Journal of Power Sources 685 (2026) 240440

Contents lists available at ScienceDirect

Journal of Power Sources

journal homepage: www.elsevier.com/locate/jpowsour

Model predictive control for spatial–temporal temperature gradient constrained dynamic operation of a marine SOFC system

Matthis H. de Lange [a,*], Pablo Segovia [b,c], Rudy R. Negenborn [a], Lindert van Biert [a]

[a] Maritime & Transport Technology, Delft University of Technology, Mekelweg 5, Delft, 2628 CD, The Netherlands
[b] Department of Automatic Control, Universitat Politècnica de Catalunya - BarcelonaTECH, Av. Diagonal 647, Barcelona, 08028, Spain
[c] Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Carrer de Llorens i Artigas, 4, Barcelona, 08028, Spain

HIGHLIGHTS

- SOFCs offer a renewable energy solution for onboard maritime power generation.
- Dynamic load induce thermal stress, increasing degradation and reducing SOFC lifetime.
- Proposed load-tracking MPC using spatial–temporal temperature constraints and cost
- A spatially resolved 1D SOFC stack model enables thermal-stress aware control.
- The strategies reduce stress from temperature gradients and improve load tracking.

ARTICLE INFO

Keywords:
Solid oxide fuel cells
Model predictive control
Thermal stress
Dynamic operation
Maritime energy system
Alternative fuels

ABSTRACT

The use of solid oxide fuel cells (SOFC) offers an alternative energy-conversion technology for the maritime sector, supporting the transition to renewable fuels. However, operating SOFCs for onboard power generation requires them to accommodate dynamic load changes, which introduces thermal stress, accelerates degradation, and reduces their operational lifetime. This work introduces a set of load-tracking model predictive control (MPC) strategies that reduce thermal stress by introducing spatial temperature gradient constraints (STGC), temporal temperature gradient constraints (TTGC) and temporal temperature gradient cost (TTGQ) components. The development of a spatially resolved one-dimensional prediction model for the SOFC stack is essential for incorporating these components into MPC strategies. The strategies are evaluated via simulations across multiple scenarios using key performance indicators (KPIs) for thermal stress, load-tracking performance and electrical efficiency, and benchmarked against a baseline MPC and a current-ramping-limit (CRL) approach. The results show that the STGC effectively reduces and constrains the spatial temperature gradient while maximising electrical efficiency. Furthermore, the TTGC and TTGQ strategies improve dynamic load-tracking response while resulting in lower temporal temperature gradients than a CRL.

1. Introduction

In 2018, the maritime sector accounted for 2.89% of global Greenhouse Gas (GHG) emissions [1], making it a substantial contributor. To reduce emissions, the sector is exploring alternative fuels and powertrain concepts. This remains challenging due to the high variability in operational profiles, which significantly affects power system design [2]. Electrification through batteries is mainly interesting for short voyages. Hydrogen can play a significant role for the bulk of voyages but lacks sufficient energy density for deep-sea voyages. As a result, fuels such as *Liquefied Natural Gas* (LNG), methanol and ammonia are expected to play a key role [3,4]. These renewable fuels enable the adoption of alternative energy-conversion technologies. In particular, the *Solid Oxide Fuel Cell* (SOFC) has gained attention due to its ability to internally reform these alternative fuels, resulting in high electrical efficiencies of 50%–65% and minimal system complexity [5]. This makes the SOFC a promising candidate for long-voyage shipping, where the volumetric energy density of batteries and pure hydrogen is insufficient to meet mission requirements. However, the SOFC technology is not fully mature and is associated with high investment costs, poor dynamic response to load changes and a limited lifetime, which prevent its large-scale adoption [6]. The high investment cost is expected to decrease as

* Corresponding author.
E-mail address: m.h.delange@tudelft.nl (M.H. de Lange).

https://doi.org/10.1016/j.jpowsour.2026.240440
Received 19 March 2026; Received in revised form 30 April 2026; Accepted 16 May 2026
Available online 27 May 2026

production scales up [7]. The limited dynamic response to load changes and the limited lifetime require a more sophisticated solution.

1.1. Operational challenges

The first commercial SOFCs were originally intended for static operation [8], which is reflected in their sensitivity to dynamic operation. Their operation requires tight management of gas flows and power demand to avoid fuel or air starvation, under-voltage, overheating and large internal temperature variations within the cell, as these factors can damage the *Positive-electrode–Electrolyte–Negative-electrode* (PEN) microstructure [9]. While many of these risks are mitigated through appropriate control strategies, temperature dynamics remain difficult to manage because they are spatially distributed and locally diverge substantially from the average stack temperature. Furthermore, internal temperature measurements are impractical. Local temperature variations can be reduced by gradually transitioning between operating points, which inherently limits the SOFC's responsiveness to dynamic load changes [10].

The sensitivity to dynamic operation originates from thermo-mechanical stresses within the PEN structure, whose bonded layers have mismatched *Thermal Expansion Coefficients* (TECs). During temperature changes, these mismatches restrict independent expansion and generate internal stresses that can cause plastic deformation, cracking, or delamination, thereby reducing membrane conductivity [11]. Although stresses are inherent at the high operating temperatures of SOFCs and contribute to long-term degradation, they are intensified by local hotspots and high spatial or temporal temperature gradients, which can lead to severe degradation or even complete conductivity loss [12]. Transient operation alters the temperature distribution, increasing the risk of local hotspots and steep gradients. Careful load-change management is therefore essential to mitigate degradation.

Local hotspots result from the cell overheating, where the maximum temperature in the PEN locally exceeds the upper-temperature limit due to insufficient cooling. The SOFC is typically operated with the maximum PEN temperature near its limit. Therefore, the limit is easily exceeded if power and airflow are not carefully adjusted together. A local hotspot results in a high spatial temperature gradient and causes the PEN material to expand above its design point, leading to excessive material stress [11]. Manufacturers specify limits for the stack temperature within which the stack is safely operated, but only for some measurable temperatures at the edge of the stack.

In the literature, the spatial temperature gradient in the PEN structure is most often reported in relation to thermal stress and degradation. The spatial temperature variation exists due to the non-uniform distribution between endothermal and exothermal reactions in the stack [13]. These result in different expansion and contraction within a material, further increasing internal stresses. The maximum spatial temperature gradient an SOFC can withstand varies per cell design. Influencing factors include operating temperature, material, manufacturing and cell configuration. Zeng et al. [11] report various maximum operating gradients of stacks ranging from 10 [K/cm] to 55 [K/cm].

SOFCs are designed to operate at a constant temperature, even under load variations. Heating up and cooling down from this operation temperature is known as thermal cycling. Often, the SOFC can withstand only a limited number of thermal cycles at a slow overall temperature rate-of-change. Quickly changing the temperature can thermally shock the PEN material due to dynamic expansion and contraction, leading to plastic deformation and degradation of the microstructure [12]. However, during dynamic operation, the spatial temperature distribution changes depending on the heat generated in the stack and on how effectively the balance of plant and control strategy provide heating and cooling. The rate at which the local temperature changes is known as the temporal temperature gradient. Dynamic load changes do not generate stress levels that cause immediate damage, but the PEN structure experiences fatigue over time due to long-term dynamic load stresses [13].

1.2. Control strategies

The amount of thermal stress the SOFC experiences depends on the manner of operation. The operational envelope of the SOFC comprises a set of steady-state operating points spanning the SOFC system's power range. The choice of these points and the way the system is dynamically operated from one point to another determine the occurrence of local hotspots and the magnitude of spatial/temporal temperature gradients experienced by the SOFC [9]. Therefore, choosing an appropriate control strategy is crucial for safely operating the SOFC.

Initial research aimed at steady-state control of the SOFC at a predefined setpoint for grid applications [8]. The controllers stabilise the conversion process, ensuring the operation remains within a small range around the setpoint. The temperature, fuel and air utilisation remain relatively constant, so additional safety measures are not needed for these limits. When designing for dynamic load tracking, the control strategy is often split into power control [14,15] and temperature management [16]. *Proportional–Integral–Derivative* (PID) control is commonly investigated in this research. Safety limits are only taken into account through proper controller tuning and input saturation [17]. Some of these thermal management strategies include consideration of spatial temperature gradients. For instance, Fardadi et al. [17] designed an H_∞-controller to minimise variations of the temperature distribution from a nominal distribution.

As transient performance demands grow, the operating constraints become increasingly critical, and power and thermal management must be addressed together. Sequential loop-shaping methods are employed in [18,19] based on PID control, and MPC methods are discussed in [20–22]. The advantage is that MPC naturally addresses the coupling in power and thermal management. Furthermore, MPC can account for hard constraints, thereby addressing the tight operating space and critical constraints of the SOFC system.

