PROCEEDINGS OF ECOS 2023 - THE 36TH INTERNATIONAL CONFERENCE ON
EFFICIENCY, COST, OPTIMIZATION, SIMULATION AND ENVIRONMENTAL IMPACT OF ENERGY SYSTEMS
25-30 JUNE 2023, LAS PALMAS DE GRAN CANARIA, SPAIN

Real-time data assimilation for the thermodynamic modeling of a cryogenic fuel tank

Pedro Afonso Marques[a,b], Samuel Ahizi[b], Miguel Alfonso Mendez[b]

[a] *von Karman Institute for Fluid Dynamics, Sint-Genesius-Rhode, Belgium, pedro.marques@vki.ac.be, CA Transfers, Interfaces & Processes (TIPs) laboratory, Université libre de Bruxelles, Brussels, Belgium*
[b] *von Karman Institute for Fluid Dynamics, Sint-Genesius-Rhode, Belgium,*

Abstract:

The current energy crisis has accelerated the interest in sustainable energy sources and cryogenic propellants, liquid hydrogen (LH_2) or liquefied natural gas (LNG). The thermal management of these systems requires advanced control strategies to minimize the boil-off losses produced by heat leakages and sloshing-enhanced heat and mass transfer. This work presents a data-assimilation approach to calibrate a 0D thermodynamic model for cryogenic fuel tanks from data collected in real time from multiple tanks. The proposed approach combines ideas from traditional data assimilation with multi-environment reinforcement learning, where an agent's training (i.e., model assimilation) is carried out simultaneously on multiple systems (environments). The real-time assimilation uses a mini-batch version of the well-known Limited-memory Broyden–Fletcher–Goldfarb–Shanno with bounds (L-BFGS-B) and adjoint-based gradient computation for solving the underlying optimization problem. The approach is tested on synthetic datasets simulating multiple tanks undergoing different operation phases (pressurization, hold, long-term storage and sloshing). We present the general closure problem for the tank's thermodynamics and the multi-environment assimilation framework. The results show that the simultaneous sampling from multiple environments and loading scenarios significantly accelerates the assimilation. Moreover, it is shown that the time scale for the training process is shorter than the relevant time scales of the tank's thermodynamics. This opens the path towards model predictive control based on data-driven engineering models.

Keywords:

Thermodynamics; Cryogenics; Sloshing; Modeling; Machine learning; Data assimilation.

1. Introduction

Cryogenic fuels such as liquid hydrogen (LH_2) or liquefied natural gas (LNG) are managed and stored at extremely low temperatures (typically $\approx$-170 °C for LNG and $\approx$-250 °C LH_2). LH_2 has been mainly used in rocket engines [1] but is now actively explored as alternatives to fossil fuels in many applications, including naval [2] and aeronautical [3] industries. The need for cryogenic temperatures stems from the need to maximize the fuel density without resorting to extreme operating pressures ($>$ 300 bars for gaseous H_2) [4]. In addition, the higher thrust-to-weight ratio compared to classical propulsive solutions [4] and the absence of pollutant emission make LH_2 a promising energy carrier for a carbon-neutral future [3, 5]. LNG is not entirely emission-free but can still bridge the gap between traditional fossil fuels and a fully renewable energy system [3].

Nevertheless, storage at cryogenic temperatures requires a complex thermal management system, which poses significant challenges to any application requiring long holding times, from marine to aeronautical, from ground transportation to deep space exploration. No insulating system can entirely prevent heat exchanges with the surroundings. Thus some of the fuel unavoidably evaporates over time and increases the tank pressure [6]. Additional challenges are faced in tanks installed on vehicles, as the external accelerations induce sloshing. Sloshing, defined as the movement of the free liquid surface, can (suddenly) increase heat and mass transfer rates between the liquid and ullage gasses and thus produce significant variations of the tank's pressure [7, 8].

This work proposes a data assimilation framework for identifying closure laws in a 0D thermodynamic model from real-time measurements of pressure, temperature, and liquid level in cryogenic tanks. The underlying optimization problem is efficiently solved via gradient-based optimization using the L-BFGS-B (Limited-memory Broyden–Fletcher–Goldfarb–Shanno with bounds) algorithm [9] coupled with the adjoint method [10] for evaluating the cost function gradient. The main novelty in our approach lies in the simultaneous assimilation of observations from multiple fuel tanks (i.e., multiple environments) undergoing distinct thermodynamic evolution in potentially widely different operating conditions.

The proposed approach takes inspiration from multi-agent reinforcement learning (MARL) [11]. MARL is a subfield of reinforcement learning where multiple agents learn to interact with each other and with the environment

https://doi.org/10.52202/069564-0095

to maximize a global reward. This is an active area of research with many open questions and challenges [12] on handling communication and coordination between agents and learning effective policies in large-scale multi-agent systems. Nevertheless, we believe that the potential impact of this concept in data assimilation is enormous and worth investigating. This work moves the first steps towards its application to the problem of data-driven calibration of thermodynamic models for cryogenic tanks. Leveraging synthetic data, we aim to quantify the performance of the model's calibration as the number of *environments* increases, using a single *agent* to learn the closure terms. As a complementary investigation, the single-environment scenario is also investigated to assess the impact of noise on the training data and the uniqueness of the solution.

This article is structured as follows. Section 2. outlines the physical modeling of the cryogenic fuel tank implemented in this work. The real-time data assimilation and inverse method strategies are outlined in Section 3.. Furthermore, this section also describes how the synthetic test cases were generated. Lastly, Section 4. overviews the results of the model calibration in single-environment and multi-environment conditions. Concluding remarks and future outlook are discussed in Section 5..

2. Thermodynamic modeling

The cryogenic propellant tank considered in this work is a single-species system composed of a liquid and its vapor enclosed in an insulated tank Fig. 1 provides a schematic of the problem with the relevant parameters involved. Subscripts l, v, and w are used to distinguish variables related to the liquid, the vapor, and the wall, respectively. The gas-liquid interface separating vapor and liquid is treated as an infinitesimally thin region where heat and mass transfer occurs. The system exchanges heat and mass, both on the liquid and the vapor side, during various operations (e.g., pressurization, venting, filling). The reader is referred to the list of symbols at the end of the article for the nomenclature.

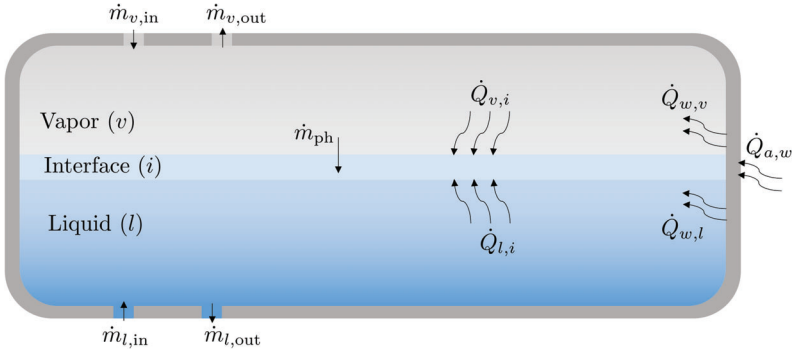

Figure 1: Schematic of the cryogenic fuel tank and its subsystems: vapor (v), liquid (l), and insulating walls (w). The heat and mass exchanges between these control volumes are expressed through the $\dot{m}$ and $\dot{Q}$ fluxes.

