operated by another organization, is depicted in Figure 11.

Via the connection to the DHG, the SDH plant can generate additional profit by feeding the solar heat into the grid in case the production exceeds the needs of the on-site consumers. Vice-versa, in case there is not enough solar heat available, the grid acts as auxiliary heating, and additional heat can be purchased to avoid comfort losses for the on-site consumers. Heat flows and corresponding costs / revenues are shown in Figure 12.




Figure 11: Schematic representation of the Solar District Heating system discussed in Sec. 4.2.3




Figure 12: Heat flows and corresponding costs / revenues for the system discussed in Sec. 4.2.3.

Therefore, simply speaking, this kind of SDH plant can run in two general operation modes:

- **HSt**: transfer the solar heat to the local heat storage in order to sell it to the on-site consumers.
- **DHG**: transfer the solar heat to the district heating grid to directly generate profit.

The goal of the algorithm is to decide when to run in mode = HSt or mode = DHG. The number of mode switches is to be minimized, while the storage is managed in a way such so that ...

- the on-site demand can be satisfied with solar heat, if possible,
- no excess heat is sold to the DHG and must be bought back later for a higher price,
- only required heat is stored, in order to minimize storage losses.

In order to determine the optimal operation schedule, the following steps shown in Figure 13 are executed for a forecast horizon of 24 hours (and periodically repeated to update the strategy in order to incorporate new information). An example of its application is shown in Figure 14 and Figure 15.

Of course there are many possibilities to improve this algorithm, e.g. by explicitly taking into account storage losses or by augmenting the discrete decision by prescribed heat flux levels.