Despite its significance, thermal stress-aware operation remains inadequately addressed. Most existing works do not address thermal stress and rely on lumped-parameter models for control design, which fail to capture local thermal behaviour and associated stress. In practice, strategies rely on a *Current-Ramping Limit* (CRL) to prevent excessive thermal stress, which leads to the well-known issue of poor dynamic load response [6]. In an earlier work, a temperature-rate constraint based on a lumped model is used. While this addresses the peak local temperature and reduces thermal stress from rapid temperature changes, it does not directly act on the temporal or spatial temperature gradients [23]. Jie et al. [24] and Xia et al. [25] use spatially distributed models and constrain both the local peak temperature and the spatial temperature gradient, thereby preventing the resulting thermal stress from exceeding critical limits. However, this approach overlooks long-term fatigue effects. While brief stress exceedances may not be immediately critical, sustained or repeated stresses still lead to mechanical damage over time [13]. Additionally, although earlier work has suggested that lowering spatial temperature gradients can reduce electrical efficiency due to higher cooling-air demand [26], any specific trade-offs have not yet been examined in detail.

Finally, the temporal temperature gradient has received little attention, even though it strongly affects the transient power response. Safe limits on local temperature and its gradients are rarely defined, largely because they are difficult to measure, which makes it challenging to design thermal stress-aware control strategies. Nevertheless, with spatially distributed models and knowledge of the system's operating envelope and response, it is possible to estimate a maximum allowable temperature and gradient [11].

1.3. Contribution

This research develops a thermal stress-aware MPC strategy that explicitly limits thermal stress by addressing local hotspots, spatial and temporal temperature gradients. Specifically, this work proposes to:

- Design an MPC framework based on a spatially distributed SOFC prediction model, allowing the local temperature and both temporal and spatial temperature gradients to be represented within the constraints and cost function.
- Investigate how spatial and temporal temperature gradients develop across the SOFC stack and identify which parts of the PEN structure are most exposed to high thermal stresses.
- Assess to what extent the proposed MPC strategy reduces thermal stress compared to the conventional CRL.
- Evaluate the broader impact of thermal stress-aware operation on the dynamic load-response capability and overall efficiency of the SOFC system.

1.4. Outline

The methods are described in Sections 2–6. The time-dependent thermal stress *Key Performance Indicators* (KPIs) based on spatial and temporal temperature gradients are established in Section 2. Section 3 presents the SOFC system and describes the modelling of its various components. A time-scale–approximated prediction model is developed in Section 4, followed by the baseline MPC formulation. The various approaches for incorporating thermal stress considerations into the baseline MPC are presented in Section 5. Section 6 provides the implementation details of both the MPC strategy and the associated models. In Section 7, the proposed MPC strategy is evaluated through simulations and compared against a CRL-MPC approach with the results analysed and discussed. Finally, Section 8 concludes this paper and offers recommendations for future work.

2. Thermal stress performance indicators

Prior to discussing the models and proposed control strategies, KPIs are defined to evaluate the performance of the control strategies with respect to thermal stress, load-following capability, and efficiency. In prior work, the maximum gradients occurring over a simulation are constrained in the control strategy [24] and analysed as a performance indicator [23]. This approach is valid if the primary focus is a sudden failure of the SOFC due to exceeding critical limits. Although these critical limits are definitely present, they are not the only sources of failure in the SOFC system. In the long term, fatigue and degradation also play a crucial role. These effects accumulate over time and are location-dependent. Ideally, one would be able to quantify mechanical stress, fatigue, probability of failure or degradation. However, a 1D spatial model is insufficient for a comprehensive understanding of mechanical stress, as the 3D geometry typically plays a key role, requiring in-depth finite element modelling with thermal–mechanical coupled dynamics [27]. These types of models are computationally intensive, making them inherently unsuitable for control purposes. As an approximation, this work proposes integrating the temperature gradients quadratically over time as a performance indicator of the thermal stress at a specific location of the PEN structure. A quadratic function is chosen to emphasise the impact of higher gradients. The temporal temperature gradient is the time derivative of the temperature in the active PEN structure ($dT_{\mathrm{PEN}}(t,l)/dt$). The spatial temperature gradient is defined as the temperature change over the spatial dimension ($dT_{\mathrm{PEN}}(t,l)/dl$). Here $T_{\mathrm{PEN}}(t,l)$ is the PEN temperature in the active area, t and l are the time and spatial dependency, respectively. The KPIs for the thermal stress due to the accumulated temporal temperature gradient ($\nabla_{\mathrm{t}}T(l)$) and the accumulated spatial temperature gradient ($\nabla_{\mathrm{l}}T(l)$) are formulated as

$$\nabla_{\mathrm{t}}T(l) := \frac{1}{t_{\mathrm{sim}}\nabla_{\mathrm{t}}\overline{T}}\int_{t=0}^{t_{\mathrm{sim}}}\left(\frac{dT_{\mathrm{PEN}}(t,l)}{dt}\right)^2 dt, \tag{1}$$

$$\nabla_{\mathrm{l}}T(l) := \frac{1}{t_{\mathrm{sim}}\nabla_{\mathrm{l}}\overline{T}}\int_{t=0}^{t_{\mathrm{sim}}}\left(\frac{dT_{\mathrm{PEN}}(t,l)}{dl}\right)^2 dt, \tag{2}$$

with t_{sim} as the total simulation time. These stress KPIs are evaluated along the active PEN area, which extends over a length L and the KPIs are normalised with respect to the maximum KPIs $\nabla_{\mathrm{t}}\overline{T}$ and $\nabla_{\mathrm{l}}\overline{T}$.

The main objective of the SOFC is to deliver the requested power ($r_{\mathrm{P}}(t)$), which is captured as

$$\Pi := 1 - \frac{\int_{t=0}^{t_{\mathrm{sim}}}\left|r_{\mathrm{P}}(t) - P_{\mathrm{EPC}}(t)\right|}{\sigma_{\mathrm{p}}}dt, \tag{3}$$

where $P_{\mathrm{EPC}}(t)$ is the electrical power that comes out of the power converter. The tracking error norm is normalised by $\sigma_{\mathrm{p}} = \int_{t=0}^{t_{\mathrm{sim}}}\left|r_{\mathrm{P}}(t) - \bar{r}_{\mathrm{p}}\right|dt$, such the median of the power reference ($\bar{r}_{\mathrm{p}}$) is the static value that minimises this norm. As a result, Π is close to 1 for a perfect power match and approaches 0 when $P_{\mathrm{EPC}}(t) = \bar{r}_{\mathrm{p}}$. A negative KPI means that the dynamic load response performs worse than the best static power.

The overall efficiency ($\bar{\eta}$) at which the system produces electricity is defined as

$$\bar{\eta} := \frac{\int_{t=0}^{t_{\mathrm{sim}}} P_{\mathrm{EPC}}dt}{\int_{t=0}^{t_{\mathrm{sim}}}\Delta\bar{H}_{\mathrm{LHV}}\dot{n}_{\mathrm{f}}dt}, \tag{4}$$

with $\dot{n}_{\mathrm{f}}$ as the molar fuel flow and $\Delta\bar{H}_{\mathrm{LHV}}$ as the lower heating value of the fuel at the inlet of the system. The output power ($P_{\mathrm{EPC}}(t)$) is the electrical SOFC stack power minus the power of both blowers and losses in the power converter.

3. SOFC system model

To evaluate the presented KPIs, the SOFC system model is introduced in this section. This work considers a generic stand-alone SOFC system as presented in [23]. The system includes an SOFC stack, two blowers, two heat exchangers, a burner, and an electrical power converter, and is configured as illustrated in Fig. 1. The two blowers regulate the mass flow of both fuel and air. The heat exchangers preheat the mass flows before they enter the SOFC stack, and the electrical power converter determines the electrical current drawn from the system. The burner increases the useable heat in the heat exchangers and ensures that no fuel remains in the exhaust. As the system does not contain any *Anode Off-Gassing Recirculation* (AOGR) or pre-reforming, a partially reformed methane mixture is pre-calculated, based on an AOGR SOFC system with an allothermal pre-reformer for a single pass fuel utilisation of 0.75 [–], a recirculation ratio of 0.65 [–], an oxygen-to-carbon ratio of 2.4 [–] and a reforming equilibrium temperature of 752.15 [K]. This results in an inlet composition of CH4 12.33%, CO2 26.96% CO 2.26%, H2O 40.57% and H2 17.88%.