The 0D thermodynamic model used in the assimilation expresses the conservation of mass and energy applied to three control volumes: the liquid, the vapor, and the solid. Considering mass-averaged thermodynamic properties, these balances result in a system of ordinary differential equations (ODEs), which must be closed with empirical relations for the heat and mass transfer rates. Finding closure from sampled data is the objective of the assimilation.

The mass conservation between the liquid and the vapor is given by

$$\frac{dm_v}{dt} = \sum \dot{m}_{v,\text{in}} - \sum \dot{m}_{v,\text{out}} - \dot{m}_{\text{ph}} \quad ; \quad \frac{dm_l}{dt} = \sum \dot{m}_{l,\text{in}} - \sum \dot{m}_{l,\text{out}} + \dot{m}_{\text{ph}} \,, \tag{1}$$

where $\dot{m}_{\text{ph}} = \dot{m}_{\text{cond}} - \dot{m}_{\text{evap}}$ is the net mass flux through the interface, accounting for the balance of condensation and evaporation. The mass conservation in the vapor and liquid phases give

$$\frac{dU_v}{dt} = \sum \dot{m}_{v,\text{in}} \hbar_{v,\text{in}} - \sum \dot{m}_{v,\text{out}} \hbar_{v,\text{out}} - \dot{m}_{\text{ph}} \hbar_{v,\text{sat}} - \dot{Q}_{v,i} + \dot{Q}_{w,v} + \dot{W}_v \tag{2}$$

$$\frac{dU_l}{dt} = \sum \dot{m}_{l,\text{in}} \hbar_{l,\text{in}} - \sum \dot{m}_{l,\text{out}} \hbar_{l,\text{out}} + \dot{m}_{\text{ph}} \hbar_{v,\text{sat}} - \dot{Q}_{l,i} + \dot{Q}_{w,l} + \dot{W}_l, \tag{3}$$

where U_v and U_l are the internal energies of the vapor and liquid phases, $\dot{Q}$ denotes the general heat flux, $\hbar$ is the specific enthalpy, and $\dot{W} = -pdV/dt$ is the expansion/compression work due to changes in filling level.

The model closure is required to link the heat transfer rates at the interface. Assuming that the heat transfer occurs at much larger time scales than the interface dynamics, we consider quasi-steady formulation and use Newton's cooling law for the closure relation:

$$\dot{Q}_{v,i} = A_i h_{v,i}(T_v - T_i) \;\; ; \;\; \dot{Q}_{w,v} = A_{w,v} h_{w,v}(T_w - T_v) \;\; ; \;\; \dot{Q}_{l,i} = A_i h_{l,i}(T_l - T_i) \;\; ; \;\; \dot{Q}_{w,l} = A_{w,l} h_{w,l}(T_w - T_l), \tag{4}$$

where $h_{v,i}$, $h_{l,i}$, $h_{w,v}$, $h_{w,l}$ are the heat transfer coefficients, A_i is the gas-liquid interface area, $A_{w,v}$ and $A_{w,l}$ are the surface exchange areas between the walls-vapor and walls-liquid phases. The heat transfer coefficients are unknown and must be identified from the data. The interface temperature in (4) is assumed to be the saturation temperature evaluated at vapor pressure p_v; hence the energy balance at the interface provides the mass flux due to phase change as $\dot{m}_{\text{ph}} = (\dot{Q}_{l,i} - \dot{Q}_{v,i})/\mathcal{L}_v$, with $\mathcal{L}_v$ the latent heat of vaporization.

Concerning the exchange areas in (4), we do not account for the time variation of A_i (due to, e.g., sloshing), and take it as the tank's cross-section when the tank is half-filled. On the other hand, given the tank's geometry, the areas $A_{w,v}$ and $A_{w,l}$ are updated at each time step depending on the liquid level. This can be computed from the liquid and vapor masses and their properties. In particular, treating the liquid phase as incompressible (as in [13]), its density is solely a function of temperature $\rho_l(p, T) \approx \rho(T)$, hence

$$\frac{dV_l}{dt} = -\frac{dV_v}{dt}. \approx \frac{1}{\rho_l}\frac{dm_l}{dt} - \frac{m_l}{c_{p,l}\rho_l^2}\left(\frac{d\rho_l}{dT}\right)_p \frac{du_l}{dt}, \tag{5}$$

where c_p is the specific heat at constant pressure, ρ is the density, and $u = U/m$ is the specific internal energy. Finally, in the 0D formulation, the insulating walls are treated as a single control volume which can exchange heat with the vapor and liquid phases, as well as with the external environment at ambient conditions. For a single control volume with mass-averaged properties, the internal energy of the walls is described through

$$\frac{dU_w}{dt} = \dot{Q}_{a,w} - \dot{Q}_{w,v} - \dot{Q}_{w,l} \tag{6}$$

where $\dot{Q}_{a,w}$ is the heat entering the tank from the environment. This term could be computed from existing literature correlations [14], but in this work, this is taken as a user-defined function that depends on the specific scenario experienced by the tank (described in the following section). It is worth noticing that (6) can be written as a function of the mass-averaged solid temperature by introducing $dU_w = m_w c_w dT_w$.

Finally, all thermodynamic properties (c_p, ρ), as well as the link between internal energies, temperatures and pressures, were evaluated using the CoolProp package [15] in Python. This library implements pure fluid equations of state and transport properties using the Helmholtz energy formulations.

To conclude this section, we note that the thermodynamic model is constituted of equations (1)-(6) and can be cast in the form of a parametric initial value problem:

$$\begin{cases} \dfrac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t; \boldsymbol{\theta}) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases} \tag{7}$$

where $\mathbf{x} = [m_v, m_l, u_v, u_l, T_w, V_v, V_l] \in \mathbb{R}^7$ is the state vector (composed by the masses, volumes and internal energies of each control volume), describing the thermodynamic condition of the system at time-instant t, $\mathbf{x}_0 \in \mathbb{R}^7$ is the (known) initial condition and $\boldsymbol{\theta} \in \mathbb{R}^p$ is the vector of unknown parameters. These identify a function mapping the state vector to the heat transfer coefficients in (4) as further detailed in Section 3.1..

3. Real-time data assimilation framework

We describe the generation of the synthetic database in section 3.1. and the multi-environment real-time assimilation approach in section 3.2..