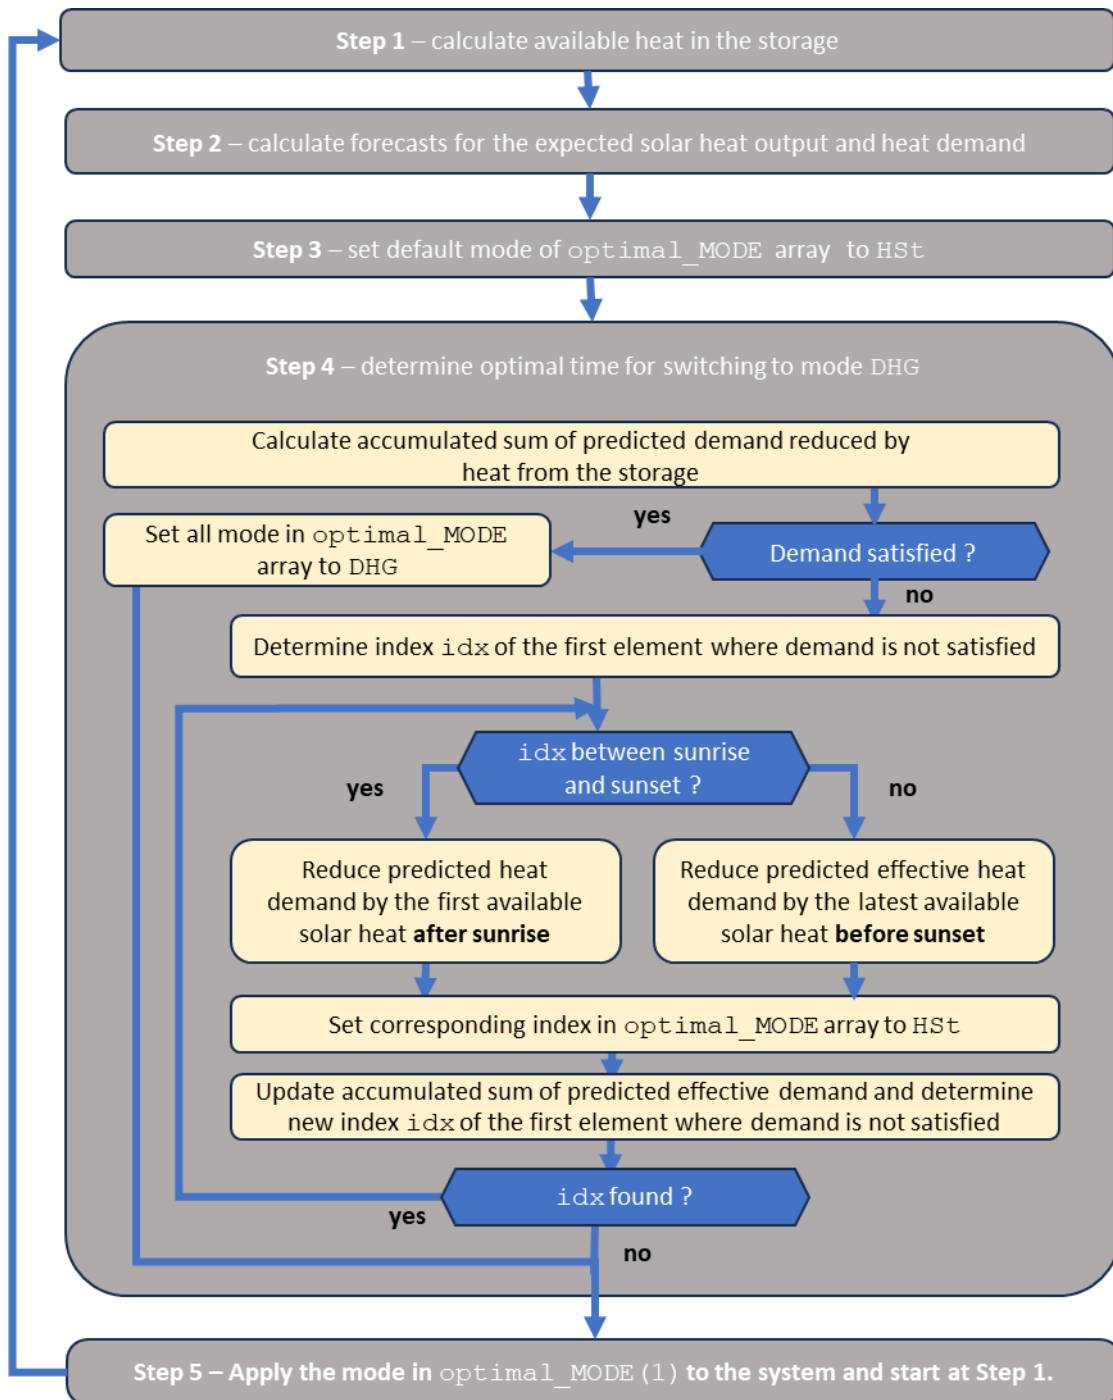


Figure 13: Flow-chart for the proposed algorithm, where the mode values HSt or DHG are stored in an array `optimal_MODE`




Figure 14: Example for the action of the proposed algorithm, Step 3 (all heat from storage is used to reduce the effective consumption as early as possible, in order to reduce storage losses)

