

Validation of a Dynamic Process Model of a Thermochemical System for Solar Fuel Production

Falko Schneider¹  <https://orcid.org/0000-0002-9830-6707>, Constantin Peters¹  <https://orcid.org/0000-2222-333-4444>,
Cristiano José Teixeira Boura¹  <https://orcid.org/0000-0003-1854-1019>, Moritz Leibauer², Lukas Geiss-
buehler² and Ulf Herrmann¹  <https://orcid.org/0000-0002-6938-0860>

¹ Solar-Institut Jülich of FH Aachen University of Applied Sciences, Germany

² Synhelion AG, Dufourstrasse 101, 8008 Zurich, Switzerland

*Correspondence: Falko Schneider, f.schneider@sij.fh-aachen.de

Abstract. A dynamic process model of a solar-driven fuel production pilot plant has previously been developed to aid in plant operation and control. This work presents the performance of the model's key components – a high-temperature thermal energy storage (TES), a solar-absorbing gas receiver (SAGR), and a reforming reactor – on operational data from Synhelion's DAWN pilot plant. The TES model plausibly reproduces charge–discharge transients. For the SAGR, a data-driven surrogate solar field model was created to provide the unmeasured solar input. The coupled surrogate–receiver simulation captures the receiver's thermal response well, achieving outlet gas temperatures near 1000 °C in good agreement with experiments. A slight overestimation of ~30–40 K in receiver outlet temperature is observed, consistent with a minor bias in the surrogate model. In contrast, the reforming reactor model, which assumes thermodynamic equilibrium at the reactor outlet, could not be fully validated. Discrepancies in syngas composition (e.g. residual CH₄ and lower CO₂ conversion than predicted) indicate that equilibrium assumptions are not valid under partial-load conditions. The reactor model overestimates reaction extent and heat consumption, suggesting that a kinetic modeling approach or adjusted equilibrium model is needed for accuracy. Overall, the validated components form a promising digital twin of the solar fuel process, but refining the chemical reactor model is essential for reliable full-process simulation.

Keywords: Dynamic Simulation, Model Validation, Modelica, Solar Fuels, Thermochemical Equilibrium, CSP, Solar-Absorbing Gas Receiver

1. Introduction

Sustainable fuels produced via concentrated solar power (CSP) are a promising pathway to decarbonize the transportation sector. In a solar fuel process, solar thermal energy drives an endothermic chemical reaction (in this case, reforming methane, carbon dioxide and steam) to generate synthesis gas (syngas), which can be converted to liquid hydrocarbons (e.g. kerosene) via Fischer–Tropsch synthesis. Such solar thermochemical plants aim to provide drop-in renewable fuels for hard-to-electrify sectors like aviation, leveraging the high energy density of liquid fuels. A pilot plant implementing this concept – named DAWN – has been built and commissioned in Jülich (Germany) by Synhelion, demonstrating the entire production chain under real solar operation, including the usage of a thermal energy storage.

Controlling such a solar-driven thermochemical system is challenging due to the highly transient and volatile solar input, coupled with the complex kinetics of the reforming reactions.

Dynamic simulations are a valuable tool for designing control strategies and understanding transient behavior in CSP plants. In prior work, a detailed dynamic process model of the solar fuel pilot plant using the Modelica has been developed [1]. This model incorporates custom component models for the thermal energy storage (TES), the solar-absorbing gas receiver (SAGR), and the reforming reactor, along with standard fluid network components for pipes and heat exchangers. However, before such a model can reliably assist in plant operation, it must be validated against experimental data.

This paper focuses on component-level validation of the dynamic models using operational data from DAWN's initial test campaigns. The focus is on the three main components: TES, SAGR and reforming reactor. Simulations are run with available operational data as boundaries and its results compared to the actual measurements. The following sections describe the plant data and model setup, present the validation results for each component and discuss implications for model improvements.

