Politecnico di Torino

Politecnico di Torino

Department of Mechanical and Aerospace Engineering

Master's Degree in Automotive Engineering

Academic Year 2025/2026

CFD parametric study of gas venting and pressure management in an immersion-cooled lithium-ion battery module during thermal runaway

Supervisors:
Prof. Angelo Bonfitto
Prof. Vittorio Ravello

Candidate:
Filippo Cardegna

RINGRAZIAMENTI

Grazie a tutte le luci che mi hanno permesso di non perdermi nella Selva Oscura.

Un grazie speciale ai miei genitori Monica e Rocco che mi hanno garantito la possibilità di studiare nonostante tutto.

Grazie ai miei zii, alle mie sorelle Petra e Charlotte e a Paolo, che hanno sempre fatto il tifo per me.

Grazie anche a tutte le persone che mi hanno accompagnato lungo il mio percorso e che in qualche modo hanno contribuito alla mia crescita come persona.

E un grazie anche a me perché, alla fine, ce l'ho fatta.

Dichiarazione sull'uso di strumenti di Intelligenza Artificiale generativa

In conformità con le linee guida del Politecnico di Torino sulla trasparenza dell'uso di strumenti di Intelligenza Artificiale generativa nella redazione degli elaborati di tesi, dichiaro quanto segue.

Strumento utilizzato: Claude (Anthropic), accessibile tramite l'applicazione claude.ai.

Perimetro di applicazione: revisione finale del documento, successiva al completamento della modellazione fisico-matematica, delle simulazioni CFD in COMSOL Multiphysics e dell'analisi dei risultati.

Modalità d'impiego:

- revisione della coerenza tra sezioni, capitoli e appendici;
- verifica e cross-checking dei riferimenti bibliografici (autori, titoli, DOI, URL) e individuazione di citazioni non correttamente attribuite;
- riformulazione linguistica di paragrafi selezionati per migliorare chiarezza, concisione e uniformità terminologica;
- suggerimenti su struttura, formattazione e riorganizzazione di sezioni.

Attività non svolte tramite lo strumento: la scelta della strategia di modellazione, l'implementazione del modello CFD in COMSOL Multiphysics, l'esecuzione delle simulazioni numeriche, l'interpretazione fisica dei risultati, la progettazione dello studio parametrico e le conclusioni ingegneristiche sono opera esclusiva dell'autore.

La responsabilità sull'accuratezza dei contenuti, sulla correttezza delle citazioni e sulla conformità agli standard etici della ricerca resta interamente in capo all'autore.

Filippo Cardegna

Abstract

Immersion cooling with dielectric oils is a candidate architecture for the next generation of battery thermal management systems (BTMS), promising uniform temperature control and inherent electrical insulation of lithium-ion cells. During a thermal runaway (TR) event, however, the failing cell releases a substantial amount of high-temperature gas that must be safely expelled from the module through a dedicated pressure-relief architecture. The coupling between the gas released by the runaway cell, its buoyancy-driven and oil-flow-driven redistribution inside the immersion loop, and the dynamic response of a module-level pressure-relief valve on a closed hydraulic loop has not previously been resolved in a three-dimensional CFD model on the reference 21700 NMC cell format.

The present work develops a three-dimensional CFD model of a nine-cell 21700 NMC module immersed in a synthetic hydrocarbon dielectric coolant. The model, implemented in COMSOL Multiphysics 6.2, couples three physics interfaces (laminar single-phase flow, the two-phase Level Set method for the gas-oil interface, and heat transfer in fluids and solids) with six global ODEs, of which the two physically essential ones are the closed-loop compliance driven by the excess volume V_excess and the hysteretic valve state s, modelled as a pressure-threshold vent with separate opening and closing pressures; the remaining four (V_gas, V_produced, phase_TR, phase_done) act as diagnostic integrals and temporal latches for the heat-source staging.

A reference simulation reproduces the full sequence of the venting event and quantifies the peak loop pressure, the residual gas trapped inside the loop and the amount of oil expelled together with the gas. Three parametric sub-studies exercise the sensitivity to the module-vent diameter, to the closed-loop stiffness and to the position of the module vent combined with the location of the TR cell within the array. The results reveal a critical trade-off between safety and gas removal efficiency in the stiffness sweep — a stiff loop can reach a peak pressure well above the burst threshold of typical automotive plumbing — and identify a viable baseline configuration for the experimental scale-up of the immersion-cooled architecture, with the module vent positioned towards the outlet of the module to intercept the buoyant gas plume more efficiently.