3.1. Synthetic database generation

Borrowing from the reinforcement learning terminology, the assimilation is carried out by an *agent* interacting with an *environment* to achieve a goal. In our context, the *agent* is a function that must predict the closure parameters (heat transfer coefficients) with the goal of having the model prediction as close as possible to the available data. In our context, the *environment* definition requires defining (1) the geometry of the fuel tanks, (2) the working fluid, (3) the initial conditions of the system, (4) the temporal evolution of the heat transfer coefficients (unknown to the agent), and (5) the sequence of operations applied to the tank. These are summarised in Fig. 2.

The proposed multi-environment formulation involves learning the *same* closure law from multiple environments. In this work, we consider multiple cryogenic fuel tanks, partly filled with LH_2. The environment parameters in terms of dimensions and initial conditions are presented in Table 4 in Appendix A. Tank 1 is the

https://doi.org/10.52202/069564-0095

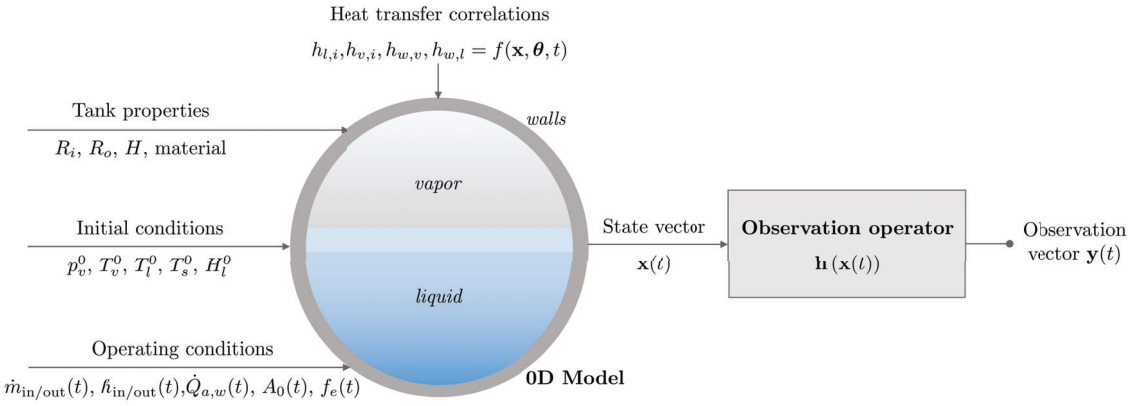

Figure 2: Schematic of the parameters required to define an environment for the data assimilation of the 0D thermodynamic model of the cryogenic tank. Interaction with the environment produces the observation vector $\mathbf{y}(t)$ simulating measurements from the tank.

reference case used to assess the single-environment performance of the assimilation framework (Section 4.1.). The remaining tanks are added to the analysis of the multi-environment scenario (Section 4.2.).

In all environments, the heat transfer coefficients in (4) were modeled as

$$h_{v,i} = \frac{k_v}{R}\left(\theta_0 \mathrm{Re}_{s,v}^{0.69}\mathrm{Pr}_v^{1/3} + \theta_1 \mathrm{Ra}_v^{0.15}\right) \;;\; h_{l,i} = \frac{k_l}{R}\left(\theta_2 \mathrm{Re}_{s,l}^{0.69}\mathrm{Pr}_v^{1/3} + \theta_3 \mathrm{Ra}_l^{0.15}\right) \;;\; h_{w,v} = \theta_4 h_{v,i} \;;\; h_{w,l} = \theta_5 h_{l,i} \tag{8}$$

where Re_s is the Reynolds number due to sloshing, $\mathrm{Pr} = \nu/\alpha$ is the Prandtl number, Ra is the Rayleigh number, and $\theta_0 - \theta_5$ are the set of parameters governing the heat and mass transfer evolution in the tank.

The dimensionless numbers in (8) are defined as

$$\mathrm{Re}_s = \frac{f_e}{f_{11}}\left(\frac{b}{R}\right)^2 \frac{\left(gR^3\right)^{1/2}}{\nu}\sqrt{1.841} \quad ; \quad \mathrm{Ra} = \frac{g\beta\Delta T R^3}{\nu\alpha} \tag{9}$$

where f_e is the frequency of the forcing motion acting on the tank, f_{11} is the natural frequency of the tank [1], b is the maximum expected wave-height during sloshing, which is a function of the forcing amplitude A_0 [8], R is the tank's radius, ν is the kinematic viscosity, g is the gravitational acceleration, β is the volumetric thermal expansion coefficient, and α is thermal diffusivity.

The $h_{l,i}$ and $h_{v,i}$ coefficients simultaneously account for forced convection through the Reynolds and Prandtl numbers, as well as buoyancy-driven fluxes through the Rayleigh number. All synthetic test cases were generated using the set of parameters $\boldsymbol{\theta} = [\theta_0, \theta_1 \ldots]$ presented in Table 3 in Appendix A. These were tuned to roughly portray the experimental pressure and temperature reported in [7, 16].

Assuming that the functional relations in (8) are prescribed beforehand, the scope of data assimilation consists in identifying the parameters $\boldsymbol{\theta}$ from multiple environments. A general formulation of the problem could be based on general parametric models (e.g. using radial basis functions expansions or artificial neural networks) to learn the closure laws. These are currently being explored and will be presented in future work.

The interaction between the agent and the environment is carried out by monitoring some observations of the system, here denoted as $\mathbf{y}(t) = \mathbf{h}(\mathbf{x}(t))$, with $\mathbf{h}(\cdot)$ the observation function simulating a measurement process. The observations considered in this work are the vapor pressure $p_v(t)$, the mass-averaged vapor $T_v(t)$, liquid $T_l(t)$, and wall $T_w(t)$ temperatures, and the fill-level $H_l(t)$. These quantities were retrieved from the state vector. Hence the observation function is $\mathbf{h}(\cdot) : \mathbb{R}^7 \to \mathbb{R}^5$.

3.2. Real-time inverse method for multiple environments

Identifying the parameters $\boldsymbol{\theta}$ requires solving an optimization problem. The optimal parameters are those that minimize a cost function $\mathcal{J}(\boldsymbol{\theta})$ measuring the discrepancy between model prediction and data across all the available environments. Denoting as $\mathbf{x}_{\{j\}}(t)$ the time evolution of the thermodynamic state in the j^{th} environment and as $\mathbf{y}_{\{j\}}(t)$ the evolution of the corresponding observations, the cost function considered in this work is

$$\mathcal{J}(\boldsymbol{\theta}) = \frac{1}{N}\sum_{j=1}^{N}\int_{t_i}^{t_i+T} g(\mathbf{y}_{\{j\}}, \mathbf{x}_{\{j\}}, t; \boldsymbol{\theta})dt = \frac{1}{2N}\sum_{j=1}^{N}\int_{t_i}^{t_i+T}\left(\mathbf{y}_{\{j\}}(t) - \mathbf{h}(\mathbf{x}_{\{j\}}(t))\right)^T \mathbf{R}_{\{j\}}^{-1}\left(\mathbf{y}_{\{j\}}(t) - \mathbf{h}(\mathbf{x}_{\{j\}}(t))\right) dt \tag{10}$$

where N is the number of environments, T is the observation time and $\mathbf{R}_{\{j\}}$ is the covariance matrix accounting for measurement noise. Since the assimilation seeks to learn the full set of parameters from all environments, the proposed formulation is a multiple-environment but single-agent framework, as opposed to a multi-agent formulation in which different agents could be assigned to learn different coefficients in the same environment.