This paper provides a general description of the system. Detailed modelling equations for the *Balance of Plant* (BOP) components can be found in [18,19,23], and the SOFC stack itself is modelled using the 1D spatially distributed model described in [28]. Furthermore, a summary of the modelling equations for the BOP and SOFC stack is provided in the supplementary material: *SOFC system dynamics*, Section 1. The blower dynamics are described by a first-order state–space model that relates the motor torque input ($u_{\tau}(t)$) to the blower rotational speed ($\omega(t)$), which in turn determines the mass flow rate ($\dot{n}(t)$) [18]. The blower power is determined via ideal compression work as in [29]. It is assumed that the SOFC system controller interfaces with a lower-level mass flow controller, where the blower inertia sets the dominant time constant of that control loop. Heat exchangers are modelled using three control volumes: one for the cold fuel or air stream, one for the hot exhaust gas, and one representing the separating wall [18]. Heat is transferred from the hot exhaust gases to the wall and subsequently from the wall to the cold inlet stream by conduction. The thermal dynamics of each control volume are governed by enthalpy balance equations, including heat losses to the environment. The burner is modelled as a single control volume, with all remaining fuel assumed to combust adiabatically [30]. The resulting temperature dynamics ($T_{\mathrm{b}}(t)$) are determined by an enthalpy balance that includes environmental

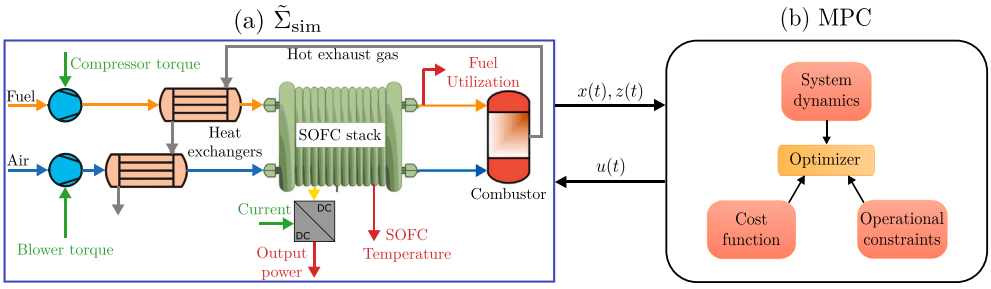

Fig. 1. (a) Stand alone generic SOFC system as presented in [23]. (b) MPC and closed-loop connection with the system.

heat losses. The electrical power converter operates on a time scale several orders of magnitude faster (on the order of microseconds) than the thermal dynamics of the SOFC system. Consequently, its behaviour is approximated using averaged dynamics that include resistive load losses [30]. To match the time scale of interest for system-level control, the current response ($i(t)$) is further filters the requested current ($u_{\text{i}}(t)$) using a low-pass filter.

The SOFC stack is modelled in detail using a one-dimensional spatial representation along the flow direction, in accordance with [28], and represents the co-flow electrolyte-supported Staxera/Sunfire ISM V3.3 short stack. The stack is discretised along the spatial dimension l into multiple control volumes, each consisting of four layers: the fuel channel, the PEN structure, the air channel, and the interconnect. Within each control volume, enthalpy balance equations describe the temperature dynamics of the fuel and air channel ($T_{\text{f}}(t,l)$, $T_{\text{a}}(t,l)$), the PEN structure ($T_{\text{PEN}}(t,l)$) and the interconnect ($T_{\text{IC}}(t,l)$). The PEN structure consists of an active region where the chemical reaction occurs, and inactive regions on both sides. The active region is discretised into five control volumes, while each inactive region is discretised into two control volumes, resulting in a total of nine control volumes for the SOFC stack. The spatial gradients are determined in each control volume via finite difference. Mass balance equations capture the evolution of molar species concentration in both the fuel ($\chi_{\text{f}}(t,l)$) and air channels ($\chi_{\text{a}}(t,l)$). In addition, electrochemical relations are used to determine the current density distribution ($J_{\text{PEN}}(t,l)$) for a uniform cell voltage ($\frac{dU_{\text{cell}}(t,l)}{dl} = 0$) and corresponding total current $i(t) = \int_0^L J_{\text{PEN}}(t,l)dl$. In [28], the 1D spatial SOFC model is validated by means of current power curves from manufacturers' data for various fuel compositions and the spatial distribution of the reaction rates is validated with a 3D CFD model. Both show the validity of the SOFC stack model. Further details on the stack model equations and parametrisation can be found in [28] as well.

The overall system is described as a continuous-time, nonlinear index-1 *Differential Algebraic Equations* (DAE) system as

$$\tilde{\Sigma}_{\text{sim}} : \begin{cases} \dot{\tilde{x}}(t) = & f(\tilde{x}(t), \tilde{z}(t), u(t)) \\ 0 = & g(\tilde{x}(t), \tilde{z}(t)) \\ y(t) = & h(\tilde{x}(t), \tilde{z}(t)) \end{cases}, \tag{5}$$

where $\tilde{x}(t) \in \mathbb{R}^{\text{n}}_{\text{x}}$ is the composite differential state vector, i.e., $\tilde{x}(t) := [\omega^\top(t), T^\top_{\text{HEX}}(t), T^\top_{\text{b}}(t,l), T^\top_{\text{SOFC}}(t,l), \chi^\top_{\text{SOFC}}(t), i(t)]^\top$, with $\omega(t)$ comprising the states of both the fuel and air blower and T_{HEX} representing the temperatures of both the fuel and air heat exchanger. Furthermore, $\tilde{z}(t) \in \mathbb{R}^{n_z}$ is the algebraic state vector, consisting of $J_{\text{PEN}}(t,l)$, $u(t) \in \mathbb{R}^{\text{n}}_{\text{u}}$ is the input vector defined as $u(t) := [u_{\tau,\text{f}}(t), u_{\tau,\text{a}}(t), u_{\text{i}}(t)]^\top$, and the output vector is defined as $y(t) := [\mu_{\text{f}}(t), T_{\text{a}}(t,0), T_{\text{a}}(t,L), P_{\text{EPC}}(t)]^\top$. f captures the differential equations, g represents the algebraic equations and h captures the output equation. Due to the spatial discretisation, the variables that depend on l become a vector within the state vectors, eliminating the explicit dependence on l in the system description.

4. Baseline MPC

The proposed MPC strategies consist of a baseline formulation, introduced in this section, and additional thermal stress–aware constraints and cost terms, which are incorporated into this baseline in the next section. The MPC solves an *Optimal Control Problem* (OCP) at each time step. The OCP is a general optimisation problem consisting of a cost function, equality constraints and inequality constraints. The cost function defines the system's desired behaviour. For the SOFC system, this means that the output $y(t)$ tracks the setpoint $r(t)$ while rapid state changes are discouraged by penalising the state derivative. The system dynamics are captured in the prediction model and incorporated into the equality constraints along with the initial states. The inequality constraints define the feasible operating region.

4.1. Prediction model

The prediction model represents the system dynamics within the OCP and should accurately reflect the SOFC system model of Section 3. Since it is evaluated on a fixed time grid, the dynamics must match the grid's timescale. A horizon of at least 100 [s] with a resolution of seconds is required to capture both rapid local PEN temperature changes and long-term PEN temperature evolution. However, the SOFC system model $\tilde{\Sigma}_{\text{sim}}$ includes much faster dynamics that cannot be integrated accurately on such a time grid. These fast states, the molar gas species ($\chi_{\text{f}}(t,l)$, $\chi_{\text{a}}(t,l)$) and the gas channel temperatures ($T_{\text{f}}(t,l), T_{\text{a}}(t,l)$), are therefore assumed to converge instantaneously. The corresponding differential equations are set to steady state, and are treated as algebraic equations. Since the algebraic species concentrations are uniquely determined by the current density ($J_{\text{PEN}}(t,l)$) and the molar reaction rates ($r_{\text{MSR}}(t,l), r_{\text{WGS}}(t,l)$), the reaction rates are used as algebraic variables instead of $\chi_{\text{f}}(t,l)$, $\chi_{\text{a}}(t,l)$, yielding a minimal and linearly independent set of algebraic equations. An overview of these derivations for the prediction model is provided in the supplementary material: *SOFC system dynamics*, Section 1.7. This assumption limits the model's validity to situations where the current changes significantly more slowly than the approximated time scales. If the current varies at comparable or faster rates, the model no longer holds. The low-pass filter in the electrical power converter mitigates this issue by suppressing such rapid variations. The resulting prediction model is an index-1 continuous-time DAE system with state vector $x(t)$ and algebraic vector $z(t)$, and evaluated in the OCP in implicit form as

$$\Sigma_{\text{MPC}} : \begin{cases} 0 = & F(\dot{x}(t), x(t), z(t), u(t)) \\ y(t) = & H(x(t), z(t)) \end{cases}. \tag{6}$$

4.2. Collocation

In the OCP, the DAE system given by (6) is integrated using an implicit collocation method over a time grid containing one point per sampling instant, with each interval subdivided by Gauss–Legendre collocation points [31]. The number of collocation points (n_{t}) and the sampling time are parameters that determine the integration accuracy,

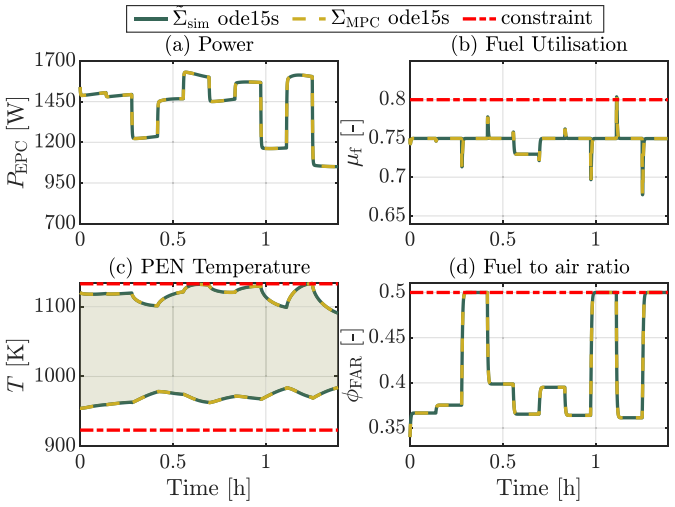

Fig. 2. Open-loop system simulation of the simulation model $\bar{\Sigma}_{sim}$ and prediction model Σ_{MPC} integrated with `ode15s` and a maximal time-step of 0.1 [s]. The spatial PEN temperature distribution is depicted as a time-varying temperature region.