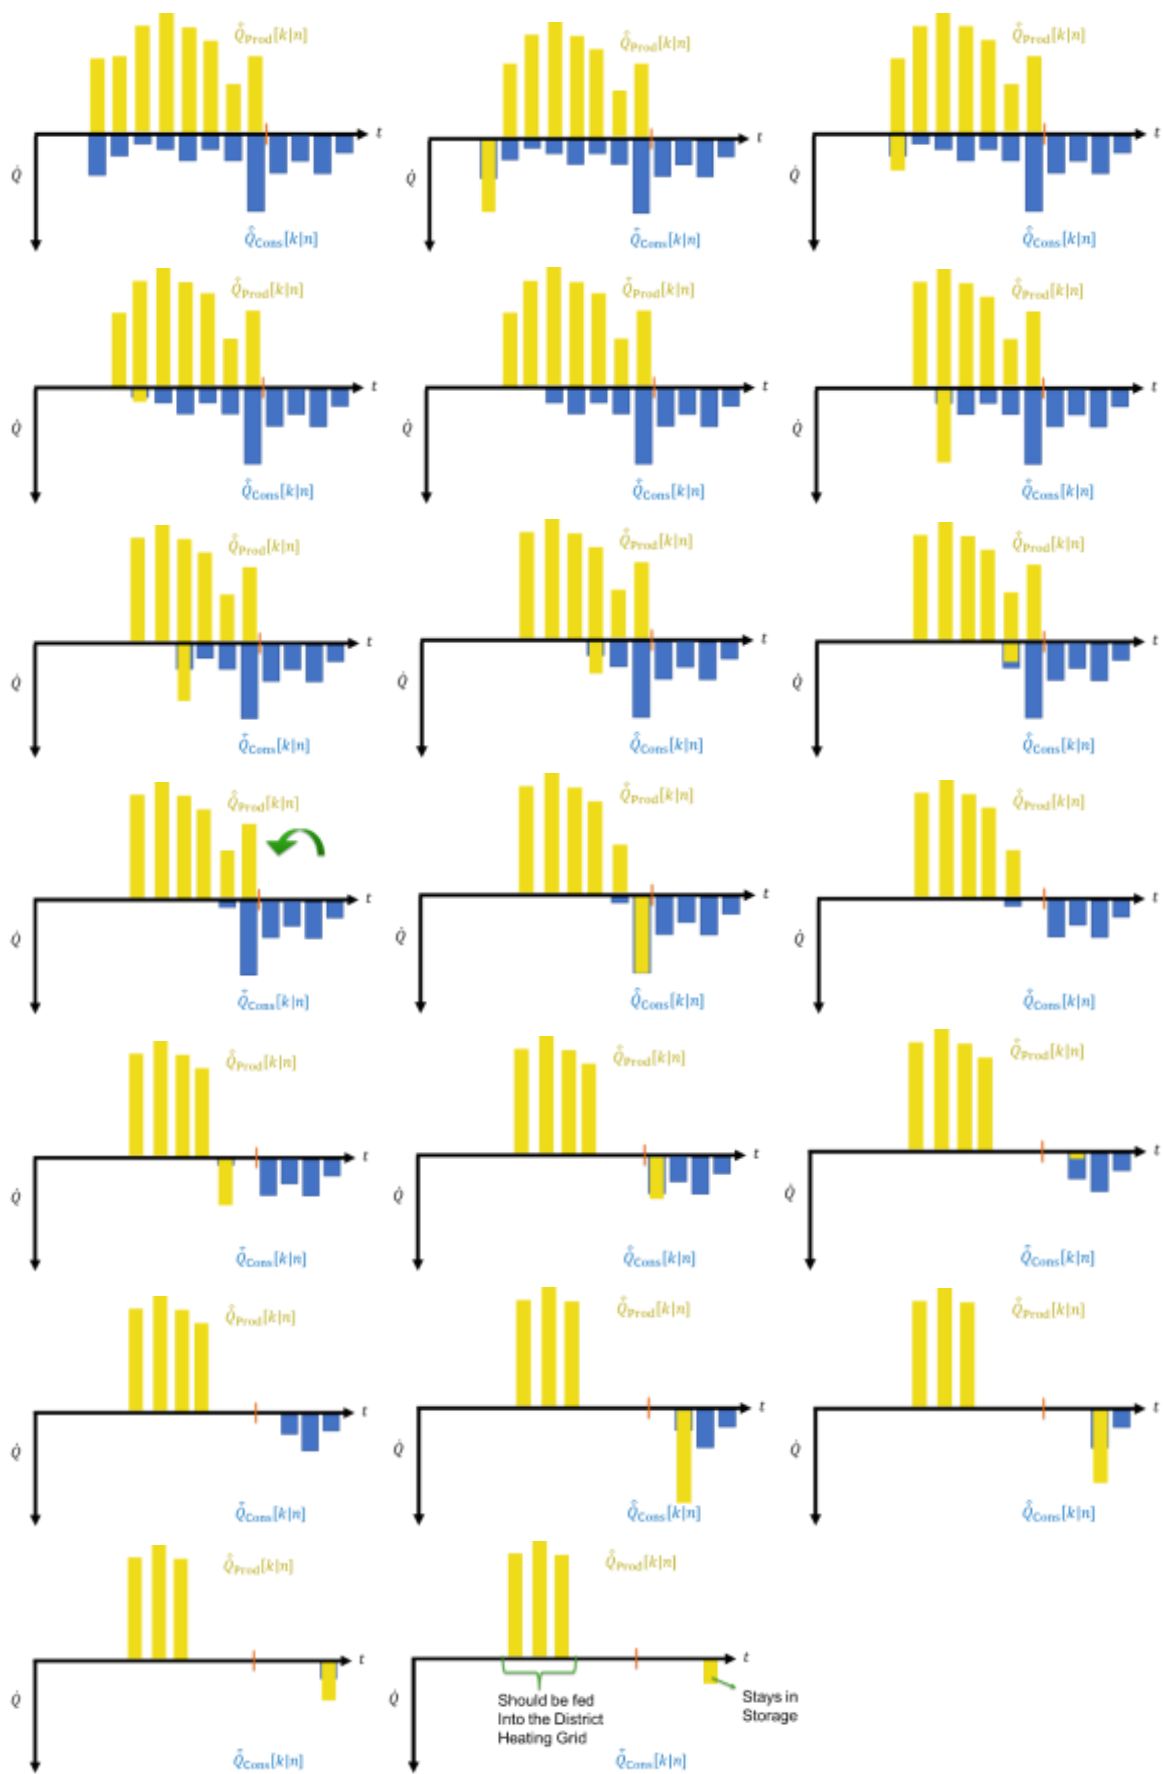


Figure 15: Example for the action of the proposed algorithm, Step 4 (predicted production \hat{Q}_{Prod} is matched against predictive consumption, in order to reduce the effective consumption \hat{Q}_{Cons})

4.3 Optimization-based Control Strategies

4.3.1 Centralized Optimization-based Predictive Supervisory Control

A systematic approach for the supervisory control is the application of optimization-based predictive supervisory controllers, i.e. control strategies based on solving a mathematical optimization problem and the consideration of knowledge on future boundary conditions. The available approaches for optimization-based predictive supervisory control are not limited to solar thermal systems; rather they aim to cover as many technologies, energy sectors, storages, etc., as possible. The coupling of the different energy sectors will continue to increase, and it is obvious that a joined consideration of all producers and consumers, storages, distribution grids, and all coupled sectors theoretically must lead to the best overall operating behaviour. Because of the many possibilities of the detailed structure and configuration of the different hybrid energy systems to be considered, it is not reasonable to develop specific approaches for single sectors, but instead holistic approaches should be pursued. This is already the case for most of the approaches currently available. Even if the different approaches are often clearly derived from specific sectors, the general principles used have a common basis. For this reason, the approaches for optimization-based predictive supervisory control are presented in a very general manner in the following, and the aspects specifically important for large-scale solar thermal systems are highlighted appropriately.

These optimization-based predictive supervisory controllers are often referred to as energy management systems (EMS). Their main task is to control the entire energy production (among all plants and sectors) while fulfilling the demand of the different consumers and ensuring that all restrictions are fulfilled, e.g. the storage capacities of thermal buffer storages or electrical batteries and the transport capacities of pipes or power lines.

The idea of optimization-based control is to formulate the control problem as an optimization problem, which then is periodically solved. Casting the control problem into an optimization problem requires the formulation of two main parts: First, the dynamics of the considered (hybrid) energy system must be described through the constraints of the optimization problem, and second a so-called cost-function of the optimization problem penalizing or rewarding a certain operating strategy must be formulated. This optimization problem is then solved, leading to an optimal control schedule for a given horizon, e.g. 24 or 48 hours. However, only the first instance of the control signals is applied to the system, strictly speaking as set points of the lower-level controllers of the integrated plants and components. After a certain period, e.g. 15 min or 1 hour, the solution of the optimisation problem is repeated, using updated measurement values, i.e. the current state of the systems, and forecasts. This repeated update and solution of an optimization problem is typically referred to as *moving horizon model predictive control (MPC)* approach in control theory.

A schematic overview of the structure of such an optimization-based predictive supervisory control (Energy Management System) is given in Figure 16.