Table of Contents

1. Introduction

1.1 Context and motivation

The diffusion of battery electric vehicles and of stationary energy storage systems has placed the lithium-ion battery at the centre of the ongoing energy transition in the transport sector. Two market trends are pushing thermal management in particular to the limit. The first is the steady increase in energy density at the cell and at the pack level, driven by larger-format cylindrical and prismatic cells and by tighter packaging. The second is the increase in peak DC fast-charging power: several premium models on the market now accept more than 320 kW in DC fast charging, with the Lotus Emeya reaching 350 kW (Lotus Cars, 2024) and the Porsche Taycan family operating around 320 kW (Porsche AG, 2024), roughly twice the figures of only a few years ago. Both trends translate, at the pack level, into higher heat dissipation per unit volume and tighter safety margins against thermal runaway.

Thermal runaway (TR) is the most severe failure mode a Li-ion cell can experience. A chain of exothermic, auto-catalytic reactions, triggered by mechanical, electrical or thermal abuse, leads to the decomposition of the solid electrolyte interphase (SEI), of the cathode structure, of the electrolyte and of the binder, with massive release of heat and of combustible gases (H_2, CO, CH_4, light hydrocarbons) and with opening of the cell safety vent (Feng et al., 2018). In automotive packs, where cells are integrated into tight, poorly ventilated modules, TR can propagate to adjacent cells; if combustible gases accumulate in a confined volume, deflagration or fire can follow (Feng et al., 2020).

Conventional battery thermal management systems (BTMS) based on forced air, on indirect liquid cooling via cold plates, or on phase-change materials, have proved either insufficient for sustained fast charging or unable to suppress the propagation of thermal runaway once it has been triggered (Wahab et al., 2025). Against this background, immersion cooling (IC), in which the cells are in direct contact with a dielectric fluid, has emerged in recent years as a promising route. It offers convective heat transfer coefficients significantly higher than indirect cold-plate cooling, excellent thermal homogeneity across the pack and, as shown by recent experimental work, a strong inhibiting effect on TR propagation (Roe et al., 2022).

Despite this body of evidence, the literature on immersion cooling has so far focused mainly on heat removal during normal operation and on the suppression of TR propagation. Comparatively little attention has been paid to a more specific question, which becomes critical in a closed and pressurised loop architecture: how is the gas released by a vented cell handled inside an immersion-cooled module? The dielectric oil is incompressible to a very good approximation; the vent gas is not. As soon as a single cell vents, a non-trivial volume of gas enters the closed loop, the system pressure rises sharply, and the question of where the gas should go, and at what threshold a passive safety device should release it to the outside, becomes a design problem in its own right.

1.2 Objectives of this thesis

The objective of this work is to develop, calibrate and exercise a three-dimensional CFD model of a single immersion-cooled battery module subject to a thermal runaway event in one of its cells, and to use it to derive design indications for the module-level venting system.

More specifically, the work addresses the following questions:

- How should the diameter of the module-level safety vent be sized, given the gas inventory of a single 21700 NMC cell and a given hydraulic stiffness of the closed loop, in order to keep the loop pressure within the design limits during a venting event?
- How does the hydraulic stiffness of the external closed loop, lumped into a single coefficient K, control the amplitude and the timing of the pressure transient, and how does the valve sizing have to be adapted to different values of K?
- How does the position of the module vent on the module ceiling affect the residual gas inventory at the end of the venting event, given that buoyancy tends to accumulate the gas in dead zones away from the inlet/outlet flow paths?

A two-dimensional model is developed first, as a verification step of the multiphysics coupling, the valve formulation and the lumped-loop closure (Section 4.1). The deliverable of the work is the three-dimensional parametric design study that quantifies the dependencies listed above and identifies the configurations for which the gas removal is most effective.

1.3 Document structure

The remainder of the thesis is structured as follows.

Chapter 2 reviews the state of the art on heat generation and degradation in lithium-ion cells, on the staged nature of the thermal runaway phenomenon, on battery thermal management strategies with a focus on immersion cooling, on the recent experimental evidence of TR under immersion cooling, on the vent technologies in use in commercial battery packs, and on the numerical modelling of cell venting.