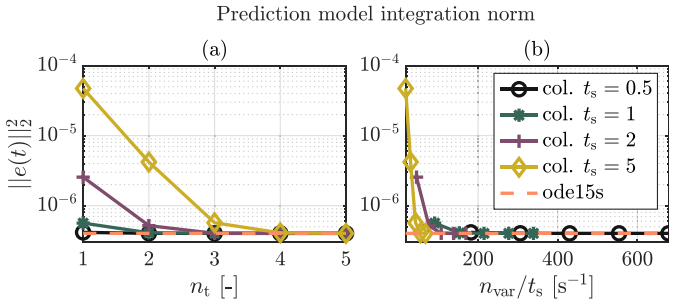

Fig. 3. Validation of the prediction model in the state norm comparing the integration accuracy of the collocation-based prediction model and the DAE simulation model as a function of (a) sampling times and collocation points, and (b) the resulting number of integration variables per second of prediction horizon in the optimal control problem.

control flexibility and the number of optimisation variables. The prediction model is validated against the original simulation model $\bar{\Sigma}_{sim}$, which is integrated with `ode15s` with a maximum step size 0.1 [s] in a 5000 s open-loop simulation with various setpoints, yielding $5 \cdot 10^4$ data points. Simulated under the same conditions, the prediction model shows good agreement with the simulation model (Fig. 2) and achieves a state-error norm of $\|\bar{x}(t) - D(x(t), z(t))\|_2^2 = 4.0 \cdot 10^{-7}$. Here, $D(x(t), z(t))$ maps the prediction-model states to the simulation-model states and contains system equations. The prediction model is integrated using the collocation method for $t_s \in 0.5, 1, 2, 5$ [s] and $n_t \in 1, 2, 3, 4, 5$. Fig. 3(a) shows the state-error norm, illustrating convergence to `ode15s` accuracy as n_t increases. Fig. 3(b) highlights how sampling time and number of collocation points affect the number of integration variables in the OCP. In general, a larger sampling time with more collocation points provides similar accuracy with fewer optimisation variables but reduces control flexibility since the input is held constant over each sampling interval.

4.3. Safety constraints

Bounds on the states and inputs are applied according to the valid operating range of the system. Additional safety-related operational constraints are included as listed in Table 1. Upper bounds are placed

Table 1
Safety constraints for the SOFC system applied in the baseline MPC, as given in the operating manual [32].

Constraint name	Symbol	Value	Unit
Fuel utilisation	$\mu_f(t)$	≤0.80	[–]
Air utilisation	$\mu_a(t)$	≤0.60	[–]
Fuel-to-air equivalence ratio	$\phi_{FAR}(t)$	≤0.50	[-]
PEN temperature	$T_{PEN}(t, l)$	≥923 ≤1133	[K]
Cell voltage	$U_{cell}(t)$	≥0.6	[V]

on the fuel utilisation ($\mu_f(t)$), the fuel-to-air equivalence ratio ($\phi_{FAR}(t)$) and various thermal stress constraints as discussed in Section 5. The air-utilisation constraint is omitted because it is implicitly enforced through the combination of the fuel-utilisation limit and the fuel-to-air equivalence ratio. The fuel utilisation is a fraction and therefore can be represented exact in linear form as

$$i(t) - \bar{\mu}_f a_\mu \dot{n}_f(t) \leq 0, \tag{7}$$

with $\bar{\mu}_f$ as the maximum fuel utilisation and a_μ as a stoichiometric conversion coefficient from molar fuel flow to current. Similarly, the Fuel-to-Air equivalence ratio constraint is given as

$$\dot{n}_f(t) - \bar{\phi}_{FAR} a_\phi \dot{n}_a(t) \leq 0, \tag{8}$$

with $\bar{\phi}_{FAR}$ as the maximum fuel-to-air equivalence ratio and a_ϕ as a stoichiometric conversion coefficient from molar air flow to molar fuel flow. All constraints are given as linear half-space constraints, resulting in the feasible operation regions $\mathcal{X} := \{x(t) | G_x^\top x(t) \leq b_x\}$, $\mathcal{U} := \{u(t) | G_u^\top u(t) \leq b_u\}$ and $\mathcal{X}_\delta := \{\dot{x}(t) | G_{x_\delta}^\top \dot{x}(t) \leq b_{x_\delta}\}$.

4.4. Cost function

The cost function ($J(\cdot)$) is given as a combination of tracking cost ($J_e(\cdot)$) and cost on the state derivative ($J_\delta(\cdot)$) as

$$J\left(r_t, \dot{x}_{\tau|t}, x_{\tau|t}, z_{\tau|t}\right) = J_e(r_t, x_{\tau|t}, z_{\tau|t}) + J_\delta(\dot{x}_{\tau|t}), \tag{9}$$

with the quadratic tracking cost as

$$J_e(r_t, x_{\tau|t}, z_{\tau|t}) = \|r_t - H(x_{\tau|t}, z_{\tau|t})\|_S^2, \tag{10}$$

with output reference setpoint r_t and the positive definite weighting matrix S. The cost on the state derivative is quadratic as well

$$J_\delta(\dot{x}_{\tau|t}) = \begin{bmatrix} \dot{x}_u \\ \dot{x}_{BOP} \\ \dot{T}_{PEN} \\ \dot{T}_{IC} \end{bmatrix}^\top \begin{bmatrix} Q_1 & 0 & 0 & 0 \\ 0 & Q_2 & 0 & 0 \\ 0 & 0 & Q_3 & 0 \\ 0 & 0 & 0 & Q_4 \end{bmatrix} \begin{bmatrix} \dot{x}_u \\ \dot{x}_{BOP} \\ \dot{T}_{PEN} \\ \dot{T}_{IC} \end{bmatrix}, \tag{11}$$

with $\dot{x}_u(t) = [\dot{\omega}^\top(t), \dot{i}(t)]^\top$, $\dot{x}_{BOP} = [\dot{T}_{HEX}^\top(t), \dot{T}_b(t)]^\top$ and the positive definite weighting matrices $Q_k = q_k I$ for $k = \{1, 2, 3, 4\}$, with I as identity matrix of suitable dimension, accounting for various options to tune the MPC.

4.5. Optimal control problem

The OCP is formulated in continuous time following the standard formulation as

$$\min_{x_{\cdot|t}, u_{\cdot|t}} \quad \int_0^{t_h} J\left(r_t, \dot{x}_{\tau|t}, x_{\tau|t}, z_{\tau|t}\right) d\tau, \tag{12a}$$

$$\text{s.t.} \quad 0 = F\left(\dot{x}_{\tau|t}, x_{\tau|t}, z_{\tau|t}, u_{\tau|t}\right), \tag{12b}$$

$$x_{\tau|t} \in \mathcal{X}, \tag{12c}$$

$$u_{\tau|t} \in \mathcal{U}, \tag{12d}$$

$$\dot{x}_{\tau|t} \in \mathcal{X}_\delta, \tag{12e}$$

$$x_{0|t} = x(t), \tag{12f}$$

$$\tau \in [0, t_h], \tag{12g}$$

with t_h as the prediction horizon, and τ as the time prediction index. This formulation represents a setpoint-tracking OCP in which only the current setpoint reference is known and assumed to remain constant over the prediction horizon.

The state derivative can be expressed as a linear combination of the collocation points and the shooting node, making the cost $J_\delta(\dot{x}_{\tau|t})$ and constraint $\dot{x}_{\tau|t} \in \mathcal{X}_\delta$ effectively state-dependent.

5. Thermal stress awareness

Both the spatial and temporal temperature gradients in the SOFC stack cause thermal stress, as discussed in the introduction. Extending the baseline MPC formulation, this section introduces cost terms and constraints for several thermal stress–aware MPC formulations, either addressing the temporal or spatial temperature gradients.

5.1. Spatial temperature gradient constraint (STGC)

The first MPC variation accounts for thermal stress by constraining the spatial temperature gradient as

$$\delta_l T^- \leq \frac{dT_{\text{PEN}}(t,l)}{dl} \leq \delta_l T^+, \tag{13}$$

with $\delta_l T^-$ and $\delta_l T^+$ as the minimum and the maximum spatial constraint. This introduces additional constraints in the compact set $\mathcal{X}$. There are no constraints on the derivative ($\mathcal{X}_\delta = \mathbb{R}^{n}_{x}$). Cost is imposed on the input (q_1) and BOP state (q_2) derivatives, but not explicitly on the local temperature variations inside the SOFC stack as $q_3 = 0$ and $q_4 = 0$.

5.2. Current ramping limit (CRL)

The current ramping limit reflects standard practice, assuming that the SOFC stack's spatial distribution is unknown. Therefore, the current ramping rate is limited as

$$\mathcal{X}_\delta := \delta_t i^- \leq \dot{i}(t) \leq \delta_t i^+, \tag{14}$$

with $\delta_t i^-$ and $\delta_t i^+$ as the minimum and the maximum limit, respectively. A cost is imposed on the input (q_1) and the BOP state (q_2) derivatives, while no cost is assigned to the spatially distributed temperatures ($q_3 = 0,\ q_4 = 0$).