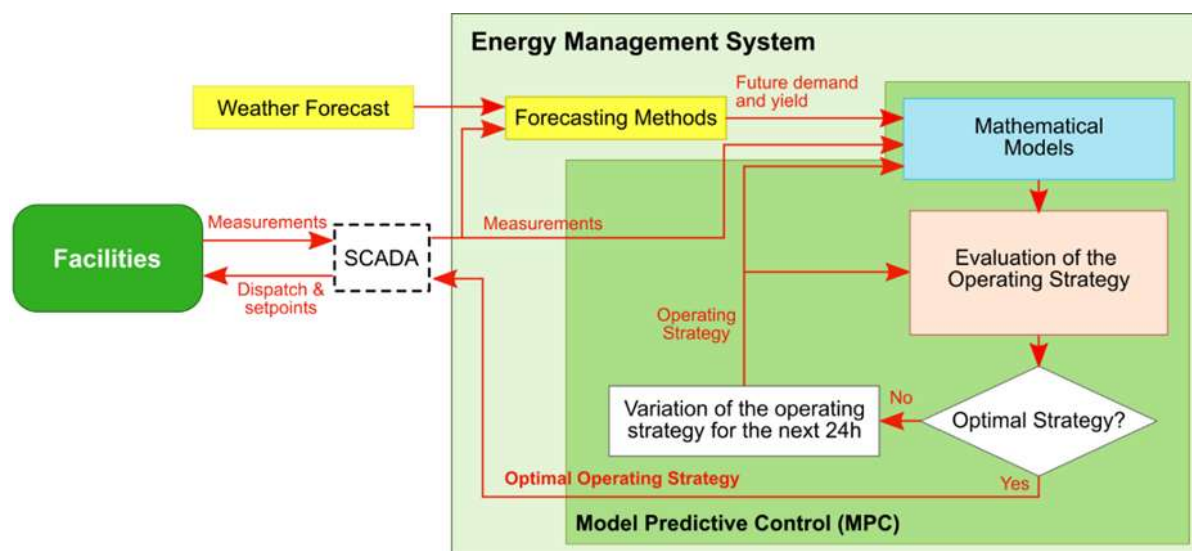


Figure 16: Structure of an optimization-based predictive supervisory control (Energy Management System) and its interfaces to forecasting and the supervisory control and data acquisition (SCADA)

The key difference between different approaches consists of the underlying **mathematical model**:

First, *linear* or *non-linear models* could be used - consequently leading to linear and respectively far more complex and computationally expensive non-linear optimization problems. Since the control must be real-time capable, even though the step-sizes are comparatively large, a trade-off between model complexity and speed must be made. Non-linear approaches would allow for more complex and thus more precise models, however, obtaining optimal solutions for the resulting non-linear optimization problems is very hard. For most non-linear optimization problems, the optimality of the result cannot be guaranteed. Therefore, linear models are more widely used in both industry and academia.

Second, certain discrete decisions must be taken by the supervisory controller, e.g. the decision between two operating modes or simply the decision whether a certain controllable plant, e.g. a gas boiler, should be in operation or not. To do so additional Boolean or integer variables need to be added to the models, leading to *mixed-integer linear programmes* (MILP), which are currently the most common approach for optimization-based predictive supervisory control of hybrid energy systems in both industry and academia, see e.g. (Moser, et al., 2020), (Bergh, Bruninx, Delarue, & D'haeseleer, 2014), (Verrilli, et al., 2017). Many years of research and development on the application of MILP have devised fast algorithms like *branch-and-bound* or *branch-and-cut* as well as efficient heuristics that are able to derive proven optimal solutions for such MILPs under tight time constraints, see e.g. (Achterberg & Wunderling, 2013), (Gurobi-Optimization, 2019), (IBM, 2019). To allow for a wide and efficient application the derivation of the mathematical models and the optimization problem needs to be largely automated, see e.g. (Moser, et al., 2020), in order to decrease the effort necessary for the development and implementation of such a control at a certain system.

The reduced complexity and thus accuracy of the models, however, goes along with disadvantages. The decision variables correlating to the different production units are typically the demanded power / heat flow or the mass flows corresponding to a certain temperature spread. In the first step, constant efficiencies over the entire operating range are assumed. Especially in the thermal sector these two assumptions, fixed temperature levels and constant efficiencies imply significant inaccuracies:

1. First, efficiencies of certain producers often depend on the load they are operated at. However, this can be modelled sufficiently well within the MILP framework by using piece-wise-affine (PWA) functions, see e.g. (Moser, et al., 2020).
2. The use of fixed temperature levels implies a much bigger problem, since efficiencies of certain producers, losses, and in particular the yield of solar thermal plants strongly depend on the temperature levels they are operated at. For sure, the most accurate approach to model all these phenomena would be to go for non-linear models and consequently to mixed-integer non-linear programming (MINLP). The authors, however, do not consider this to be a practically suitable, widely and systematically applicable approach, since the complexity and computationally effort increases dramatically, and what is still more problematic, the optimality of the result cannot be guaranteed.

A reasonable approach to overcome this problem while remaining within the MILP framework is presented in (Muschick, Zlabinger, Moser, Lichtenegger, & Gölles, 2022), describing the idea of heat flows with a set of mass flows at different, constant temperatures. A promising approach could also be to benefit from the extension of commercial solvers to non-convex mixed-integer quadratically-constrained programming (MIQCP) problems, e.g. (Gurobi-Optimization, 2019), allowing for direct modelling of heat flows at variable temperatures and therefore allowing the EMS to compute optimal control schedules not only in terms of energy flows, but also at which specific temperatures. However, these investigations are at an early stage of research at the moment.

Of course it is, in principle, also possible to use any (black-box) model together with an arbitrary minimization algorithm (gradient-based methods, Newton and Newton-like descent methods, Nelder-Mead, swarm methods, simulated annealing, evolutionary algorithms, ...). While approaches are very flexible, the performance is usually far worse than for those with rather simple models and solvers designed and optimized for them.

The optimal prediction horizons and time intervals the optimization problem should be resolved strongly depend on the detailed configuration, in particular on the storage capacities. In most cases longer horizons lead to better results, but also more complex and difficult to solve optimization problems. Especially for very large storages, i.e. seasonal storages, this must be considered. In this case a cascading combination of a mathematically less complex, long-term optimization solved at longer intervals, and a more detailed and more frequently solved short-term optimization may be a sufficient approach.

An increasingly important aspect of predictive optimization is robustness of operation. *Stochastic optimization* allows to identify strategies that may perform slightly worse in terms of efficiency, but can cope better with disturbances, faulty data and modelling errors.

4.3.2 Decentralized and Agent-based Methods of Supervisory Control

Due to the rather unfavorable scaling properties of most optimization methods (like MILP or MIQP), centralized optimization is often not applicable to complex systems. If a single entity is in the authority of all producers, this is essentially an algorithmic problem, which can be addressed with methods from distributed optimization, like the Alternating direction method of multipliers (ADMM), which links subsystems via constraints that are considered with Lagrange multipliers, see (Boyd, Parikh, Chu, Peleato, & Eckstein, 2010).