Chapter 3 develops the physical and mathematical framework of the model: the governing equations of the two-phase flow, the Level Set formulation of the gas–oil interface, the surface tension and heat transfer formulations, the lumped representation of the closed external loop, the vent valve model with hysteresis, and the heat source and cell-vent opening trigger.

Chapter 4 documents the implementation of the model in COMSOL Multiphysics 6.2. It opens with a condensed account of a preliminary two-dimensional verification of the multiphysics coupling (Section 4.1), then describes the three-dimensional geometry, materials, physics interfaces, global definitions, ordinary differential equations, boundary conditions, mesh and solver setup that constitute the operational 3D model.

Chapter 5 presents the three-dimensional results: a reference simulation that reproduces the full sequence of a TR-driven venting event in the closed immersion-cooled loop, and a

parametric study that explores the sensitivity of the response to the four design variables identified in Section 1.2.

Chapter 6 summarises the main findings, discusses the limitations of the model and outlines future developments.

Two appendices close the document. Appendix A documents alternative modelling choices that were considered and discarded, and the main numerical difficulties encountered during the development of the model. Appendix B lists the nomenclature.

2. State of the art

2.1 Heat generation and degradation in Li-ion cells

During normal operation, a lithium-ion cell dissipates heat through three main mechanisms: ohmic (Joule) heating, due to the internal electrical resistance; polarisation heating, due to the irreversible kinetics of the electrochemical reactions; and reversible entropic heating, associated with the change in entropy during intercalation and deintercalation. The classical formulation by Bernardi et al. partitions the energy balance into irreversible and reversible contributions, the irreversible terms typically dominating (around 70 %) at the high C-rates of fast charging and acceleration.

The kinetics of degradation and abuse reactions is governed by Arrhenius-type laws. As an empirical rule, a sustained 1 °C rise in operating temperature in the 30–40 °C range reduces useful cell life by approximately two months, as reported in recent reviews on immersion cooling for automotive applications (Wahab et al., 2025). Under normal operation, degradation is dominated by SEI layer growth on the graphite anode, fed by the consumption of cyclable lithium and by the decomposition of the electrolyte. Non-linear SEI growth, amplified by micro-cracks induced by the volume changes of graphite during intercalation, is the leading degradation mechanism under fast-charging conditions.

Above approximately 90–120 °C the SEI layer starts to decompose. This exposes intercalated lithium directly to the electrolyte, triggering a sequence of exothermic reactions that, under unfavourable thermal boundary conditions, can lead the cell into thermal runaway.

2.2 Thermal runaway: mechanism, stages, gas composition

Thermal runaway is a self-sustaining thermochemical instability that occurs when the internal heat generation of a cell exceeds its capacity to dissipate heat to the environment, producing an exponential thermal acceleration. The literature identifies three main stages, separated by characteristic temperature thresholds, although the precise values depend on the cathode chemistry, on the state of charge and on the triggering condition.

Stage 1, SEI decomposition (≈ 90–130 °C): the SEI layer decomposes and exposes intercalated lithium in the graphite to the electrolyte, igniting exothermic reactions between LiC_6, the binder and the solvents. The $LiPF_6$ salt starts to decompose, releasing PF_5 that, in the presence of trace moisture, generates hydrogen fluoride (HF).

Stage 2, separator melting and internal short circuit (≈ 130–200 °C): the polyolefin separator melts and collapses, leading to large-area internal contact between the electrodes. Joule heating from the internal short adds to the chemical exothermic contributions, and the cell safety vent typically opens during this stage, releasing a first burst of gas. Wang et al. (2026) report this for LFP 26700 cells around 129–135 °C in dimethyl silicone oil immersion.

Stage 3, full runaway and venting (≥ 200 °C for LFP, lower for NMC/NCA): cathode decomposition, electrolyte oxidation and, in the most severe cases, partial reduction of the

cathode oxides accelerate the temperature rise. A large volume of gas (typically tens of litres at STP for a 21700 cell, hundreds of litres for large-format prismatic cells) is released through the vent disc on a time scale of tens of seconds. The composition is dominated by H_2, CO, CO_2, CH_4 and light hydrocarbons, with traces of HF; the H_2 and CO fractions make the mixture flammable and, in confined volumes, prone to deflagration.