5.3. Temporal temperature gradient constraint (TTGC)

The first proposed option to limit the time-dependent thermal stress is to constrain the local temporal temperature gradient as

$$\mathcal{X}_\delta := \delta_t T^- \leq \dot{T}_{\text{PEN}}(t) \leq \delta_t T^+, \tag{15}$$

with $\delta_t T^-$ and $\delta_t T^+$ as the minimum and the maximum constraint, respectively. The state-derivative costs q_1 and q_2 are employed. No further cost is given to the spatially distributed temperatures ($q_3 = 0, q_4 = 0$).

5.4. Temporal temperature gradient cost (TTGQ)

The second proposed strategy to limit the time-dependent thermal stress involves an additional cost associated with the temporal temperature gradient in the PEN structure. In addition to q_1 and q_2, a positive value is chosen for q_3. q_4 stays equal to zero. Furthermore, it can be beneficial to keep the constraints on the temporal temperature gradient but with more permissive bounds than for the TTGC. The values for these costs and constraints are given in Section 6.

6. Implementation

The SOFC system model and MPC formulations are implemented in MATLAB R2024b. The simulation model is integrated using `ode15s` for each time step.

The MPC strategies are implemented using the CasADi-3.7.0 toolbox and the accompanying FATROP solver [33]. FATROP implements a nonlinear primal–dual interior-point algorithm that exploits the OCP structure [34]. A sampling time of 5 [s] with 4 collocation points is selected for the MPC prediction model, providing good accuracy with minimal OCP variables as shown in Section 4.2. This avoids an excessively large and slow OCP while maintaining adequate control flexibility. The horizon length is set to 100 [s] to capture the long-term temperature dynamics.

Table 2 summarises the various strategies as described in Sections 4 and 5. The results analyse the spatial and temporal temperature gradients separately. For the spatial gradient, four strategies are evaluated: Baseline, Baseline-25, STGC-25 and STGC-20. The Baseline is the MPC formulation as defined in Section 4 and tracks reference setpoints from the nominal operating curve. Baseline-25 follows the same MPC formulation, cost and constraints as the Baseline strategy. Only this strategy tracks setpoints that keep the spatial gradient below 25 [K/cm] in the associated steady-state. STGC-25 and STGC-20 enforce the spatial gradient constraint from Section 5.1 at 25 [K/cm] and 20 [K/cm], respectively, while tracking reference setpoints that are feasible within these limits. The weighting matrices (S, Q_k, $k \in 1,2,3,4$) are tuned per strategy to ensure fast power tracking and steady-state convergence, as shown in Table 3. The choice of these spatial temperature gradient limits is not based on material properties such as critical stress or fatigue limits, as these vary significantly across stacks and applications. Instead, these limits reflect the point at which a change of operating mode becomes necessary. This shows how the spatial temperature gradient can be reduced relative to nominal operation, regardless of stack- and application-specific limits.

The temporal temperature gradient is analysed using the CRL, TTGC and TTGQ strategies across a static reference setpoint scenario for the three cases Stat-1/2/3, with constraints that allow operation from slow to fast. The ramping limits for the CRL are doubled in each case: $-\delta_t i^- = \delta_t i^+ \in \{1/30, 1/15, 1/6\}$ [A/s], where 1/30 [A/s] is the typical limit for the SOFC stack [32]. In addition, these strategies are also assessed under a dynamic reference-setpoint scenario for the slow case Dyn-4 $-\delta_t i^- = \delta_t i^+ = 1/30$ [A/s]. TTGC and TTGCQ limits are tuned to match the CRL strategy results, ensuring comparable power-tracking and temporal-gradient KPIs. The constraints specific to each case are given in Table 4. These weighting matrices are the same for each case and strategy, and are given as: $q_1 = 100$, $q_2 = 1$, $q_4 = 0$, $s = (1, 0, 0.1, 10)$, $q_3 = 0$ for the CRL and TTGC strategies and $q_3 = 10^5$ for the TTGQ strategy.

7. Results

The proposed MPC strategies are evaluated using the SOFC system simulation model and the defined KPIs. The results are presented in this section, beginning with the Baseline and STGC strategies, which assess spatial temperature gradients. Then the temporal temperature gradient is assessed through the evaluation of the CRL, TTGC and TTGQ strategies. These three strategies are first evaluated for the static setpoint reference cases Stat-1/2/3 and then for the dynamic setpoint reference case Dyn-4.

7.1. Spatial temperature gradients

The spatial temperature gradients result from variations in the temperature distribution within the SOFC. Although this distribution changes during transient operation, it is largely determined by the operating point. Consequently, the spatial temperature gradient can be

Table 2
Summary of the introduced MPC strategies, corresponding acronyms and for which scenarios the strategies are evaluated.

Acronym	Corresponding strategy	Evaluation scenario
Baseline	Baseline MPC as presented in Section 4	Spatial temperature gradient
Baseline-25	Baseline with adapted tracking setpoints	
STGC	Spatial temperature gradient constraint	
CRL	Current ramping limit	Temporal temperature gradient
TTGC	Temporal temperature gradient constraint	
TTGQ	Temporal temperature gradient cost	

Table 3
Weights used for the spatial temperature gradient MPC strategies with output weighting matrix $S = \mathrm{diag}(s)$, responding to the output $y(t) := [\mu_f(t), T_a(t,0), T_a(t,L), P_{EPC}(t)]^\top$.

Strategy	q_1	q_2	q_3	q_4	s
Baseline	100	1	0	0	(1, 0, 0.1, 10)
Baseline-25	100	1	0	0	(1, 0, 0.1, 10)
STGC-25	10	1	0	0	(1, 100, 0, 10)
STGC-20	10	1	0	0	(1, 20, 0, 10)

Table 4
Specific constraints and weights for CRL, TTGC and TTGQ MPC strategies per case.

	CRL		TTGC		TTGQ		
	$\delta_t i^-$	$\delta_t i^+$	$\delta_t T^-$	$\delta_t T^+$	$\delta_t T^-$	$\delta_t T^+$	q_3
Stat-1	−1/30	1/30	−0.11	0.11	−0.15	0.12	10^5
Stat-2	−1/15	1/15	−0.16	0.16	−0.16	0.16	10^5
Stat-3	−1/6	1/6	−0.25	0.19	−0.50	0.19	10^5
Dyn-4	−1/30	1/30	−0.08	0.08	−0.13	0.10	10^5

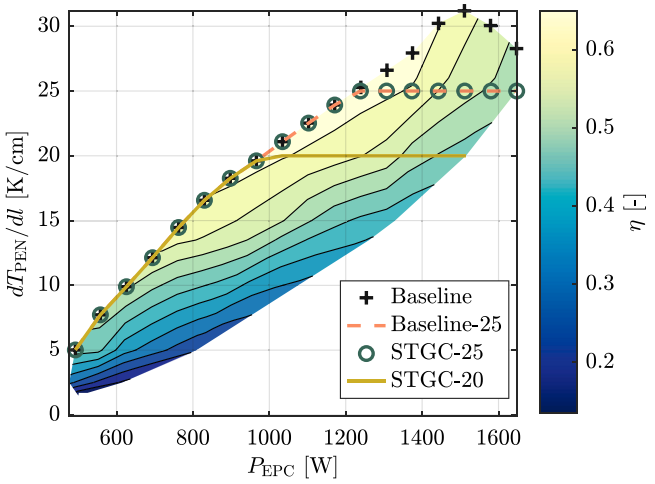

Fig. 4. Static operating map of the SOFC system, highlighting the achievable efficiency for a certain power and a spatial temperature gradient with the steady-state operating envelope of the spatially constrained MPC policies.

represented in a steady-state operating map as shown in Fig. 4. The map shows the SOFC's operational envelope and electrical efficiency. Each power level spans a range of spatial temperature gradients, highlighting the trade-off that reducing the spatial temperature gradient at a given power level also reduces the electrical efficiency. Conversely, reducing the gradient below 25 [K/cm] also limits the maximum achievable power. The four spatial temperature gradient operating strategies are indicated in the map. The Baseline strategy follows the maximum-efficiency setpoints and therefore exhibits the highest gradients. Baseline-25 and STGC-25 operate along the same power curve but restrict the gradient to 25 [K/cm], resulting in a reduced efficiency. The STGC-20 strategy further limits the gradient to 20 [K/cm], which decreases both efficiency and the maximum attainable power.