In this work, the function $\mathcal{J}(\boldsymbol{\theta})$ is minimized through gradient-based optimization using the L-BFGS-B optimizer [9] coupled with the adjoint method [10] to compute the gradient $\nabla_{\boldsymbol{\theta}}\mathcal{J}$. The adjoint method allows computing this gradient without computing the sensitivity of the state with respect to the closure parameters, i.e. $d\boldsymbol{x}/d\boldsymbol{\theta}$.

The gradient is computed by relying on the augmented Lagrangian function $\mathcal{L}(\boldsymbol{\theta})$

$$\mathcal{L}(\boldsymbol{\theta}) = \mathcal{J}(\boldsymbol{\theta}) + \sum_{j=1}^{N} \left(\int_{t_j}^{t_j+T} \left(\mathbf{f}(\mathbf{x}_{\{j\}}, t; \boldsymbol{\theta}) - \frac{d\mathbf{x}_{\{j\}}}{dt} \right) \boldsymbol{\lambda}_{\{j\}}(t)^T dt \right) \tag{11}$$

where the (row) vector $\boldsymbol{\lambda}_{\{j\}}(t) \in \mathbb{R}^7$ is the vector of adjoint variables for each environment. The cost function (10) and the augmented cost function (11) are equivalent because the additional term is identically null for any finite choice of the variable $\boldsymbol{\lambda}_{\{j\}}$ (by definition of the underlying dynamics in (7)). Therefore, it is possible to make the gradient computation of both cost functions independent of the sensitivities if the adjoint variables are taken as the solution of the following terminal value problem [17]:

$$\begin{cases} \dfrac{d\boldsymbol{\lambda}_{\{j\}}}{dt} = -\left(\dfrac{\partial g}{\partial \mathbf{x}}\right)_{\{j\}} - \boldsymbol{\lambda}_{\{j\}}(t) \left(\dfrac{\partial \mathbf{f}}{\partial \mathbf{x}}\right)_{\{j\}} \\ \boldsymbol{\lambda}_{\{j\}}(T) = 0 \end{cases} \tag{12}$$

This is a linear system of equations that must be integrated backward in time to solve for $\boldsymbol{\lambda}_{\{j\}}$. The gradient can then be computed as

$$\nabla_{\boldsymbol{\theta}}\mathcal{J}(\boldsymbol{\theta}) = \sum_{j=1}^{N} \left(\int_{t_j}^{t_j+T} \left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)_{\{j\}} + \boldsymbol{\lambda}_{\{j\}}(t) \left(\frac{\partial \mathbf{f}}{\partial \boldsymbol{\theta}}\right)_{\{j\}} dt + \boldsymbol{\lambda}_{\{j\}}(0) \frac{d\mathbf{x}_{\{j\}}}{d\boldsymbol{\theta}}(0) \right) \tag{13}$$

where $d\mathbf{x}_{\{j\}}/d\boldsymbol{\theta}(0)$ is known from the initial conditions of each environment (tank). Thus, in the adjoint-based approach, the gradient of the loss function is evaluated by solving two systems of equations for each environment in the ensemble (i.e., one forward pass to obtain $\mathbf{x}_{\{j\}}(t)$, and one backward pass to obtain $\boldsymbol{\lambda}_{\{j\}}(t)$).

The vector of parameters is iteratively updated throughout the optimization loop as

$$\boldsymbol{\theta}^{(k+1)} = \boldsymbol{\theta}^{(k)} - \mathbf{B}^{(k)} \nabla_{\boldsymbol{\theta}}\mathcal{J}(\boldsymbol{\theta}^{(k)}) \text{ for } k \in [0, 1 \dots n_k] \tag{14}$$

where the superscript $^{(k)}$ indicates the iteration counter, and $\mathbf{B}$ is the approximation of the Hessian matrix inverse according to [9]. The gradient computation in (13) and the updates in (14) define two-time scales of the assimilation problem. The first scale is the observation time T, linked to the rate at which new information is collected. The second is the learning time scale implicitly defined by the number of optimization iterations n_k carried out before an update on the gradient is requested (or is available). The first time scale defines the rate at which the cost function changes because of the dynamics of the system and the potential occurrence of unseen scenarios. The second time scale defines the rate at which the optimizer travels along the parameter space before the cost function changes.

The optimal setting of these scales poses a fundamental question on the impact of the observation time T in relation to the observed scenarios and the quality of the gradient computation: one might ask, for example, whether the observation time should be long enough to observe at least two or three pressurizations or sloshing events. This work investigates the impact of this parameter by making it vary from T = 10 min until T = 60 min, with the five environments observing different scenarios and different combinations of events.

4. Results and discussion

We split the presentation of results and the discussion into a section dedicated to the performance of a single environment (Section 4.1.) and a section dedicated to multiple environments (Section 4.2.).

4.1. Single-environment performance

We first describe the investigated scenario in 4.1.1.. Section 4.1.2. studies the impact of measurement noise in the collected data for different observation times T. Finally, section 4.1.3. reports on the impact of mini-batching the assimilation using only a portion of the data.

https://doi.org/10.52202/069564-0095

4.1.1. Test case description

We consider the environment denoted as 'case 1' in Table 4. Figure 3 illustrates the loading and operation scenario of the tank in an observation of 1 hour. The first row of plots illustrates the time evolution of the ullage gas pressure (on the left), and the mass averaged temperatures of the vapor, the liquid and the solid volumes (on the right). The second row plots the time amplitude (left) and the frequency (right) of sinusoidal sloshing events. The third row plots the mass inflow/outflow of vapor (left) and liquid (right) as a function of time. Finally, the last row shows the heat flux exchanged through the walls (left) and the liquid level (right).