NMC-811, considered in this work, is the chemistry of reference for high-energy density automotive applications, where the higher cell-level energy and the more reactive Ni-rich cathode make gas management during thermal runaway particularly critical. Compared to LFP, NMC-811 exhibits a lower onset temperature for safety venting (≈ 125 °C at 100 % SoC for the LG INR21700 M50, Abbott et al., 2022) and a substantially higher peak cell temperature (≈ 492 °C versus ≈ 360 °C for LFP under nail penetration without immersion, Liu et al., 2025). The volume of gas released, normalised by capacity, is also higher (≈ 1.30 L/Ah in inert atmosphere for NMC-811, Abbott et al., 2022). The selection of NMC-811 as the reference chemistry is therefore a worst-case engineering choice with respect to the design of the safety venting system: a vent architecture sized for the NMC scenario will also accommodate the milder dynamics of an LFP cell of comparable capacity. This is consistent with the approach adopted throughout this work, in which the baseline values of the main physical parameters (peak gas volume V_max, peak heat-release rate q_peak, gas injection temperature T_gas) are taken at the upper end of the experimentally observed range for 21700 NMC cells.

2.3 Battery thermal management systems

Conventional BTMS architectures for automotive applications can be grouped into four families. Forced-air cooling, used in earlier-generation EVs, offers simplicity but limited heat transfer coefficients and poor temperature uniformity. Indirect liquid cooling on cold plates, today's mainstream solution, brings the coolant in contact with the bottom or side surface of the cells via metallic plates and a glycol–water mixture; it offers better heat transfer than air but is intrinsically limited by the thermal contact resistance between cell and plate and by the small contact area. Phase-change materials provide passive heat capacity buffering for short transients but cannot sustain steady-state heat removal. Refrigerant-based direct cooling is used by some manufacturers and is effective at high cooling rates, but introduces system complexity.

Beyond steady-state heat removal, none of these architectures was originally designed to suppress thermal runaway propagation. When a cell enters TR, the small contact area between cell and cold plate is insufficient to extract the heat fast enough to prevent neighbouring cells from being heated above their onset temperature. Pack-level fire propagation is then a matter of seconds.

2.4 Immersion cooling: principles, fluids, benefits, hurdles

Immersion cooling places the cells in direct contact with a dielectric fluid that acts both as electrical insulator and as thermal vector. Two regimes are typically distinguished. Single-

phase immersion cooling uses a fluid that remains liquid over the operating temperature range; heat is removed by sensible heating of the coolant and by forced convection. Two-phase immersion cooling uses a fluid that boils at a temperature within the cell operating range; heat is removed mostly via the latent heat of vaporisation. Single-phase systems, of the kind considered in this thesis, are simpler from a sealing standpoint and are the leading candidate for first-generation automotive applications.

The dielectric fluids used belong to a few families. Synthetic esters such as MIVOLT DFK are biodegradable and have good thermal properties. Mineral oils (Shell Diala and equivalents) and dimethyl silicone oils offer high dielectric strength and chemical stability; the latter are the reference fluid in the experimental work of Wang et al. (2026) and in several recent immersion-cooling studies on dimethyl silicone oil at varying immersion depths. Engineered fluorocarbon fluids such as 3M Novec exist at the high-cost end, with the additional environmental concern of their PFAS classification. The properties required of an immersion coolant are: high dielectric strength, adequate thermal conductivity and heat capacity, moderate viscosity (to keep pump power reasonable), thermal stability over the operating range, low flammability and material compatibility with the components of the pack.

The benefits of immersion cooling over indirect liquid cooling are well documented. The wetted area is the entire surface of the cell rather than a single tab, the thermal contact resistance is essentially eliminated, the temperature uniformity across the pack is improved, the coolant inventory provides additional thermal capacity in transient events, and the cell-to-cell temperature gradients that are responsible for accelerated ageing in mismatched packs are reduced. The review by Wahab et al. (2025) reports heat transfer rates exceeding indirect-cooling values by orders of magnitude.