2. Performance on Operational Data

Operational data from the DAWN pilot plant was provided by Synhelion to validate the models. The data include measured temperatures, flow rates, pressures and other sensor readings at various locations in the plant during specific test runs. For individual validation, we isolated tests where a single component's behavior could be observed with minimal interference from other subsystems (e.g. charging/discharging the TES alone, or operating the receiver under certain conditions). The simulation models were configured to mirror the test conditions (initial states and boundary conditions), and the same inputs (such as inlet HTF flow rate, DNI or heater inputs, etc.) recorded in the experiments were fed into the simulation. Key outputs (temperatures, heat transfer rates, etc.) from the simulation were then compared to the measured data. In all cases, the measured signals were pre-processed to remove high-frequency noise – for example, temperature measurements were slightly smoothed to filter out short spikes from sensor noise. The following subsections present validation results for the TES, the SAGR and the reforming reactor.

2.1 Thermal Energy Storage (TES)

The pilot plant's TES is a high-temperature solid thermal storage consisting of ceramic refractory bricks with internal air channels (combstone geometry). Hot heat transfer fluid (HTF), such as air or steam, flows through the bricks to charge the storage, and later cold HTF is heated by the bricks to discharge it. In the model, the TES is spatially discretized into a grid of smaller cuboid cells in order to resolve the axial and radial temperature gradients within the packed bed and the surrounding insulation. Each cell represents a lumped mass of storage material with a uniform temperature state. The cells are connected in a 3D lattice to approximate the continuum temperature field. For simulations, a discretization level was chosen such that further refinement did not significantly change the results, indicating the solution had converged spatially.

The chosen operational scenario to validate the model contains 8 charging and discharging cycles over a period of 10 days. The temperatures inside the storage are measured at various locations with two sensors per spot as redundancy. For validation purposes, the measured temperatures along the center vertical line are compared with the simulated HTF and storage material temperatures of the corresponding cell in the model (cf. Figure 1). Usually, the HTF and material temperatures differ by less than 10 K. The model was initialized with a slight vertical temperature profile with linear interpolation in-between sensors to mirror the initial test conditions. Due to confidentiality restrictions, the positions as well as the temperatures are normalized.




Figure 1. Comparison of simulated temperatures along the vertical center line of the TES with measurements.

While the model shows a good fit for the upper layers of the storage, the deviations between operational and simulation data increase towards the bottom (cf. Table 1). This can be caused by various factors, both inside the model as well as the plant itself. These material properties are based on the bulk material and are implemented as temperature-dependent polynomials. They do not reflect possible changes to the density and heat capacity due to the manufacturing process of the storage bricks that might change the temperature-dependence. Additionally, during some test campaigns, the measurement of the mass flow rate proved to be unreliable and might be an error source. A lower mass flow rate results in a lower charging/discharging rate and vice versa. In the end, the simulation shows a faster charging and discharging rate than the available measurements.

Table 1. Root mean squared errors between simulated temperatures (Mat & HTF) and measured temperatures (M1 & M2), normalized to the overall TES temperature swing.

Height	M1-HTF	M1-Mat	M2-HTF	M2-Mat
96%	1,6%	1,3%	1,4%	1,3%
81%	3,0%	2,9%	3,0%	2,7%
65%	4,0%	3,8%	4,0%	3,7%
50%	6,6%	6,4%	4,9%	4,4%
35%	9,4%	8,8%	7,1%	6,4%
19%	11,1%	10,5%	9,9%	9,3%
4%	8,8%	8,4%	8,9%	8,4%

The model was used with parameters from prior design estimates, and no data-based optimization was performed for the validation. Even without calibration, the results imply that the model properly captures the energy balance and heat transfer dynamics of the storage, giving confidence that it can predict storage behavior under different charging/discharging schedules, although minimal fitting might be necessary.