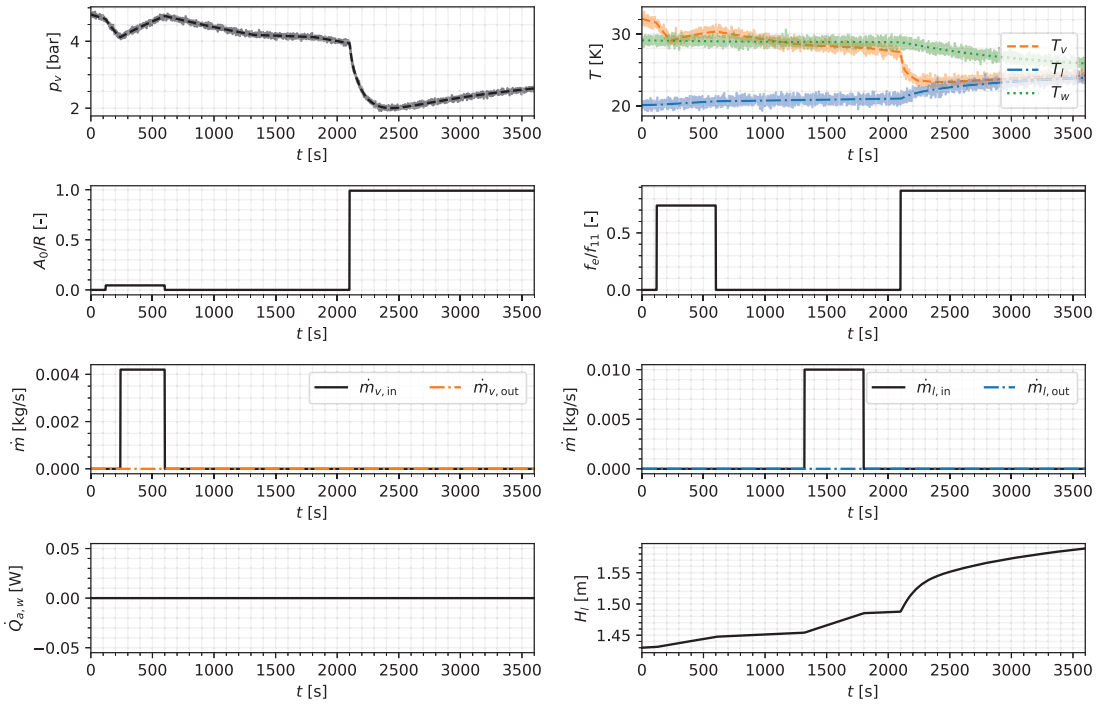

Figure 3: Synthetic data generated for Tank 1 (see Table 4) used to evaluate the single-environment performance. The components of the observation vector are displayed alongside the injected/removed mass fluxes, the non-dimensional excitation conditions, and the external heat flux.

A detailed overview of these graphs helps understand the complexity of the system control problem. In the presented scenario, the tank is undisturbed in the first 2 minutes. Then, a moderate sloshing event occurs with a dimensionless amplitude of $A_0/R = 0.045$ and dimensionless frequency $f_e/f_{11} = 0.8$ (see definitions of sloshing conditions in (9)). This triggers a visible pressure drop. To counter-balance this, at time $t = 4$ minutes, hydrogen vapor at 40 K and 2 bar is injected in the ullage at a rate of 4.2 g/s for 6 minutes. This allows the tank to recover 4.8 bar pressure at 10 minutes when the sloshing event ends. From $t = 10$ to $t = 22$ minutes, the tank is again undisturbed but the temperature difference between the wall, the gas, and the liquid results in the warming of the liquid and the cooling of the vapor. This slightly reduces the pressure in the ullage. Finally, between $t = 22$ and $t = 30$ minutes, liquid at 22 K is injected into the system at a rate of 0.01 kg/s. This reduces the ullage volume, resulting in a moderate compression that arrests the decreasing trend. Finally, at $t = 35$ minutes, a violent sloshing event occurs and continues until the end of the observation. This results in a violent pressure drop due to the significant condensation (see liquid level evolution), followed by a moderate pressure rise due to heat exchanges with the (now) warmer walls. Throughout this test, no heat exchange is assumed to occur from the environment.

4.1.2. Impact of noise in the observation data

We here consider three scenarios in terms of measurement noise in the collected observations; these are termed (1) 'clean' (2) 'low-noise', and (3) 'high-noise'. No noise is present in the first, while Gaussian noise (with zero average) is added in the other two. In (2), the noise's standard deviation is 2 kPa on the pressure, 0.2 K on the temperature, and 0.5 mm on the fill level. In (3), these values are doubled.

We considered various observation windows T of the environment described in the previous section for the

three scenarios. Figure 4 shows the results of the real-time data assimilation applied to the 'clean' data for T = 10, 20, and 30 minutes. The thick solid lines indicate observations from the virtual experiment, whereas the thin dashed lines indicate the predictions given by data assimilation. The columns of the figure show the tank pressure, vapor temperature, and liquid temperature as the observation time increases.

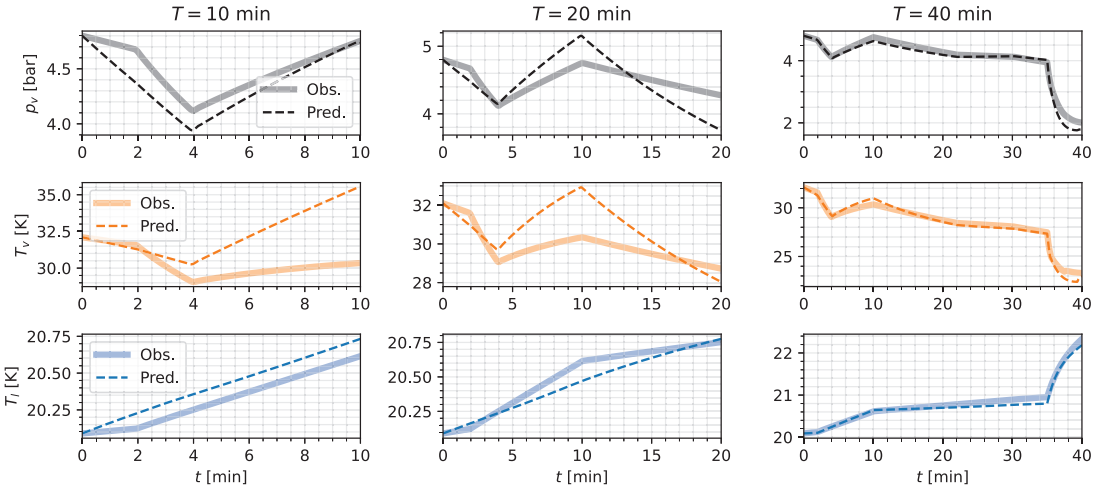

Figure 4: Model predictions and observation data extracted from the test described in Section 4.1.1. for 10, 20 and 40 minutes of observation.

Interestingly, the assimilation produces excellent matching between the model and observation if the observation time is sufficiently long. In the virtual experiment analyzed, this is of the order of T = 40 min, i.e., enough to see all the events in the environment (sloshing, filling on the vapor and the liquid side).

The analysis of the noise impact is thus completed in Fig. 5, which presents the minimum of $\mathcal{J}$ (on the left) and the l_2 relative error norm of $\boldsymbol{\theta}$ (on the right) as a function of the observation time, for the three tested scenarios as a function of the observation time. When noise is added, the minimal cost function is expected to be larger in the presence of noise since the underlying model filters it out from the data. Nevertheless, these results show that the noisy cases consistently produced smaller values of $\mathcal{J}$ when the full observation is used. Surprisingly, the measurement noise helps the optimization avoid local minima, and the parameters identified in the assimilation are more accurate (right plot in Fig. 5) and learned with shorter observation times. Interestingly, in all cases, the learning of the parameters saturates for $T > 40$ min.