The hurdles are equally well known. The most relevant for the present work are the closed and pressurised nature of the loop, the leakage and seal challenge that the larger wetted volume introduces, and — directly relevant to this thesis — the management of the gas released during a venting event of a single cell. Whereas in air-cooled and cold-plate-cooled packs the vent gas is released into a partially open compartment, in an immersion-cooled closed loop the gas has nowhere to go unless an active or passive vent path is provided at the module or pack level.

2.5 Experimental evidence of TR under immersion cooling

Two recent experimental works are particularly relevant to the present thesis. Wang et al. (2026) investigated the heat and gas removal capability of immersion cooling on 26700 LFP cells in dimethyl silicone oil, with full immersion at a depth of 75 mm. They report a maximum cell-wall temperature of approximately 123 °C during the controlled overheating test, a delayed and significantly attenuated venting event compared to the same test in air, and a gas-mixture composition dominated by CO, H_2 and CO_2 with traces of HF and light hydrocarbons. Importantly, the immersion-cooled cell did not progress to a full thermal runaway: the dielectric

oil acted as a heat sink large enough to suppress the third stage. This is the reference dataset against which the calibration of the heat source Q_Cell(t) is performed in the present work.

Zhao et al. (2025) extended the analysis to large-format 280 Ah prismatic LFP cells, comparing four different TR triggers (overheating, overcharge, nail penetration and external short circuit) under liquid immersion cooling with a fluorinated dielectric coolant. They report that immersion cooling kept the maximum surface temperature below approximately 123 °C in the overheating test and 164.2 °C in the overcharge test, against peak temperatures of 471.7 °C and 551.6 °C measured on identical cells tested in air. Under nail penetration without liquid cooling the cell exploded, whereas under nail penetration with immersion cooling the same trigger produced severe deformation of the cell casing but no explosion. External short circuit, both with and without immersion cooling, did not trigger TR. They identify CO and H_2 as the primary vented gases, with CH_4 and HF as secondary species.

Ye et al. (2024) investigated the same problem on cylindrical 18650 cells, with a particular focus on the role of the immersion depth and on the propagation between cells. For the single cell, increasing the immersion depth from 20 mm to 70 mm extended the interval between safety-valve rupture and full TR trigger from 70 s to 312 s, and lowered the peak TR temperature to 306 °C. For multi-cell modules, they showed that immersion cooling inhibits the propagation of TR between adjacent cells, but does not prevent the rupture of their safety valves: in modules at 100 %, 75 % and 50 % state of charge, adjacent cells still reached peak temperatures of 118.2 °C, 121.4 °C and 111.5 °C respectively, with venting of the safety disc in all three cases. This last result is particularly relevant to the present thesis: it documents experimentally that, even in a well-designed immersion-cooled module, multiple cells release gas during a propagating event, and the pack-level gas management problem cannot be assumed solved by the immersion strategy alone.

Liu et al. (2025) brought a complementary perspective by testing nail penetration as the trigger, on prismatic 32 Ah LFP cells immersed in perfluorohexanone (FS49), a two-phase coolant. Without immersion the cell reached a peak temperature of 363.3 °C with a maximum temperature rate of 24.8 °C/s; under immersion cooling the peak dropped to 222 °C and the temperature rate to 8.4 °C/s, a reduction by a factor of approximately three on both quantities. The duration above 60 °C was reduced from 2792 s in air to 134 s in immersion. The gas inventory was also significantly reduced: aggregating the five main flammable species (H_2, CO, CO_2, C_2H_4, CH_4), the total volume dropped from ≈ 59 L per cell without cooling to ≈ 7.2 L with cooling, an ≈ 88 % reduction. The single-species reductions reported by the authors are 87.7 % for H_2, 95.8 % for CO_2, 76.2 % for CO, 72.2 % for C_2H_4 and 79.1 % for CH_4. These figures must be used with caution as a quantitative benchmark for an NMC system in dimethyl silicone oil: the cell chemistry, the coolant chemistry (FS49 is a two-phase coolant with low boiling point, silicone oil is single-phase) and the cell format are all different. Nevertheless, the order-of-magnitude reduction is a robust qualitative indication that immersion cooling significantly attenuates both the thermal dynamics and the gas inventory of a TR event, and

motivates the design of an effective module-level venting system even in an immersion-cooled architecture.

Together, these four studies provide the empirical anchor for the present numerical work. They confirm that immersion cooling can suppress TR rather than just delay it, but they also confirm that gas is still released during the first stage of venting and must be accommodated by the pack design.