Usually, however, in increasingly complex energy system, one tends to encounter several (semi-)independent actors. For such energy systems, individual optimization can be combined with local market-based interactions to reduce the total cost, (Lichtenegger, et al., 2020), see Figure 17 for a two-level hierarchy with prosumers and clusters, connected by a superordinate grid. Such an approach can be interpreted as an optimization method, based on swarm intelligence, which takes advantage of the emergent behaviour produced by interacting agents.

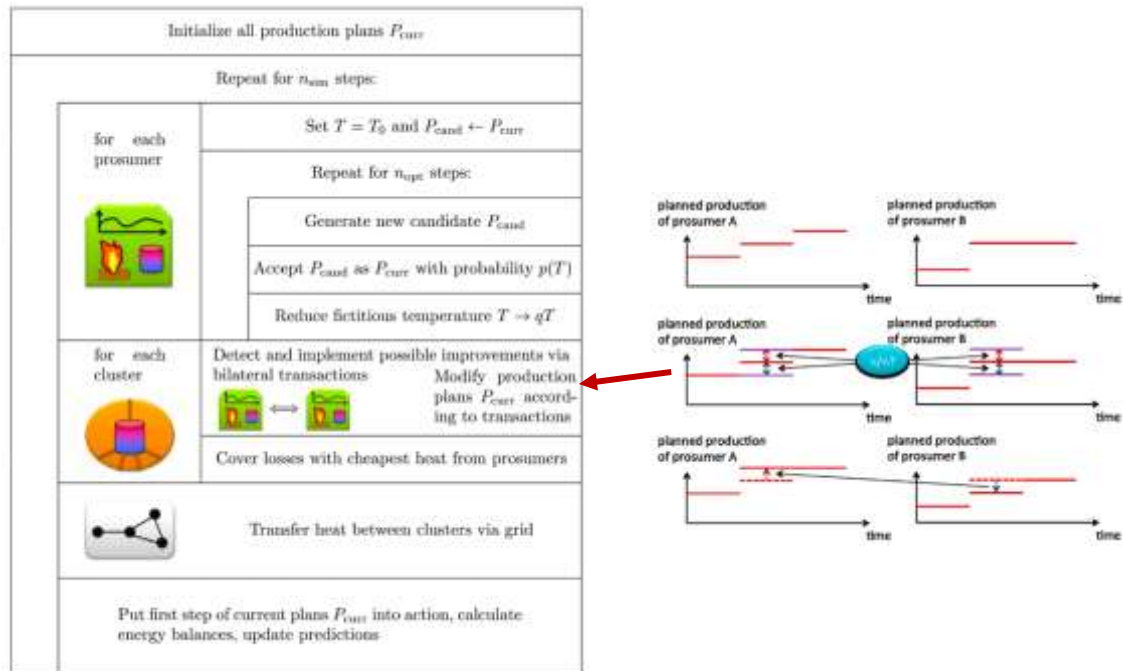


Figure 17: Combination of individual optimization (realized via simulated annealing with annealing parameter q as optimization method) with local market-based interactions, adapted from (Lichtenegger, et al., 2020).

Typically, the resulting complex, nonlinear, optimization problems lack useful properties like convexity. Often, as illustrated in Figure 18, many local minima of the cost function exist, and usually there will be no possibility to check whether the global minimum has been found. This is not a severe restriction, however, as long as a "good" minimum is found – which is usually sufficient, with differences between the costs for different minima being quite often smaller than the uncertainty occurred by modelling and forecast errors.




Figure 18: Illustration of a function with many local minima

4.4 Machine-Learning and AI-based Control Strategies

4.4.1 Supervised Learning

Two main branches of Machine Learning (ML) are *Supervised Learning* (learning to predict a target variable, either a continuous one in regression or a discrete one in classification) and *Unsupervised Learning* (finding patterns in data with no explicit target variable given). While also for control purposes, unsupervised learning can have applications, e.g. for anomaly detection or for visualization purposes (by using dimensional reduction techniques), the two main tasks connected to control belong to the branch of supervised learning:

- **Forecasts** are traditional tasks for supervised learning, see Sec. 4.1.3. In the context of solar district heating, this may be forecasts of solar production, of heat consumption or of time-variant electricity tariffs (in particular when integrating heat pumps, combined heat & power or PVT, see Sec. 3.4.1).
- **Data completion**, which is a necessary preprocessing step for many tasks, including those from control, can be regarded as a generalization of the supervised learning approach. Instead of learning a functional dependence of a single variable on all others (“glorified curve-fitting”), one learns a distribution of variables and predicts those that are missing based on those that are available.
- Since control required finding suitable actor values based on observations (sensor signals) and possibly also on predictions, **establishing a control strategy** can be regarded as a supervised learning problem. Obtaining enough data that covers all relevant situation and system states, however, is a highly nontrivial task, and typically, simulations have to be used in addition to data from plant operation. In order to learn a control strategy, a controller has already to be available. Thus, at first glance, there seems to be little use in a supervised-learning approach to the control task. Since, however, repeatedly solving an optimization problem is computationally expensive, it might make sense to replace an optimization-based controller by a properly trained ML model, called a *surrogate*, which can be evaluated much faster.

For supervised learning approaches, a whole toolbox of methods is available, including simple statistical methods like (quasi-linear) regression and Naïve Bayes classification, decision tree learning, ensemble methods like Random Forests and Boosting methods, support vector machines (SVM) and various architectures of neural networks, e.g. multi-layered perceptrons (MLP), Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), possibly with extensions like Long Short-Term Memory (LSTM), or restricted Boltzmann machines. While such methods are used with considerable success for forecasting tasks and some of them also for data completion, the barrier for successfully training and using surrogates is still high, while the level of explainability is low. Thus, this is still largely a domain of fundamental research, not a readily applicable technology.