The STGC strategies are dynamically evaluated in simulation by varying the power from 1470 [W] to its minimum of 700 [W] for 800 [s], followed by stabilisation at the maximum power 1648 [W] for 2200 [s]. The results of these strategies are compared to the Baseline MPC in Fig. 5, with (a) the dynamic power tracking, (b) PEN temperature, (c) spatial temperature gradient, (d) Power tracking KPI, (e) efficiency KPI, and (f) spatial temperature gradient KPI, with the normalised PEN length on the y-axis and the KPI value on the x-axis. These results demonstrate that the Baseline and Baseline-25 MPC are unconstrained in the load-tracking response during both ramp-down and ramp-up. The STGC-25 and STGC-20 are free to respond to ramping down and to ramping up until 1450 [W]. The temperature then increases until the STGC is reached, at which point the power has to be reduced. Finally, both temperature and power are gradually increased in compliance with the STGC to the maximum achievable power. The Baseline-25 strategy reaches the same power as the Baseline but limits the setpoint spatial temperature gradient to 25 [K/cm]. Because this strategy is unconstrained, it first overshoots the limit during the transient, reaching 29 [K/cm]. It also achieves slightly faster load-tracking performance than the Baseline, as the reference setpoint prescribes a higher temperature at minimum power. The STGC-25 strategy eventually reaches the same maximum power as the Baseline. However, increasing the overall temperature of the SOFC stack under the spatial gradient constraint takes over 1000 [s]. As a result, the STGC restricts the available power for fast ramp-ups. Furthermore, the SOFC stack operates at a slightly lower overall efficiency of 52.7% instead of 53.47% as highlighted by the KPI. The STGC-20 strategy exhibits reduced maximum power and efficiency due to further restrictions on the spatial constraints, achieving 1529 [W] with an efficiency of 48.6%. As expected, the Baseline shows the highest spatial temperature gradient KPI. The Baseline-25 and STGC-25 perform similarly, with Baseline-25 only marginally higher, while STGC-20 achieves the lowest KPI. The highest spatial temperature gradient KPI occurs between 60%–80% of the PEN length. Endothermic methane reforming cools the inlet region, while exothermic hydrogen-oxidation and water–gas-shift reactions heat the downstream section. The interaction between these cooling and heating zones produces the maximum spatial stress.

Since the STGC strategies must operate dynamically within the spatial temperature gradient, achieving maximum power with them is difficult. Therefore, it takes a considerable time to reach the maximum power. Nevertheless, the spatial gradient constraint primarily affects the steady-state operating setpoint since most of the dynamic load response remains unconstrained. To achieve a similar power level while maintaining a lower spatial temperature gradient, the heat generated in the stack is dissipated using a hotter airflow. This higher air temperature is achieved by combusting more fuel in the burner while maintaining the same level of heat dissipation by increasing the air flow rate. As a result, electrical efficiency is reduced as illustrated in Fig. 5. The system ultimately cannot dissipate all generated heat through the airflow while maintaining the specified spatial temperature gradient, resulting in a reduced maximum power.

Although the reported maximum spatial temperature gradient in these results might seem high for some SOFC stacks, these results are

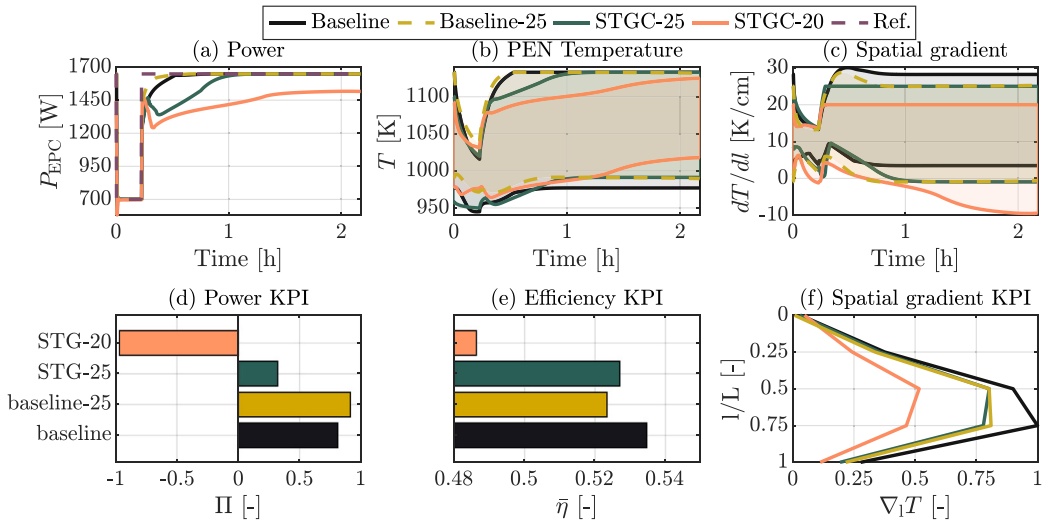

Fig. 5. Dynamic operation of the SOFC, with the STG strategies, for a ramp-down, ramp-up scenario, showing the (a) electrical power tracking, (b) PEN temperature distribution, (c) the spatial temperature gradient and (d–f) the corresponding KPIs with the KPI values shown on the x-axis, and either the control strategies (d–e) or the spatial dimension (f) on the y-axis.

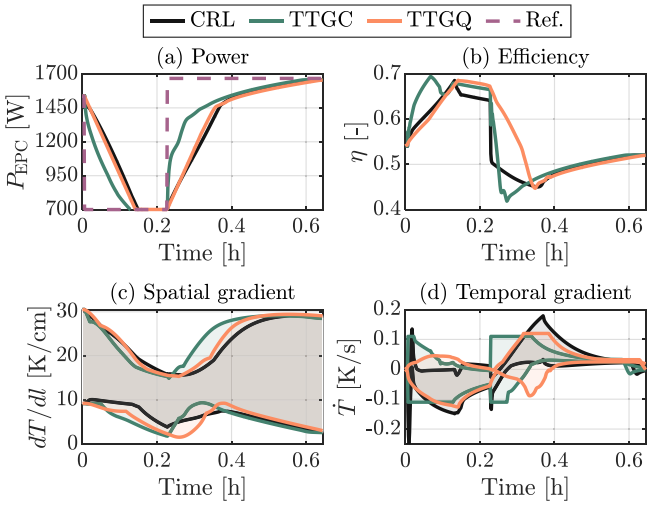

Fig. 6. Comparison of the TTGC, TTGQ with the CRL MPC strategies for case Stat-1, showing the dynamic behaviour of the (a) power, (b) efficiency, (c) spatial and (d) temporal temperature gradient.

consistent with those shown in [11,28]. Furthermore, the model is based on a Staxera stack ISM V3.3, which is known to have a high tolerance for temperature differences [32].

7.2. Temporal temperature gradient

The CRL, TTGC and TTGQ strategies are first evaluated using a simulation scenario similar to the previous scenario. The load response is tested by varying the power from 1470 [W] to its minimum for 800 [s] and to the maximum for 1500 [s]. The CRL, TTGC and TTGQ strategies are evaluated in this scenario for the three cases Stat-1/2/3 and will be discussed next.

The simulation results for case Stat-1 are presented in Fig. 6, showing the (a) power tracking, (b) electrical efficiency, (c) spatial temperature gradient, and (d) temporal temperature gradient. The CRL strategy optimises load response within the current ramping limits. The TTGC strategy optimises power tracking while respecting the temporal temperature gradient constraints. The TTGQ strategy balances temporal temperature gradients and power tracking performance in the cost function. The respective derivative constraints limit the load ramp-up until the maximum fuel-to-air equivalence constraint is reached, after which the power and temperature are increased more slowly. The maximum power setpoint is therefore reached simultaneously, but the TTGC ramps up and down significantly faster before this constraint is reached. The temporal temperature gradient clearly responds to the dynamic load changes. For the TTGC, the constraint is directly met when the power is changed, maximising the load-tracking response. For CRL, the temporal temperature gradient steadily increases as the current is ramped up, and the TTGQ follows this ramp until the relaxed TTGC maximum is reached. The spatial temperature gradient and efficiency are mostly determined by the tracking setpoint, as discussed in Section 7.1, while small variations are observed during transient operation.

The power tracking and temporal temperature gradient KPIs for all static setpoint cases are presented in Fig. 7. Analysing the KPIs for the case Stat-1 shows that the TTGC strategy achieves the fastest load response, resulting in a higher load-tracking KPI than the TTGQ and CRL strategies. The TTGQ's load-tracking KPI performs slightly better than the CRL's KPI. The cumulative temporal temperature gradient indicates that most thermal stress occurs in the last active control volume of the SOFC stack. The highest PEN temperature in the stack typically occurs at this location. Consequently, this is where temperatures change the most during dynamic temperature fluctuations. The CRL strategy exhibits the highest temporal stress KPI, while the TTGC and TTGQ strategies reduce the temporal stress KPIs by 15.50% and 18.00%, respectively. In the first two active control volumes, the TTGC and TTGQ strategies increase the temporal stress KPI by 7.62% and 3.97%, respectively, trading a small increase in stress in less critical control volumes for a reduction in stress in the more critical control volumes. The spatial gradient KPI and efficiency do not exhibit structural differences between the MPC strategies and are therefore omitted from the figure.

In case Stat-2, the constraints are tightened as specified in Table 4, which leads to an improved load-tracking response as confirmed by the power-tracking KPIs for case Stat-2 in Fig. 7. The temporal temperature gradient KPI represents the KPI value on the x-axis and the normalised PEN length on the y-axis. The TTGC strategy is tuned to achieve the fastest load response, whereas the TTCQ and CRL strategies exhibit comparable load-tracking performance. The difference in the load-tracking KPI between the CRL and TTGC strategies amounts to 11.72%.