2.6 Vent technologies for battery packs

2.6.1 ePTFE dual-stage membranes

The most widely used pack-level vent technology in commercial automotive battery packs is the dual-stage ePTFE (expanded polytetrafluoroethylene) vent membrane, of the kind manufactured by Donaldson (Tetratex product family) and Gore (PolyVent product family). Such a device performs two distinct functions on the same membrane. The first stage, passive, allows slow equalisation of pressure between the inside of the pack and the atmosphere through the microporous structure of the ePTFE, which is permeable to gas but retains water and contaminants thanks to its hydrophobic and oleophobic character. The second stage, active, occurs when the differential pressure exceeds a design threshold, typically in the range 250–300 mbar; the membrane ruptures or a safety cap is ejected, allowing the rapid release of the accumulated gas.

The design assumption behind ePTFE dual-stage vents is that the pack interior is, in nominal conditions, very close to atmospheric pressure, with the membrane handling only small variations due to temperature, altitude and moderate transients.

2.6.2 Why ePTFE membranes are not suited to a closed pressurised loop

The system considered in this thesis has features that make commercial ePTFE dual-stage vents structurally unsuitable. The cooling loop is closed and pre-pressurised at a nominal operating pressure of 1.2 bar absolute, slightly above atmospheric, in order to prevent cavitation in the pump and ingress of air through any unintended leak path. There is no air–gas exchange with the outside in normal operation, and there is no liquid–gas interface where the breathability function of the first stage could operate. During a venting event, once the membrane bursts, it no longer separates gas from liquid in any useful way: the loss-of-coolant path and the gas-escape path become the same.

Moreover, the typical 250–300 mbar opening threshold of commercial ePTFE vents is too low for the pressure dynamics expected in the present application. Automotive transient loads (pothole shocks, accelerations up to 4–5 g, pump pulsations, thermal expansion of the oil) can produce instantaneous pressure peaks in the closed loop of comparable magnitude, even in the absence of a venting event. Spurious openings of the vent would lead to coolant loss and degrade the integrity of the loop.

For these reasons, in the present work an explicit ePTFE membrane has not been modelled. The vent is instead represented as a generic pressure-threshold safety device,

functionally equivalent to a tuned rupture disc or to a spring-loaded relief valve sized for the application, with an opening overpressure higher than the worst-case transient and with a regularised, reversible behaviour suitable for parametric study (see Chapter 3).

2.6.3 Pressure-threshold valves with hysteresis

Mechanical pressure-threshold valves with hysteresis are routinely used in the process industry, in pneumatic and hydraulic systems and in pressure-relief applications. A defining feature is the presence of two distinct thresholds: an opening pressure p_open, at which the valve transitions from closed to open, and a closing pressure p_close < p_open, at which the valve returns to the closed state. The difference p_open − p_close, often called blowdown, prevents valve chatter and ensures that, once opened, the valve releases enough fluid to bring the system meaningfully below the danger threshold before reseating.

In a battery pack venting context, hysteresis has an additional benefit: it keeps the vent open long enough to remove not only the original gas surge but also any gas that has accumulated due to buoyancy in dead zones of the module. A simple non-hysteretic threshold valve, on the other hand, closes as soon as the instantaneous pressure drops below threshold, which in a closed loop happens very rapidly once the source ceases. In the present work, this hysteretic behaviour is adopted as the reference valve model, formalised in Section 3.4.3 and implemented as a global ODE on the valve state variable s in Section 4.6.

2.7 Numerical modelling of cell venting

Several numerical studies have addressed the modelling of gas venting from individual cells during TR. Kim et al. (2021) developed a CFD model of an 18650 cell, in which the flow inside the jelly roll is described by a Darcy–Forchheimer porous-medium formulation and the flow outside the cell is resolved by a k–ε RANS model. They identified two stages of venting (initial container breach and full TR gas release) and showed that the second stage dominates the gas budget. Their model is built around a single cell and does not address the pack-level closed-loop behaviour, but their kinetic and mass-flow data have been widely re-used.