2.2 Solar-Absorbing Gas Receiver (SAGR)

The boundary of the process model is at the receiver aperture, where the incoming concentrated solar irradiation is defined as a model input. Then, depending on the HTF inlet conditions (mass flow, pressure and temperature), the HTF outlet conditions are determined. This presents a challenge for the validation since the solar flux onto the receiver is not measured during operation. Indirect indicators of the solar input are included in the operational data: direct normal irradiance (DNI) measurements, number of focusing heliostats and the sun's position (azimuth and zenith). Thus, to validate the receiver model, the scope had to be expanded to include a model of the solar field.

2.2.1 Solar Field Surrogate Model

The efficiency of a heliostat field depends on many factors, e.g. irradiance, tracking errors, mirror soiling, atmospheric attenuation, shading or manufacturing defects such as misaligned mirrors or optical abnormalities. These effects are difficult to capture in detail with a first-principles Modelica model of the field. Therefore, a simple data-driven surrogate model of the solar field was developed and coupled with the receiver model to provide an estimated solar flux input. This surrogate is used to check the plausibility of the receiver model by ensuring the combined field–receiver model can reproduce the observed receiver behavior.

First, some relevant reference scenarios are selected from the plant data. In order to mitigate the effects of the heliostats conditions on the efficiency, only operational data from a timeframe of about two dry weeks are chosen. This dataset is further filtered by only including receiver operation during clear-sky conditions. Data from three days are used to train the data-

driven model with the available model inputs while a fourth day is used to run the validation simulation with the combination of the receiver model and the solar field surrogate model.

The necessary training target (the effective solar power incident on the receiver) was obtained by running the receiver model with fixed HTF inlet/outlet conditions taken from the training days. Essentially, it was simulated what solar input would be required for the receiver model to produce the measured outlet temperatures during training periods. The resulting inferred irradiation profiles appeared plausible based on operator experience and also comparable to estimates from higher-fidelity tools (CFD and ray-tracing models).

Based on the training simulation results, the solar field is modeled as a smooth, low-bias function of these four predictors: Azimuth, zenith, the number of active heliostats and DNI. Each predictor is standardized (zero-mean, unit-variance) and we expand the feature space with all monomials up to total degree 4 (including interactions). The resulting polynomial basis (70 terms for 4 inputs) is fitted with ridge regression (i.e. a linear regression with L2 regularization), which shrinks correlated coefficients and mitigates overfitting while retaining the nonlinear interaction structure. The final surrogate is therefore a 4th degree polynomial with standardized inputs, trained on multiple days of operation to improve robustness. Table 2 summarizes the quality of the surrogate model regarding the training and the independent validation data. For these KPIs, data points with DNI below 200 W/m² or receiver irradiation below 1% of the nominal solar field power are disregarded.

Table 2. Coefficient of variation of the root mean squared error (CVRMSE) and normalized mean bias error (NMBE) of the solar field surrogate model on the training data and for the validation scenario.

Dataset	CVRMSE	NMBE
Training	6.29%	0.03%
Validation (first 3h)	9.45%	3.25%

These results indicate that the 4th degree polynomial with ridge regularization captures the main nonlinear dependence of field output on the inputs while generalizing reasonably well to unseen periods. The small positive NMBE suggests a minor systematic overestimation. Overall, the error levels are low enough so the approximation can be used as a solar-field surrogate to at least verify the receiver model plausibility.

Since the surrogate model was trained in conjunction with the receiver model (using the receiver to infer solar input), any constant biases or certain inaccuracies in the receiver model could be absorbed by the surrogate (for instance, an error in assumed receiver aperture reflectivity would reflect in the inferred solar input). However, since the surrogate is a static model trained on data from different receiver operating points (including various ramp rates and target temperatures), the dynamic response of the receiver in the coupled simulation is still governed by the physical Modelica model of the receiver.

2.2.2 Receiver Simulation Results

In order to validate the receiver dynamics, the validation scenario is simulated with the combined receiver and solar field models. The simulation inputs are the HTF inlet conditions, DNI, number of active heliostats, azimuth and zenith recorded in the operational data. The HTF outlet temperature is simulated and compared to the measurement. This temperature response provides a KPI to determine the validity of the modelling approach (cf. Figure 2).