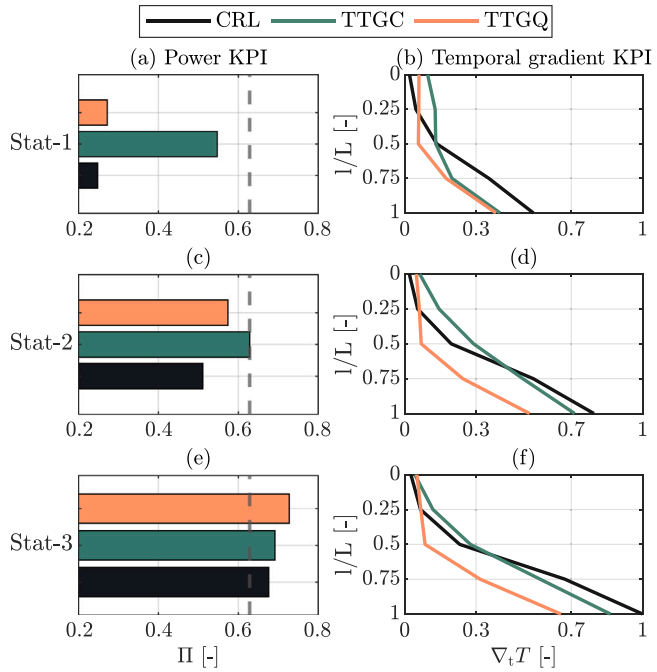

Fig. 7. Comparison of the (a, c, e) power tracking and (b, d, f) temporal temperature gradient KPIs for the CRL, TTGC, and TTGQ MPC strategies. Each row corresponds to cases Stat-1–3, while the KPIs are arranged by column. The KPI values are shown on the x-axis, and either the control strategies (a, c, e) or the spatial dimension (b, d, f) on the y-axis. The dashed grey line in the power-tracking KPI indicates the shared reference.

The TTGC strategy continues to impose more stress at the beginning of the active PEN length (4.37%) than the CRL strategy, while inducing less stress at the end of the active PEN length (8.08%). In case Stat-2, the TTGC strategy exhibits a less favourable trade-off between load-tracking response and accumulated temporal stress compared to case Stat-1. The TTGQ strategy reduces the accumulated temporal stress in the final control volume by 27.03% relative to the CRL strategy.

Case Stat-3 further increases the dynamic load response. While the load-tracking performance is now relatively similar across all strategies, the TTGC and TTGQ perform 1.55% and 5.13% better, respectively, than the CRL strategy. Where the TTGC and TTCQ reduce the accumulated temporal stress in the last control volume by 13.54% and 34.33%, respectively, compared to the CRL strategy.

The second scenario evaluates practical performance via a dynamic load profile. Case Dyn-4 uses a dynamic load profile that combines harmonics with a step reference, resulting in a reference that alternates between fast and slower changes. The CRL, TTGC and TTGQ are simulated for this scenario to show how changing the load profile impacts the results. To make the results comparable with those of the CRL strategy, the constraints of the TTGC and TTGQ strategies must be adjusted so that the power-tracking KPI and the accumulated temporal gradient are comparable to those of the CRL strategy.

The results of the dynamic load-tracking scenario are shown in Fig. 8, with (a) power tracking, (b) power tracking KPI, (c) temporal temperature gradient KPI and (d) MPC computation time statistics. The results are similar to those of the previous scenario for case Stat-1. However, as the overall load change of the dynamic reference is slower than in the case Stat-1, the overall load tracking of all three strategies performs better with lower temporal temperature gradient KPIs. The CRL strategy achieves a load-tracking KPI of 0.6219 [–] and a maximum accumulated temperature gradient KPI of 0.2389 [–]. The TTGC strategy improves the load-tracking KPI by 17.85% and exhibits a slightly smaller accumulated temporal temperature gradient KPI of 0.2054 [–] compared to CRL. The TTGQ achieves a similar load-tracking KPI of 0.6426 [–] and reduces the maximum accumulated temporal temperature gradient KPI to 0.1513 [–]. This scenario highlights that, despite the strong influence of the load profile on performance, the TTGC and TTGQ strategies can be tuned to consistently outperform the CRL approach. The computational time statistics of the MPC strategies show that the average computation time falls within the 5 [s] sampling interval, but there is a large range of variation that violates this real-time application limit. While these results demonstrate a proof of concept, further development is needed to make the implementation suitable for real-time applications. While the current implementation using MATLAB and CasADi is not yet time-efficient, its computational performance can be significantly improved for practical deployment. This can be achieved through optimised C-code generation, execution on embedded hardware platforms [35], and the use of advanced techniques such as real-time iteration schemes to meet real-time requirements [36].

8. Conclusion and future work

This research addresses thermal stress in the PEN structure associated with dynamic operation of SOFC systems by proposing several MPC strategies with thermal stress-aware cost terms and constraints that explicitly account for local spatial and temporal temperature gradients during dynamic operation. The key findings are:

- The results in Section 7.1 show that the spatial temperature gradient is primarily governed by the chosen steady-state operating setpoint and only marginally influenced by the dynamic operation strategy. However, comparing the results from Baseline-25 and STGC-25 shows that actively suppressing the spatial temperature gradient via the reference setpoint results in constraint violations, whereas the proposed STGC strategy effectively limits these violations. Still, the STGC only decreases the accumulated spatial temperature gradient KPI marginally compared to the Baseline-25.
- A clear trade-off exists where a lower spatial temperature gradient results in a lower electrical efficiency for a given operating power, following the operational map depicted in Fig. 4. Moreover, imposing strict limits on the allowable spatial temperature gradient reduces the maximum achievable electrical power.
- Following Section 7.2, constraining dynamic load-tracking capability mainly impacts the temporal temperature gradient, as faster load transitions induce larger temporal gradients. Compared with the state-of-practice CRL approach, the TTGC and TTGQ strategies leverage this trade-off more effectively by either enabling faster load tracking at comparable temporal temperature gradients or reducing temporal temperature gradients while maintaining similar load-tracking performance. In the dynamic load-tracking scenario, the TTGC strategy enhances load tracking by 17.85% while achieving a slightly lower maximum temporal temperature gradient KPI compared to the CRL strategy. Similarly, the TTGQ strategy reduces the maximum temporal temperature gradient KPI by 8.76%, alongside a 2.07% improvement in the load-tracking KPI compared to the CRL strategy.
- Assessing the spatial and temporal temperature gradient KPIs shows that overall, the most stress due to the spatial temperature gradient appears between 50%–75% of the active PEN length. The greatest stress due to the temporal temperature gradient occurs at 100% of the active PEN length. These are expected to be the most vulnerable locations of the PEN structure during dynamic operation.

Future research should investigate how the TTGC and TTGQ strategies can be tailored to specific applications, since their relative benefits depend heavily on system characteristics, hybrid system configurations and load-response requirements. A promising direction is to develop

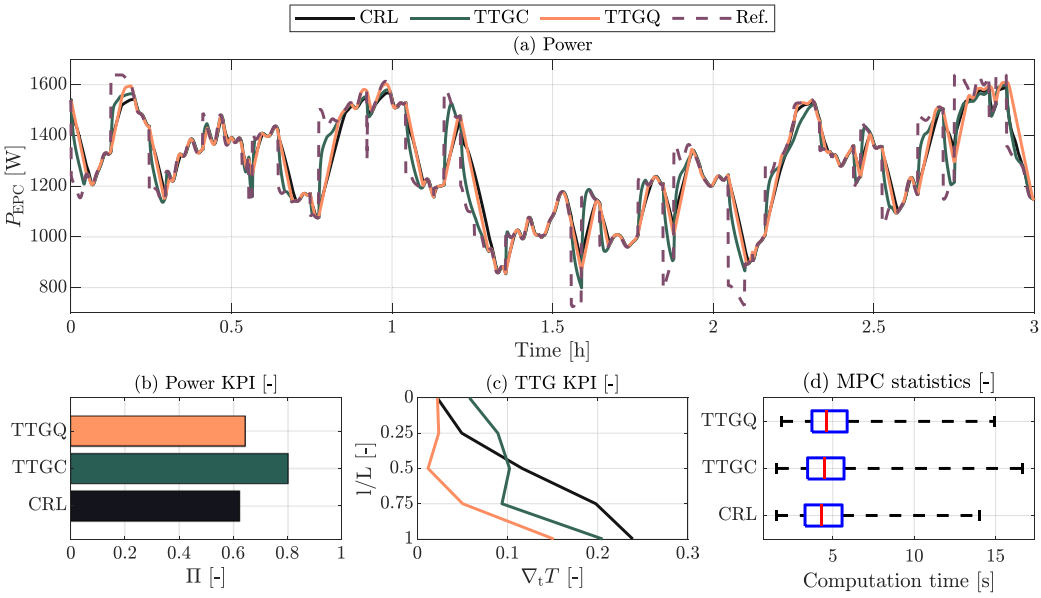

Fig. 8. Dynamic responses and corresponding KPIs of the CRL, TTGC, and TTGQ strategies, for case Dyn-4 with the(a) power tracking, (b) power KPI, (c) temporal temperature gradient KPI and (d) MPC computation time statistics.

systematic tuning methodologies that optimise these strategies with respect to dynamic performance, acceptable degradation risk and lifetime requirements. Such an approach could enable the definition of application-specific safe operating limits and cost-optimal parameter settings for TTGC and TTGQ, rather than ad-hoc tuning based on the CRL strategy for specific load profiles and scenarios.