Wang G. et al. (2023) coupled a CFD description of the gas transport with a lumped thermal resistance network for TR propagation between cells in a module. Their objective is the explosion-risk assessment in air-cooled, partially open compartments; the approach is well suited to that scenario but is not directly transferable to a closed, liquid-filled loop, since the gas transport in the CFD domain is single-phase and the pressure of the surrounding fluid is treated as atmospheric.

Kong et al. (2022), Ostanek et al. (2020) and Li et al. (2021) developed conjugate heat-transfer / CFD models for 18650 cells, focusing on jet fire and on the high-resolution 3D structure of the venting plume. These references justify the incompressible simplification of the present 3D model: they show that the high-Mach features of the venting plume relax to subsonic velocities within a few vent diameters from the cell exit, which is what allows an incompressible, lower-velocity representation downstream.

To the author's knowledge, no published work has so far combined an immersion-cooled module with an explicit closed-loop compliance model and an active vent valve, which is precisely the gap the present thesis aims to fill.

3. Physical and mathematical framework

3.1 Modelling strategy and simplifications

The phenomenon of thermal runaway with venting in an immersion-cooled battery module is intrinsically multiphysical and three-dimensional. A complete simulation would require the coupling of heat transfer, chemical kinetics of the degradation reactions, two-phase fluid dynamics, structural mechanics of the cell casing and of the module enclosure, with time scales ranging from milliseconds (vent disc rupture) to minutes (full propagation to the module). To keep the study tractable, and to maintain the interpretability of the results, a deliberately simplified but physically coherent approach has been adopted.

The main simplifications are listed below.

- Three-dimensional geometry. The model represents the full three-dimensional layout of the module, with nine cells in a 3×3 arrangement, four plastic supports, the oil cavity, the cell vent and the module vent. A two-dimensional version of the model has been used in the verification phase (Section 4.1); it is preserved in the document as a methodological step, but it is not part of the parametric study.
- No internal chemistry. The chemical kinetics of the TR reactions inside the cell is not solved. Heat generation is imposed externally as a piecewise function Q_Cell(t), calibrated against the experimental record of Wang et al. (2026), and gas injection at the cell vent is triggered by a temperature threshold.
- Incompressible two-phase flow. Both the oil and the vent gas are treated as incompressible. The compliance of the system, which physically derives from gas compressibility, from elasticity of the loop and from the presence of an expansion chamber, is collapsed into a single lumped parameter K that links the closed-loop pressure to the integrated gas volume.
- Laminar flow regime. The Reynolds number in the oil flow at the chosen inlet velocity is well below the laminar–turbulent transition, and a laminar interface is used. The gas plume itself reaches the turbulent regime in two localised regions (the centre of the cell-vent injection face and the module-vent efflux region during valve opening), but a turbulence model is not introduced because the bulk of the module remains laminar and the integrated quantities of interest are relatively insensitive to the detailed turbulent structure at the vents. This choice and its consequences are discussed in detail in Section 6.2.

3.2 Governing equations

3.2.1 Conservation of mass and momentum

The motion of the two-phase fluid is described by the incompressible Navier–Stokes equations, with local density and viscosity depending on the local phase fraction:

$$\nabla \cdot u = 0$$

$$\rho\,(\partial u/\partial t + (u\cdot\nabla)\,u) = -\nabla p + \nabla\cdot[\mu\,(\nabla u + (\nabla u)^{T})] + \rho\,g + F_{st}$$

where u is the velocity field, p the pressure, ρ and μ the locally interpolated density and dynamic viscosity, g the gravitational acceleration vector, and F_st the volumetric force associated with surface tension at the oil–gas interface.

3.2.2 Level Set method for the interface

The interface between oil and gas is tracked by the Level Set method, which introduces a smooth scalar field φ defined throughout the domain. The field takes the value 0 in one phase and 1 in the other, with a numerical transition of finite thickness ε_ls across the interface. The evolution equation is the standard COMSOL conservative form with reinitialisation:

$$\partial\varphi/\partial t + u\cdot\nabla\varphi = \gamma\,\nabla\cdot[\varepsilon_{ls}\,\nabla\varphi - \varphi\,(1-\varphi)\,\nabla\varphi/|\nabla\varphi|]$$

The right-hand side is the reinitialisation term, which keeps the interface profile of constant thickness ε_ls and of constant reinitialisation velocity γ. The volume fraction of Fluid 1 (oil) is directly given by φ (denoted ls.Vf1 in the COMSOL nomenclature), and the volume fraction of Fluid 2 (gas) by (1 − φ).