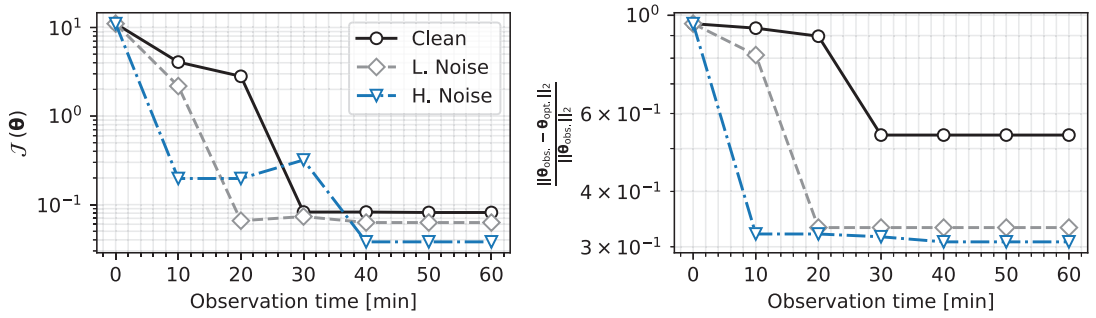

Figure 5: Performance of the data assimilation framework applied to one environment for different noise levels in the data. Left: minimum of the cost function as a function of the observation time. Right: l_2 relative error norm associated with $\boldsymbol{\theta}$.

We close this analysis by comparing the retrieved parameters from the assimilation and those used to generate the data; these are reported in the next section and presented alongside the impact of data batching for the gradient computation.

https://doi.org/10.52202/069564-0095

4.1.3. Impact of approximating the gradient of the cost function via mini-batches

The mini-batch sampling [18] is a classic approach to escape local minima and limit memory requirements in the gradient descent method. The idea consists in computing the cost function gradient using a randomly chosen subset of the observation data; this produces an approximation of the gradient $\nabla_\theta \mathcal{J}$ that might point away from local minima.

The use of mini-batch strategies on quasi-Newton methods, such as the L-BFGS-B technique used in this work, is the subject of active research (see [19]). Nevertheless, in this work, we explored its impact on the 'high-noise' configuration from the previous subsection. We test the assimilation using 80%, 50%, 20% and 5% of the observed data to evaluate $\nabla_{\boldsymbol{\theta}} \mathcal{J}$. To account for the stochasticity of the process, the assimilation is repeated one hundred times in each case and the results are averaged. The average behavior of the cost function evolution and the convergence of the parameters are shown in Fig. 6. Remarkably, the additional stochasticity of the minibatch selection further improves the result of the assimilation in both aspects. Interestingly, the relative error is nearly identical when sampling 80%, 50%, and 5% of the training data, whereas for the 20% bath size the error is reduced by roughly 6%. The cost function, on the other hand, is slightly higher.

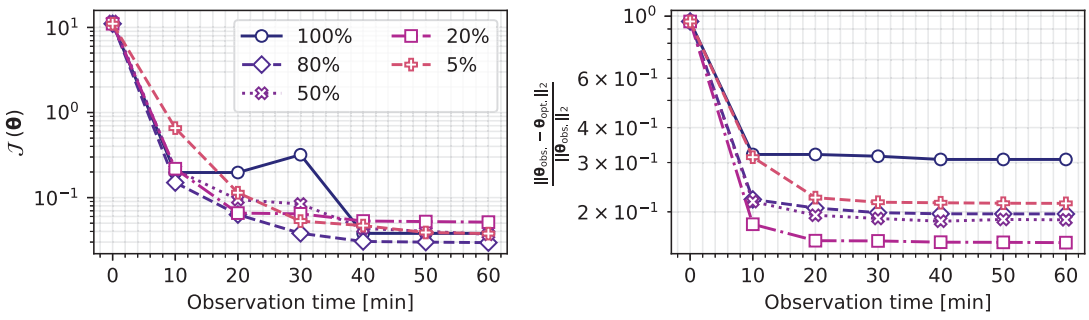

Figure 6: Performance of the data assimilation framework applied to one environment. Left: minimum of the cost function as a function of the observation time. Right: l_2 relative error norm associated with $\boldsymbol{\theta}$.

Finally, we compare the parameters obtained in the assimilation for different mini-batch sizes in table 1. For each coefficient, the table collects the average and the standard deviation (over 100 repetitions of assimilation). The reference coefficients used to generate the synthetic data are shown in the last row of the table. These results confirm that the best match of $\boldsymbol{\theta}$ was achieved when 20% of the training data was sampled at each gradient computation, even though this case reported the highest cost function evaluation on average.

Table 1: Mean coefficients obtained during the model training, alongside their respective standard deviation. The cases in which $\nabla_{\boldsymbol{\theta}} \mathcal{J}$ was approximated via mini-batches were repeated 100 times.

Data sampled	min $\mathcal{J}$	θ_0	θ_1	θ_2	θ_3	θ_4	θ_5
100%	3.8e-2	98.3	17.2	65.2	1.21	8.91	11.6
80%	2.7e-2	134.1 ± 14.6	10.7 ± 2.4	51.6 ± 4.7	1.9 ± 0.2	34.9 ± 37.5	10.1 ± 0.4
50%	2.8e-2	128.3 ± 16.4	11.1 ± 2.9	50.9 ± 3.8	1.9 ± 0.3	28.8 ± 32.0	10.2 ± 0.7
20%	3.4e-2	134.5 ± 16.3	10.0 ± 3.6	49.6 ± 5.5	2.0 ± 0.3	18.1 ± 19.2	10.0 ± 0.6
5%	3.0e-2	120.0 ± 23.8	12.8 ± 4.2	53.3 ± 7.2	1.8 ± 0.4	18.7 ± 16.6	10.0 ± 0.6
Reference	-	140.0	10.0	50.0	2.00	20.0	10.0

The table shows that the standard deviation for each component of $\boldsymbol{\theta}$ increases as the batch size decreases. This is attributed to less accurate gradient evaluation during the parameter updates. The only outlier to this trend is θ_4, which shows the opposite behavior. However, this is characterized by a significantly larger standard deviation. This parameter is associated with the heat transfer coefficient $h_{w,v}$, governing the heat exchanges between the vapor and the insulating walls, which appears to have a minor impact on the system's dynamics. Thus, given the insensitivity of the model to this parameter, the assimilation cannot determine its true value from the observed data.

In summary, these results showcased that sampling portions of the training data during the optimization can improve the assimilation results. Furthermore, this sampling provides an approximation of $\nabla_{\boldsymbol{\theta}} \mathcal{J}$, which allows for a broader exploration of the parameter space, yielding better predictions for $\boldsymbol{\theta}$, even if the value of the cost

function is not significantly affected (on average). For the current training data set, the best compromise was obtained by sampling 20% of the input data.