4.4.2 Reinforcement Learning

The branch of ML most suited for development of control strategies is *Reinforcement Learning* (RL), where training is not based on previously available data, but done in an environment (usually a simulated one). An agent takes actions, which change the environment and yield some reward, which may also be negative for some situations, (Sutton & Barto, 2018). The learning task usually is to maximize the cumulative reward on the long run. During training, a delicate balance between *exploration* (testing new strategies) and *exploitation* (making use of discovered strategies that work well) has to be made.

Even though RL has made enormous progress during the last few years, it is still rarely used in productive systems (apart from being one element in the training of Large Language Models, LLMs). This is mainly due to the following reasons:

- While modern techniques like Proximal Policy Optimization (PPO) have led to more reliable learning, as compared to older approaches like Deep-Q learning, there is still no guarantee for the method to work, and effects like catastrophic forgetting may still occur.
- It is not practical to train RL algorithms with actual installations (both due to the limited amount of training that is possible when working in real time and due to poor performance and possibly even destructive operation modes during the exploration phase). Due to modelling errors, however, algorithms trained with simulation models, will in general perform worse on the real plant.
- RL approaches are “black box” methods, lacking transparency and explainability – which is highly problematic for control of critical infrastructure, as it is energy supply.

4.5 Supervisory Control: Conclusions

The three approaches discussed in Secs. 4.2 to 4.4 have strengths and weaknesses, which are summarized in Figure 19.

Aspects	Rule-Based	Optimization-Based (Model Predictive Control – MPC)	AI-Based (Reinforcement Learning – RL)
Principle	IF-THEN logic, static functions	Minimize cost under constraints	Agent learns via simulation environment
Models	Usually none; good knowledge about the system required	Typically, simple linear models, nonlinear models possible (but costly)	Environment: flexible dynamic simulation models, usually simplified
Setup Effort	Sound domain knowledge required; low coding effort	Usually moderate (for standard systems)	Requires dynamic simulation model and interface to RL framework
Training / Startup Effort	Low (usually some run-time adaptations)	Low (e.g. few days of measurements for parameter estimates)	High – extensive training (on some powerful GPU machine) before deployment
Execution Effort	Very low	Generally considerable (repeatedly solve a new optimization problem)	Low (inference is cheap)
Runtime requirements	Simple hardware sufficient	Depend on complexity of the system (from Raspberry Pi to high-end PC)	Mid-level PC sufficient
Forecast Use (weather, load, etc.)	Possible, but requires more know-how, still rarely used	Usually integrated – Model Predictive Control	Possible, but increases training complexity (though feature engineering can help)
Adaptability to Structural Changes	Poor (all rules must be revised manually)	Moderate manual effort	Moderate manual effort (change of the environment), but retraining needed
Performance	Often reasonable, rarely optimal	High, near-optimal (model-dependent)	Potentially high (but no guarantee)
Robustness	Good, if relevant exceptions are covered by rules	Baseline low, but can be improved (safety margins, scenarios, ...)	Good, if training is truly exhaustive
Scalability for Increasing complexity	Poor for complex systems (large number of rules)	Good for setup, but the opt. problem becomes computationally heavy	Moderate manual effort (change of the environment), but retraining needed
Transparency	High (rules are explicit)	Medium (costs and constraints are known, but actions may be obscure)	Very low (black box)

Figure 19: General comparison of the aspects of rule-based, optimization-based and RL-based strategies for supervisory control, from (Lichtenegger, Muschick, Kaisermayer, Leitner, & Gölles, 2025)

From the point of view of the authors, at the moment, optimization-based control still should be regarded as the most promising single approach, though Reinforcement Learning has certainly gained a lot of attention and will most likely become more common for certain control tasks. An optimization-based approach still can and should, however, be complemented by elements from other ones as well:

- Integrate rule-based safety checks and fallback solutions,
- Use machine-learning methods for continuous updates of system parameters (online-learning),
- Possibly reduce effort for optimization by substituting some optimization steps by cheaper inference,
- Increase transparency by analysis methods adapted from explainable AI or enabled by complementary AI approaches.

5 Conclusions and Outlook

The control of solar district heating systems spans two interconnected domains: fast, robust component-level control and predictive, system-level supervisory control. Both are essential for transforming solar thermal energy from a fluctuating resource into a dependable contributor to district heating grids.

Component-level control, including collectors, hydraulic circuits, thermal storage and heat exchangers, is technically mature but increasingly influenced by new technologies such as tracking systems, hybrid collectors (PVT) and advanced storage concepts. While conventional control techniques remain effective, future improvements will likely stem from better integration with system-level forecasting and supervisory decisions.

System-level control faces greater complexity. It must coordinate multiple producers, consumers, storage units and network conditions, all while optimizing for cost, emissions, resilience, and equipment longevity. Among the evaluated strategies, optimization-based predictive control currently offers the best balance of transparency, reliability and performance. Rule-based methods remain valuable for safety, fallback logic and simple plant configurations, while AI-based methods – especially reinforcement learning – hold promise but are not yet ready for widespread deployment in critical infrastructure.

A central theme that emerges across all strategies is the need for hybrid approaches:

- combining rule-based safety mechanisms with predictive optimization,
- using machine learning to continuously update model parameters,
- replacing computationally heavy optimization steps with faster surrogate models,
- incorporating explainable AI to enhance trust and transparency.

Important issues for future investigations are:

- **Hybrid physics-based and AI-based methods:** Clear guidelines are needed to determine which approach best suits which task and how to combine both safely and efficiently.
- Improved **explainability** of advanced control strategies: As systems become more complex and decisions increasingly data-driven, transparency will be essential for operators, regulators and system planners.
- **Robustness** and **resilience:** Future supervisory controllers must better address forecast uncertainty, component failures and cybersecurity challenges in more interconnected energy systems.
- **Scalable** system architectures: As solar district heating evolves toward higher integration with electricity, mobility and industrial heat networks, control strategies must support multi-sector coupling and larger, more diverse portfolios of assets.