In this work, spatial and temporal temperature gradients are analysed separately to emphasise their distinct characteristics and the implications they entail. However, future strategies may aim to address both simultaneously, for example, by integrating the STGC approach with the TTGC or TTGQ strategies into a unified framework. This framework should evaluate dynamic strain to assess mechanical stress, thereby enabling analyses of degradation, lifetime and feasibility of dynamic operation for long-voyage shipping. Such a feasibility assessment also necessitates a detailed evaluation of maritime load profiles for specific maritime applications.

CRediT authorship contribution statement

Matthis H. de Lange: Writing – original draft, Visualization, Supervision, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Pablo Segovia:** Writing – review & editing, Supervision, Methodology, Formal analysis, Conceptualization. **Rudy R. Negenborn:** Writing – review & editing, Supervision, Methodology, Investigation, Conceptualization. **Lindert van Biert:** Writing – review & editing, Supervision, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Matthis de Lange reports financial support was provided by European Climate Infrastructure and Environment Executive Agency. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The research is supported by the European Consortium 'HELENUS' (Grant agreement ID: 1010567). The HELENUS Project aims to demonstrate the applicability, scalability and fuel-flexibility of highly efficient solid oxide fuel cells (SOFCs) in various large ship applications.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jpowsour.2026.240440.

Data availability

Data will be made available on request.

References

[1] International Maritime Organization, Fourth IMO ghg study 2020, 2020.
[2] F.M. Kanchiralla, S. Brynolf, T. Olsson, J. Ellis, J. Hansson, M. Grahn, How do variations in ship operation impact the techno-economic feasibility and environmental performance of fossil-free fuels? A life cycle study, Appl. Energy 350 (2023) 121773.
[3] B. Stolz, M. Held, G. Georges, K. Boulouchos, Techno-economic analysis of renewable fuels for ships carrying bulk cargo in Europe, Nat. Energy 7 (2022) 203–212.
[4] H. Xing, S. Spence, H. Chen, A comprehensive review on countermeasures for CO2 emissions from ships, Renew. Sustain. Energy Rev. 134 (2020) 110222.
[5] L. van Biert, M. Godjevac, K. Visser, P. Aravind, A review of fuel cell systems for maritime applications, J. Power Sources 327 (2016) 345–364.
[6] B. Van Veldhuizen, L. Van Biert, P.V. Aravind, K. Visser, Solid oxide fuel cells for marine applications, Int. J. Energy Res. 2023 (1) (2023) 5163448.
[7] R. Scataglini, M. Wei, A. Mayyas, S.H. Chan, T. Lipman, M. Santarelli, A direct manufacturing cost model for solid-oxide fuel cell stacks, Fuel Cells 17 (6) (2017) 825–842.
[8] R.J. Braun, Optimal Design and Operation of Solid Oxide Fuel Cell Systems for Small-Scale Stationary Applications (Ph.D. thesis), The University of Wisconsin-Madison, 2002.
[9] J. Hollmann, S. Kabelac, Steady-state and transient operation of solid oxide fuel cell systems with anode off-gas recirculation within a highly constrained operating range, Energies 16 (23) (2023) 7827.
[10] F. Mueller, F. Jabbari, J. Brouwer, On the intrinsic transient capability and limitations of solid oxide fuel cell systems, J. Power Sources 187 (2) (2009) 452–460.

[11] Z. Zeng, Y. Qian, Y. Zhang, C. Hao, D. Dan, W. Zhuge, A review of heat transfer and thermal management methods for temperature gradient reduction in solid oxide fuel cell (SOFC) stacks, Appl. Energy 280 (2020) 115899.
[12] T.M.M. Heenan, X. Lu, J.B. Robinson, F. Iacoviello, D.J.L. Brett, P.R. Shearing, Thermally driven SOFC degradation in 4D: Part II. Macroscale, J. Electrochem. Soc. 165 (11) (2018) F932.
[13] W. Cai, R. Zhou, C. Wang, C. Xie, L. Xiao, Z. Zhang, C. Yang, J. Yuan, On characteristics and research development of coupled fuel cell stack performance and stress, Appl. Energy 388 (2025) 125719.
[14] A. Chaisantikulwat, C. Diaz-Goano, E.S. Meadows, Dynamic modelling and control of planar anode-supported solid oxide fuel cell, Comput. Chem. Eng. 32 (10) (2008) 2365–2381.
[15] W. Frenkel, J. Kersten, R. Husmann, H. Aschemann, Design of Robust PID Controllers for SOFC Stacks, in: Proceedings of the 2022 IEEE Conference on Control Technology and Applications, CCTA, Trieste, Italy, 2022, pp. 510–515.
[16] A. Pohjoranta, M. Halinen, J. Pennanen, J. Kiviaho, Model predictive control of the solid oxide fuel cell stack temperature with models based on experimental data, J. Power Sources 277 (2015) 239–250.
[17] M. Fardadi, D.F. McLarty, F. Jabbari, Actuator limitations in spatial temperature control of SOFC, J. Fuel Cell Sci. Technol. 10 (3) (2013) 031005.
[18] F. Mueller, F. Jabbari, R. Gaynor, J. Brouwer, Novel solid oxide fuel cell system controller for rapid load following, J. Power Sources 172 (1) (2007) 308–323.
[19] C. Stiller, B. Thorud, O. Bolland, R. Kandepu, L. Imsland, Control strategy for a solid oxide fuel cell and gas turbine hybrid system, J. Power Sources 158 (1) (2006) 303–315.
[20] B.J. Spivey, T.F. Edgar, Dynamic modeling, simulation, and MIMO predictive control of a tubular solid oxide fuel cell, J. Process Control 22 (8) (2012) 1502–1520.
[21] B.M. Sanandaji, T.L. Vincent, A.M. Colclasure, R.J. Kee, Modeling and control of tubular solid-oxide fuel cell systems: II. Nonlinear model reduction and model predictive control, J. Power Sources 196 (1) (2011) 208–217.
[22] X. Zhang, S. Chan, H.K. Ho, J. Li, G. Li, Z. Feng, Nonlinear model predictive control based on the moving horizon state estimation for the solid oxide fuel cell, Int. J. Hydrog. Energy 33 (9) (2008) 2355–2366.
[23] M.H. de Lange, P. Segovia, R.R. Negenborn, L. van Biert, Model predictive control for thermal stress-aware power modulation of solid oxide fuel cell systems, Energy Convers. Manage. 347 (2026) 120395.
[24] H. Jie, J. Liao, G. Zhu, W. Hong, Nonlinear model predictive control of direct internal reforming solid oxide fuel cells via PDAE-constrained dynamic optimization, Appl. Energy 360 (2024) 122804.
[25] L. Xia, M. Zhang, M. Han, L. Sun, Model predictive control of a one-dimensional methane reforming solid oxide fuel cell system based on Modelica/SIMULINK co-simulation, Renew. Energy 256 (2026) 123954.
[26] M. Fardadi, D.F. McLarty, F. Jabbari, Controlling spatial temperature variation in a rapid load following SOFC, in: Proceedings of the International Conference on Fuel Cell Science, Engineering and Technology, vol. 55522, American Society of Mechanical Engineers, 2013, V001T02A006.
[27] A. Nakajo, Z. Wuillemin, J. Van herle, D. Favrat, Simulation of thermal stresses in anode-supported solid oxide fuel cell stacks. Part I: Probability of failure of the cells, J. Power Sources 193 (1) (2009) 203–215, Scientific Advances in Fuel Cell Systems.
[28] L. Van Biert, M. Godjevac, K. Visser, P. Aravind, Dynamic modelling of a direct internal reforming solid oxide fuel cell stack based on single cell experiments, Appl. Energy 250 (2019) 976–990.
[29] D. Marra, C. Pianese, P. Polverino, M. Sorrentino, Models for Solid Oxide Fuel Cell Systems, Springer, 2016.
[30] C. Stiller, Design, operation and control modelling of SOFC/GT hybrid systems (Ph.D. thesis), Norwegian University of Science and Technology, 2006.
[31] L.T. Biegler, Nonlinear programming: concepts, algorithms, and applications to chemical processes, SIAM, 2010.
[32] S.G. (D), Technical Documentation Integrated Stack Module (ISM), Tech. Rep..
[33] J.A.E. Andersson, J. Gillis, G. Horn, J.B. Rawlings, M. Diehl, CasADi – A software framework for nonlinear optimization and optimal control, Math. Program. Comput. 11 (1) (2019) 1–36.
[34] L. Vanroye, A. Sathya, J. De Schutter, W. Decré, Fatrop: A fast constrained optimal control problem solver for robot trajectory optimization and control, in: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, IEEE, Detroit, USA, 2023, pp. 10036–10043.
[35] P. Krupa, D. Limon, T. Alamo, Implementation of Model Predictive Control in Programmable Logic Controllers, IEEE Trans. Control Syst. Technol. 29 (3) (2021) 1117–1130.
[36] M. Diehl, H.G. Bock, J.P. Schlöder, A Real-Time Iteration Scheme for Nonlinear Optimization in Optimal Feedback Control, SIAM J. Control Optim. 43 (5) (2005) 1714–1736.