4.2. Multi-environment performance

We here extend the previous analysis to the case of multiple environments characterised by vastly different geometries, sloshing profiles, thermal loading and inflow/outflow control actions.

The first environment is the tank described in Section 4.1.1. while the geometrical parameters, and initial conditions for the others are summarised in Table 4 in Appendix A. All environments are observed for the same T = 60min, but their thermodynamic evolutions differ by imposing different sequences of filling, pressurization, venting, sloshing, and external heating.

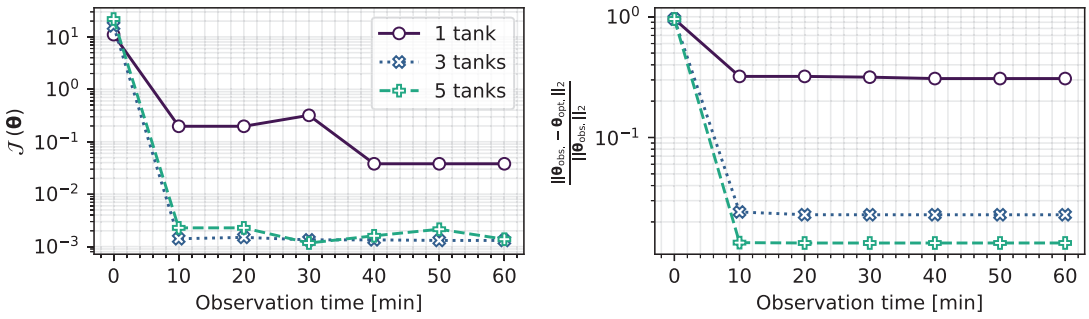

Figure 7: Performance of the data assimilation framework applied to multiple environments. Left: minimum of the cost function as a function of the observation time. Right: l_2 relative error norm associated with $\boldsymbol{\theta}$.

The multi-environment framework requires solving for N forward and backward passes for every parameter update. However, these evaluations could be performed in parallel to improve runtimes and could be batched also across multiple environments. We repeat the same analysis as in the previous section, applying high noise to the observed data and using a batch of 20% of the observation points to update $\boldsymbol{\theta}$ during the optimization loop. The results for $N = 1, 3, 5$ environments are shown in Fig. 7.

The increase in the learning convergence as the number of environments increases is evident, as the cost function drops by one order of magnitude compared to the single-environment scenario for an observation time of T =10 min. The increase in performance between 3 to 5 environments is much less significant, hinting that saturation with respect to the number of environments is to be expected.

Table 2: Mean coefficients obtained during the model training, alongside their respective standard deviation for the multi-environment network. The assimilation of each case was repeated 10 times.

No. of tanks	min $\mathcal{J}$	θ_0	θ_1	θ_2	θ_3	θ_4	θ_5
1	3.8e-2	134.5 ± 16.3	10.0 ± 3.6	49.6 ± 5.5	2.0 ± 0.3	18.1 ± 19.2	10.0 ± 0.6
3	3.9e-3	137.3 ± 19.9	10.6 ± 4.0	49.8 ± 0.7	2.0 ± 0.1	21.3 ± 22.8	10.0 ± 0.1
5	6.9e-3	138.7 ± 32.8	10.2 ± 3.4	50.0 ± 4.8	2.0 ± 1.3	19.3 ± 2.2	10.0 ± 10.7
Reference	-	140.0	10.0	50.0	2.00	20.0	10.0

Like in the previous section, Table 7 collects the average and standard deviation of the identified parameters over 100 realizations in the case of multiple environments. The accuracy in the detection is comparable, although the standard deviation grows for some coefficients (e.g. θ_4 and θ_5). This might be due to the relative importance of these terms with respect to the various scenarios considered: θ_0 strongly influences the model performance in the presence of sloshing, but the relative importance of sloshing events over the five scenarios is reduced with respect to the first environment as other physical mechanisms are more present.

5. Conclusions

We presented a multi-environment real-time data assimilation framework for the data-driven thermodynamic modeling of cryogenic fuel tanks.

The implemented approach sets up an optimization problem in which the error between system observations and model predictions must be minimized by acting on closure parameters. The model is based on a lumped

https://doi.org/10.52202/069564-0095

formulation of mass and energy balances for each of the system's control volumes (i.e., the vapor, liquid, and solid walls), and the optimization is carried out by combining a quasi-Newton approach with the adjoint-based evaluation of the cost function gradients.

The training data is collected by multiple environments, i.e. different tanks undergoing various operating conditions (e.g., sloshing, filling, etc.). To provide the first proof of concept of the proposed framework, synthetic data generated through the model was used as a reference for the learning algorithm. The heat transfer closure relations in this model depend on six closure coefficients. The assimilation algorithm aimed to identify these from noisy pressure and temperature observations, which simulate measurements sampled from multiple fuel tanks.

Various tests were performed for the data assimilation applied to a single environment: we assessed the impact of measurement noise on the training data and the effect of 'mini-batches' in the gradient computation alongside the duration of an observation within which the assimilation is carried out. Remarkably, it was found that measurement noise improves the data assimilation, allowing for better identifying the underlying coefficients in shorter observation times. Similarly, an optimal batch size of 20% of the available data produced systematic improvements in the assimilation convergence. Finally, further improvements in the assimilation were achieved in a multi-environment scenario considering three or five tanks undergoing widely different loads and control scenarios.

To conclude, this work provides a solid proof of concept for applying a multi-environment data assimilation framework for the data-driven calibration of thermodynamic models of cryogenic tanks, potentially enabling model predictive control in thermal management systems. In the era of the fourth industrial revolution, where intelligent sensors are becoming widely available, the perspective of deploying the proposed multi-environment assimilation to a large fleet of heavily instrumented tanks appears particularly promising.

Acknowledgments

This work has been funded by the Flemish Agentschap Innoveren & Ondernemen in the framework of the cSBO project HBC.2021.0680 "Clean Hydrogen Propulsion for Ships (CHyPS)". Pedro Marques is supported by the FRIA grant 40009348 from the 'Fonds de la Recherche Scientifique (F.R.S. -FNRS)'.

Appendix A Details of the multi-environment framework

Table 3 shows the vector of parameters used to close the heat transfer correlations. In addition, Table 4 summarises the geometric parameters and initial conditions associated with each fuel tank.

Table 3: Coefficients associated with the heat transfer coefficients in (8) used to produce the synthetic test cases. Each column pair indicates the c_i coefficients linked to the heat transfer term in the header.

Coef.	θ_0	θ_1	θ_2	θ_3	θ_4	θ_5
Value	140	10	50	2	20	10

Table 4: Geometrical parameters and initial conditions associated with each fuel tank. All tanks were filled with liquid hydrogen and were treated as upright cylinders.