In summary, solar district heating has reached a level of technical maturity that makes it a key pillar of low-carbon heat supply. Further advances in predictive, optimization-based and hybrid control strategies will enable even higher solar shares, greater operational reliability, and more economic operation—ultimately strengthening the role of solar thermal energy in future energy systems.

6 References

- Achterberg, T., & Wunderling, R. (2013). Mixed Integer Programming: Analyzing 12 Years of Progress. In *Facets of Combinatorial Optimization* (Bd. 9783642381, S. 449-481). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Affolter, P., Eisenmann, W., Fechner, H., Rommel, M., Schaap, A., Sorensen, H., . . . Zondag, H. (2005). *PVT ROADMAP - A European guide for the development and market introduction of PV-Thermal technology*. Von https://www.researchgate.net/publication/240629568_A_roadmap_for_the_development_and_market_introduction_of_PVT_technology abgerufen
- Bava, F. (2017). Modeling of solar collector fields for solar heating plants in district heating systems. (D. o. Technical University of Denmark, Hrsg.) *B Y G D T U. Rapport*(R-365).
- Bava, F., & Furbo, S. (2018). Impact of different improvement measures on the thermal performance of a solar collector field for district heating. *144*, 816-825.
- Bergh, K. V., Bruninx, K., Delarue, E., & D'haeseleer, W. (2014). A mixed-integer linear formulation of the unit commitment problem. *University of Leuven (KU Leuven)-Energy Institute*.
- Bishop, C. M., & Bishop, H. (2023). *Deep Learning*. Springer Cham.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1), S. 1-122.
- Camacho, E. F., Berenguel, M., Alvarado, I., & Limon, D. (2021). Control of Solar Power Systems: a survey. *43*(5), S. 817-822.
- Carballo, J. A., Bonilla, J., Berenguel, M., Fernández-Reche, J., & García, G. (2019). *New approach for solar tracking systems based on computer vision, low cost*. doi:<https://doi.org/10.1016/j.renene.2018.08.101>
- Dénarié, A., Hay, S., Kemper, M., Komoszynska, M., Leusbrock, I., Madsen, P. T., . . . Styles, A. (2025). *Guidebook for the Integration of Renewable Energy Sources into Existing District Heating and Cooling Systems*. IEA DHC. Frankfurt am Main: AGFW-Project Company.
- Duffie, J. A., & Beckman, W. A. (2013). *Solar Engineering of Thermal Processes* (4 Ausg.). Hoboken, NJ, USA: John Wiley & Sons, Inc. doi:10.1002/9781118671603
- Frederiksen, S., & Werner, S. (2013). *District Heating and Cooling*. Lund: Studentlitteratur.
- Gölles, M., & Unterberger, V. (2021). *Control of large-scale solar thermal plants*. IEA SHC 55.B-D3.1.
- Gölles, M., Muschick, D., Unterberger, V., Leoni, P., Schmidt, R.-R., & Lennermo, G. (2021). *Control of DHC networks and Reduction of the operating temperatures in DH systems*. IEA SHC 55.A-D4.2.
- Gölles, M., Reiter, S., Brunner, T., Dourdoumas, N., & Obernberger, I. (2014). Model based control of a small-scale biomass boiler. *Control Engineering Practice*, 22(1), S. 94-102.
- Gölles, M., Unterberger, V., Kaisermayer, V., Nigitz, T., & Muschick, D. (2021). *Supervisory control of large-scale solar thermal systems*. IEA SHC 55.A-D4.1.
- Goyal, A., Staedter, M. A., & Garimella, S. (2019). A review of control methodologies for vapor compression and absorption heat pumps. *International Journal of Refrigeration*, 97, S. 1-20.
- Gurobi-Optimization. (8. 5 2019). Gurobi Optimizer Reference Manual. Von <https://www.gurobi.com> abgerufen
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.
- IBM. (8. 5 2019). IBM ILOG CPLEX User Manual. Von <https://www.cplex.com> abgerufen