Figure 2. Comparison of measured and simulated HTF outlet temperatures of the receiver during start-up and operation (clear-sky until ~11:00, followed by cloud transients).

The model's outlet temperature lags behind the measurement during the initial heat-up, which indicates a slower receiver heat-up. This can be caused by a simplified initialization with a uniform cavity surface temperature and a steady-state temperature distribution inside the insulation. The initial surface temperature is based on a measurement a few centimeter below the actual surface. There are no temperature sensors directly on the surface due to the harsh conditions during irradiation. These effects are exacerbated by the geometrical simplification of the complex receiver geometry to a cylindrical shape. Additionally, some discrepancies can stem from the surrogate model, which does not consider which individual heliostat is aiming at the receiver. At lower receiver temperatures (early in the start-up), less heliostats are in operation to maintain a safe ramp-up to target temperature. Individual mirror characteristics and tracking errors have a larger impact on the receiver irradiation at these lower solar field utilizations. They average out at higher utilizations, where the surrogate model's aggregate prediction is more accurate.

This is evident as the HTF outlet temperature starts to rise after about 20 minutes when more heliostats are aiming at the receiver. After this, the simulated data is more closely aligned to the operational data and reaches a quasi-steady outlet temperature slightly above 1000 °C at roughly the same time. The model does overestimate the temperature by approx. 30-40 K, which is in line with the previously identified systematic overestimation of the surrogate model. Higher receiver outlet temperatures are possible and have been achieved at DAWN. Around 3 hours into the test, clouds started to move in and reduced the receiver irradiation. An effect, which is not contained in the training data and model inputs, that leads to diverging temperatures since the surrogate model is operating out of its boundaries.

Considering only the clear-sky portion of the test (avoiding the unmodeled cloud impact), the root mean square error (RMSE) between simulated and measured receiver outlet temperature is 48.1 K. Including the cloud disturbance period, the overall RMSE rises to 59.5 K. Given an operating temperature on the order of 1000 °C, the clear-sky error corresponds to about 4.8% and the overall error ~6%. These results are deemed acceptable for an initial validation, and qualitatively the model's response and steady-state are very plausible. The modelling approach thus appears suitable to capture the characteristic thermal inertia and efficiency of the SAGR under typical conditions.

2.3 Reforming Reactor

The reforming reactor model was initially based on the simplifying assumption that the product gas composition reaches thermochemical equilibrium (TCE) at the reactor outlet temperature. During operation, the dry gas composition is analysed downstream of the reactor. This data is compared to the thermochemical equilibrium at the measured temperatures over a span of almost two weeks. Due to confidentiality restrictions, Figure 3 only shows the relative deviation of the measured composition from the respective thermochemical equilibrium state, rather than the absolute concentrations.




Figure 3. Relative deviations of measured syngas composition from the thermodynamic equilibrium (TCE) composition at the reactor outlet. Positive values indicate the simulated species fraction is higher than the measured one and negative means lower. (Data for H_2O is excluded since measurements are post water-removal.)

There is a fairly good agreement between simulation and measurement regarding H_2 production. But according to the TCE, the amount of CO in the product stream should be higher and CO_2 should be lower. Additionally, complete CH_4 conversion should occur according to equilibrium, but there are still traces present in the reformer output. So, the carbon conversion in the reformer is lower than expected and the initial modelling assumption of thermochemical equilibrium conditions at the reformer outlet is inaccurate (cf. Table 3). It is important to note though, that the reformer operated mostly in partial load during the investigated validation scenario.

Table 3. Relative and absolute root mean square errors between measurements and TCE

Species	CH_4	CO	CO_2	H_2
abs. RMSE [mol-%]	2.1	4.2	4.3	1.5
rel. RMSE [%]	100	18.6	16.3	3.1

During the validation scenarios, the reforming reactions may not reach full equilibrium in the available residence time or the catalyst effectiveness might be lower than assumed. The model in its current form appears to overestimate the reaction extent and thereby

heat consumption from the HTF, leading to a different thermal profile than measured, on the chemical as well as the HTF side. Thus, while trends are qualitatively consistent, the equilibrium assumption is quantitatively inaccurate and a direct validation was not possible.