Tank	Geometry			Initial conditions					Wall properties
case	R_i [m]	R_o [m]	H [m]	p_v^0 [bar]	T_v^0 [K]	T_l^0 [K]	T_w^0 [K]	H_l^0 [m]	$m_w c_w$ [MJ/K]
1	0.9	1.00	2.2	4.80	32.1	20.1	29.1	1.43	2.44
2	2.7	2.90	14	3.10	30.1	20.1	25	6.16	91.4
3	1.4	1.50	6.2	4.40	30.1	20.1	28.1	4.03	10.5
4	0.4	0.50	1.2	0.94	20.1	20.1	21.1	0.43	0.10
5	1.0	1.05	2.1	0.94	20.1	20.1	22.1	1.56	0.63

Nomenclature

Letter symbols

- A area, m^2
- A_0 forcing amplitude, m
- b maximum wave amplitude, m

$\mathbf{B}$ approximated Hessian inverse

C specific heat, $\mathrm{J}/(\mathrm{kgK})$

f frequency, Hz

g gravity, m/s^2

g integrand of the objective function

H height, m

h heat transfer coefficient, $\mathrm{W}/(\mathrm{m}^2\mathrm{K})$

$\hbar$ mass-specific enthalpy, J/kg

$\mathcal{J}$ cost function, $[-]$

$\mathcal{L}$ Lagrangian function, $[-]$

$\mathcal{L}_v$ latent heat of vaporisation, J/kg

m mass, kg

$\dot{m}$ mass flow rate, kg/s

$\dot{Q}$ heat transfer rate, W

R radius, m

T temperature, K

t time, s

U internal energy, J

u mass-specific internal energy, J/kg

V volume, m^3

$\dot{W}$ work, W

$\mathbf{y}$ observation vector

Greek symbols

α thermal diffusivity, m^2/s

β volumetric thermal expansion coefficient, K^{-1}

$\boldsymbol{\theta}$ vector of parameters

$\boldsymbol{\lambda}$ adjoint variable

$\boldsymbol{\nu}$ kinematic viscosity, m^2/s

$\boldsymbol{\rho}$ mass-specific density, kg/m^3

Subscripts and superscripts

a ambient

i interface

l liquid

$\{j\}$ environment counter

(k) iteration counter

ph phase change

v vapor

w walls

https://doi.org/10.52202/069564-0095

References

[1] Abramson H. N., Dodge F. T., Bauer H. F., Brooks G. W., Chu W., Dalzell J. F., Kana D. D., Reynolds W. C., Satterlee H. M., Silverman S. *The Dynamic Behavior of Liquids in Moving Containers*. Washington DC, USA: National Aeronautics and Space Administration, NASA SP 106, Scientific and Technical Information Division; 1966.

[2] Grotle E. L., Æsøy, V., *Dynamic modeling of the thermal response enhanced by sloshing in marine LNG fuel tanks*. Applied Thermal Engineering 2018;135:512-520.

[3] Michael B., Wietschel M., *The Future of Hydrogen—Opportunities and Challenges*. International Journal of Hydrogen Energy 2009;34:615-627. 10.1016/j.ijhydene.2008.11.014.

[4] Swinerd G., Stark J., *Propulsion systems*. In: Fortescue P., Swinerd G., Stark J., editors. Spacecraft Systems Engineering. West Sussex, England: Wiley. 2003. p. 169-210.

[5] Janic M., *The potential of liquid hydrogen for the future 'carbon-neutral' air transport system*. Transportation Research Part D: Transport and EnvironmentTransportation Research Part D: Transport and Environment 2008;13(7):512-520. doi.org/10.1016/j.trd.2008.07.005

[6] Petitpas G., *Boil-off losses along LH2 pathway.* United States: Lawrence Livermore National Lab. (LLNL), USDOE National Nuclear Security Administration (NNSA); 2018 April, Technical Report No.: 1466121.

[7] Arndt T. *Sloshing of Cryogenic Liquids in a Cylindrical Tank under normal Gravity Conditions [dissertation]* Bremen, Germany: Universitat Bremen; 2011.

[8] Ludwig C., Dreyer M.E., Hopfinger E.J., *Pressure variations in a cryogenic liquid storage tank subjected to periodic excitations*. International Journal of Heat and Mass Transfer 2013. 66. 223-234. doi.org/10.1016/j.ijheatmasstransfer.2013.06.072.

[9] Byrd R., Lu P., Nocedal J., Zhu C., *A Limited Memory Algorithm for Bound Constrained Optimization*. SIAM Journal on Scientific Computing 2003;16. doi.org/10.1137/0916069.

[10] Granzow G. D., *A tutorial on adjoint methods and their use for data assimilation in glaciology*. Journal of Glaciology 2014;60(221):440-446. doi.org/10.3189/2014JoG13J205.

[11] Zhang K., Yang Z., Başar T., *Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms*. In: Vamvoudakis K., Wan Y., Lewis F., Cansever D., editors. Handbook of Reinforcement Learning and Control. Springer International Publishing. 2021. p. 321-374. doi.org/10.1007/978-3-030-60990-0_12

[12] Buşoniu L., Babuška R., De Schutter B., *Multi-Agent Reinforcement Learning: An Overview*. In: Srinivasan D., Jain L., editors. Innovations in Multi-Agent Systems and Applications - 1. Berlin, Heidelberg: Springer Berlin Heidelberg. 2010. p. 183-221. doi.org/10.1007/978-3-642-14435-6_7

[13] Foreest A., *Modeling of cryogenic sloshing including heat and mass transfer [dissertation]* Bremen, Germany: Universitat Bremen; 2010. 10.2514/6.2010-6891.

[14] Bergman T., Lavine A., Incropera F., DeWitt D., *Fundamentals of Heat and Mass Transfer*. Hoboken, New Jersey. Wiley. 2007.

[15] Bell I. H., Wronski J., Quoilin S., Lemort V., *Pure and Pseudo-pure Fluid Thermophysical Property Evaluation and the Open-Source Thermophysical Property Library CoolProp*. Industrial & Engineering Chemistry Research 2014. 53. 2498–2508. doi.org/10.1021/ie4033999.

[16] Moran M., McNelis N., Kudlac M., Haberbusch M., Satornino G., *Experimental results of hydrogen slosh in a 62 cubic foot (1750 liter) tank.* 30th Joint Propulsion Conference and Exhibit 1994. doi.org/10.2514/6.1994-3259.

[17] Navon I. M., *ractical and theoretical aspects of adjoint parameter estimation and identifiability in meteorology and oceanography*. Dynamics of Atmospheres and Oceans 1998;27(1):55-55-79. doi.org/10.1016/S0377-0265(97)00032-8.

[18] Li M., Zhang T., Chen Y., Smola A. J., *Efficient Mini-Batch Training for Stochastic Optimization.* Association for Computing Machinery 2014;661-670. doi.org/10.1145/2623330.2623612.

[19] Bollapragada R., Mudigere D., Nocedal J., Shi H., Tang P., *A Progressive Batching L-BFGS Method for Machine Learning.* 2018.