- International Organization for Standardization. (2022). *Physical and mechanical properties of wood - Test methods for small clear wood specimens - Part 18: Vocabulary*. Von <https://www.iso.org/standard/79913.html> abgerufen
- Jensen, A., & Sifnaios, I. (2022). *Modeling, Validation, and Analysis of a Concentrating Solar Collector Field Integrated with a District Heating Network*. doi:<https://doi.org/10.3390/solar2020013>
- Jentsch, A. W. (2020). *DELFIN - Decentralized Feed-In: Prognose der Auswirkungen dezentraler Einbindung von Wärme aus erneuerbaren Energien und anderen Wärmeerzeugern in Fernwärmenetze*. Frankfurt am Main: AGFW - Der Energieeffizienzverband für Wärme, Kälte und KWK.
- Krüger, D., Fischer, S., Hirsch, T., & Labairu, I. (2020). Concentrating Solar Systems in Moderate Climates. *SolarPACES 2020*. Albuquerque.
- Kruse, R., Borgelt, C., Klawonn, F., Moewes, C., Steinbrecher, M., & Held, P. (2013). *Computational Intelligence A Methodological Introduction*. London: Springer. doi:10.1007/978-1-4471-5013-8
- Lemos, J. M., Neves-Silva, R., & Igreja, J. M. (2014). *Adaptive Control of Solar Energy Collector Systems*. Switzerland: Springer International Publishing.
- Lenermo, G. (2020). *"Solvärmeteknik - en handbok"*. Stockholm: AB Svensk Byggtjänst och författaren.
- Lenermo, G. (13. 10 2020). Dynamics and control of decentral ST feed-in. Presentation at the 8th and FINAL Expert Meeting. 13th to 15th of October 2020: IEA SHC Task 55 Meeting.
- Leoni, P., Geyer, R., & Schmidt, R.-R. (2020). Developing innovative business models for reducing return temperatures in district heating systems: approach and first results. *Energy 2020;195:116963*, <https://doi.org/10.1016/j.energy.2020.116963>.
- Lichtenegger, K., Berberich, M., Moser, M., Emberger, L., Beurskens, L., & Byström, J. (2025). *Efficient Solar District Heating Systems – Considering higher temperatures and digitalization measures (Technology Position Paper)*. IEA SHC.
- Lichtenegger, K., Leitner, A., Märzinger, T., Mair, C., Moser, A., Wöss, D., . . . Pröll, T. (2020). Decentralized heating grid operation: A comparison of centralized and agentbased optimization. *Sustainable Energy, Grids and Networks, 21*.
- Lichtenegger, K., Muschick, D., Kaisermayer, V., Leitner, A., & Göllles, M. (17. September 2025). Approaches to Supervisory Control: Rules, Optimization, Reinforcement Learning. *BEST Conference*. Graz: BEST.
- Lund, H., Werner, S., Wiltshire, R., Svendsen, S., Thorsen, J. E., Hvelplund, F., & Mathiesen, B. V. (2014). 4th Generation District Heating (4GDH): Integrating smart thermal grids into future sustainable energy systems. *Energy, 68*, S. 1-11.
- Ma, R., Li, J., Jin, F., Cao, Y., & Zhao, X. (2022). Fuzzy theory-based energy-saving control of solar thermal supplemental multi-energy heating system. *Energy Reports Volume 8, Supplement 6*, S. 636-646. doi:<https://doi.org/10.1016/j.egy.2022.03.065>
- Mallikarathne, T., Abeysignhe, H., Rathnayake, C., & Perera, M. (2023). *Design a Beach Cleaning Robot Based on AI and Node-RED Interface for Debris Sorting and Monitor the Parameters*. Sri Lanka. doi:10.1109/SLAAI-ICAI59257.2023.10365015
- Moser, A., Muschick, D., Göllles, M., Nageler, P., Schranzhofer, H., Mach, T., . . . Hofer, A. (2020). A MILP-based modular energy management system for urban multi-energy systems: Performance and sensitivity analysis. *Applied Energy, 261*, 114342. doi:10.1016/j.apenergy.2019.114342
- Murphy, K. P. (2022). *Probabilistic Machine Learning: An introduction*. MIT Press.
- Murphy, K. P. (2023). *Probabilistic Machine Learning: Advanced Topics*. MIT Press.
- Muschick, D., Zlabinger, S., Moser, A., Lichtenegger, K., & Göllles, M. (2022). A multi-layer model of stratified thermal storage for MILP-based energy management systems. *Applied Energy Volume 314*. doi:10.1016/j.apenergy.2022.118890

- Nigitz, T., & Göllles, M. (2019). A generally applicable, simple and adaptive forecasting method for the short-term heat load of consumers. *Applied Energy*, 241, S. 73–81.
- Nsengiyumva, W., Chen, S. G., Hu, L., & Chen, X. (2020). Recent advancements and challenges in Solar Tracking Systems (STS): A review. *Renewable and Sustainable Energy Reviews*, 81(1).
- Schörghuber, C., Göllles, M., Reichhartinger, M., & Horn, M. (2020). Control of biomass grate boilers using internal model control. *Control Engineering Practice*, 96.
- Staudt, S. (2023). *Model-based control of absorption heat pumping devices*. PhD thesis, Graz University of Technology, Graz.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2. Ausg.). Cambridge, Massachusetts; London, England: MIT Press.
- Tamm, S., & Berberich, M. (2024). *Solar Collector Technologies for District Heating*. IEA SHC Task 68. IEA SHC. doi:10.18777/ieashc-task68-2024-0002
- Unterberger, V. (2021). *Modelling and control of large-scale solar thermal systems*.
- Unterberger, V., Innerhofer, P., Lichtenegger, K., & Göllles, M. (2018). *SPC - Solar Predictive Control - Publishable final report (in German)*. Tech. rep., Österreichische Forschungsförderungsgesellschaft mbH (FFG) - Klima- und Energiefonds - Energieforschungsprogramm. Von <https://www.energieforschung.at/projekte/903/solar-predictive-control> abgerufen
- Unterberger, V., Lichtenegger, K., & Göllles, M. (2024). Predictive Rule-Based Control Strategy for Optimizing the Operation of Solar District Heating Plants. doi:10.18086/eurosun.2024.03.03
- Unterberger, V., Muschick, D., & Göllles, M. (2017). Model-based control strategies for an efficient integration of solar thermal plants into district heating grids. *ISES Solar World Congress*. doi:doi:10.18086/swc.2017.06.15
- Unterberger, V., Muschick, D., Loidl, A., Poms, U., Göllles, M., & Horn, M. (2020). Model-based control of hydraulic heat distribution systems - Theory and application. *Control Engineering Practice*, 101.
- Unterberger, V., Muschick, D., Loidl, A., Poms, U., Göllles, M., & Horn, M. (2020). Model-based control of hydraulic heat distribution systems - Theory and application. *Control Engineering Practice*, 101.
- Verrilli, F., Srinivasan, S., Gambino, G., Canelli, M., Himanka, M., Del Vecchio, C., . . . Glielmo, L. (2017). Model Predictive Control-Based Optimal Operations of District Heating System With Thermal Energy Storage and Flexible Loads. *IEEE Transactions on Automation Science and Engineering*, 14, 547-557. doi:10.1109/tase.2016.2618948
- Weiss, W., & Spörk-Dür, M. (2024). *Solar Heat Worldwide - Global Market Development and Trends 2023*. resreport, IEA SHC Programme. doi:10.18777/ieashc-shww-2024-0001