3. Conclusion & Outlook

Previous work presented modelling approaches to create a comprehensive dynamic process model for a solar-driven fuel production plant. It focused on three main components: TES, SAGR and reforming reactor. This work offers a performance analysis of these approaches using real operational data gathered at plant DAWN by Synhelion in Jülich, Germany. The models were simulated based on the previously published approaches without any fitting to the operational data.

The TES model demonstrated good predictive capability for transient charge/discharge behavior, with errors under 3% of the temperature swing, confirming that the discretization and heat transfer sub-models are sound. Because the receiver irradiation is not measured online, a surrogate model had to be developed in order to validate the solar absorbing gas receiver (SAGR) model. Then, it was shown to reproduce the receiver's dynamic response and achieve high outlet temperatures in line with measured values. The reforming reactor model however could not be conclusively validated due to the current limitations in chemical modeling. The assumption of thermodynamic equilibrium proved inaccurate and led to an overestimation of reaction extent and heat usage, which was inconsistent with the partial load scenarios observed during pilot plant operation. The equilibrium model likely needs refinement or replacement with a kinetic model to match the real reactor.

As more operational data becomes available the models can be further refined, e.g. modifying the temperature dependence of material properties. Especially the chemical modelling will profit from the growing amount of data, which could enable the determination of empirical kinetic reaction rates or by adjusting the equilibrium approach (e.g. using effectiveness factors or approach-to-equilibrium parameters) to more accurately predict the reactors behaviour in full and partial load. Additionally, lab scale experiments are being carried out on various catalysts which will aid in that effort as well. Due to the complexity of inputs determining the efficiency of a solar field, machine learning systems are being investigated to generate more accurate results than the basic surrogate model used for validating the receiver.

Without an accurate reformer model, we cannot yet fully validate the overall process model, since errors in the reformer propagate to the rest of the system. The purpose of these validations is to establish a digital twin of the solar fuel plant, which can mirror the plant behaviour in real-time and serve as powerful tools for operational support. With a reliable digital twin, one can implement model predictive control, perform what-if analyses for extreme conditions or test new control algorithms virtually before deploying them. The validated models can also be used to design scaled-up systems, as they capture the essential physics needed to predict performance under different sizes and throughput. Initial applications of the validated models, such as virtual testing of control algorithms, have already been carried out. The next step is to implement real-time simulations of validated models during plant operation, effectively enabling an operational digital twin.

Data availability statement

The operational data and specifications are subject to third-party IP and not publically available. Some data had to be de-dimensioned to protect sensitive data.

Author contributions

Conceptualization¹, Data curation¹, Formal Analysis¹, Funding acquisition^{1,3,4}, Investigation^{1,2}, Methodology^{1,2}, Project administration^{1,3}, Software^{1,2}, Supervision^{3,4}, Validation¹, Visualization¹, Writing – original draft¹, Writing – review & editing^{2,3,4}

Competing interests

The authors declare that they have no competing interests.

Funding

This work was carried out with financial support from the German government through its 7th Energy Research Programme (FKZ: 03EE5085C, SolarFuels).

Acknowledgement

The authors gratefully acknowledge the financial support from the German Federal Ministry for Economic Affairs and Energy (BMWE). We also thank our colleagues at Synhelion for their continued support and providing the operational data.

References

- [1] F. Schneider, C. Schwager, C. J. Teixeira Boura, N. Oellers, A. Villegas, and U. Herrmann, “Dynamic Modeling of a Thermochemical System for Solar Fuel Production Based on an Open-Source Framework”, *SolarPACES Conf Proc*, vol. 2, Aug. 2024, doi: <https://doi.org/10.52825/solarpaces.v2i.908>