It is important to stress that the transition region between the two phases is numerical, not physical. Oil and vent gas are immiscible and the real interface is sharp; the diffuse numerical interface is what allows the method to handle topologically complex interfaces (bubble formation, fragmentation, coalescence) robustly. The price is some loss of mass conservation, which in the present model is controlled and monitored through the global ODE on V_gas (Section 3.3).

3.2.3 Surface tension

Surface tension at the oil–gas interface generates a force per unit volume of the standard continuum form:

$$F_{st} = \sigma\,\kappa\,n\,\delta(\varphi)$$

where σ is the surface tension coefficient, κ the local interface curvature, n the unit normal to the interface, and δ a Dirac distribution concentrated on the interface. For the AmpCool AC-110 / vent gas interface, σ = 0.025 N/m has been adopted, in the typical range for synthetic hydrocarbon dielectric fluids at room-to-moderate temperatures (Coelho de Sousa Marques et al., 2021). See Section 4.5.1 for further discussion.

3.2.4 Heat transfer in solids and fluids

The temperature field in the entire computational domain is governed by the energy equation, solved by the COMSOL Heat Transfer in Solids and Fluids interface with appropriate constitutive laws for the solid sub-domains (the nine cells and the four plastic supports) and for the fluid sub-domain (the oil–gas mixture). In the fluid, the multiphysics Non-Isothermal Flow node couples the laminar-flow and the heat-transfer interfaces, so that the convective heat transport is driven by the velocity field computed by the laminar-flow interface. In a unified form:

$$\rho\, c_p\, (\partial T/\partial t + u \cdot \nabla T) = \nabla \cdot (k\, \nabla T) + Q$$

where T is the temperature field, u the velocity field (zero in the solid sub-domains, so that the convective term vanishes and only conduction is retained), ρ the density, c_p the specific heat capacity at constant pressure, k the thermal conductivity, and Q a volumetric heat source. In the fluid, ρ, c_p and k are interpolated via the Level Set volume fractions; in the solid sub-domains they are tabulated constants for the NMC 21700 cell and for the plastic support. The heat source Q is non-zero only inside the cell driven into thermal runaway, where the source term Q_Cell defined in Section 3.5 is applied.

3.3 Closed-loop lumped compliance

The external cooling loop is closed: the oil that exits the module returns to it via a pump, a heat exchanger and possibly other components (filters, expansion vessel). The total volume of the internal loop is not known a priori and depends on the specific installation. To represent the effect of the closed loop without simulating its geometry explicitly, a lumped-parameter model of the compliance is adopted.

The pressure at the module outlet is assumed to depend linearly on a lumped variable V_excess that tracks the net volume of fluid in excess in the closed loop with respect to the initial nominal state:

$$p_{outlet} = p_0 + K \cdot V_{excess}$$

where p_0 is the nominal operating pressure (1.2 bar absolute) and K is the hydraulic stiffness of the internal loop, defined as the pressure rise per unit volume of fluid injected. The K parameter aggregates the contributions of oil compressibility, of wall elasticity, of the expansion chamber (if any) and of any residual air trapped in the loop. Its value is not generally known for a specific installation, and it has been treated as a study parameter rather than as a fixed datum. A reference value $K = 1 \times 10^8$ Pa/m³ has been chosen for the baseline runs; sensitivity to K will be quantified in the parametric study (Section 5.2).

The variable V_excess itself evolves in time according to a balance between the gas injected into the module from the cell vent and the fluid (gas or oil) that leaves through the module-level vent wh en the valve is open:

$$dV_{excess}/dt = Q_{inlet}(t) - Q_{vent}(t)$$

where Q_inlet(t) is the rate of gas injection from the cell vent into the module domain and Q_vent(t) is the rate of fluid efflux through the module-level vent valve, counted regardless of phase. Both terms are computed by line integration on the corresponding boundaries (Section 4.5). A separate variable V_gas is also tracked, by the same balance but with the efflux term restricted to the gas fraction, as a diagnostic of the gas content of the loop; the discussion of the distinction between the two is deferred to Section 4.6.

3.4 Vent valve model

3.4.1 Boundary-condition representation

The module vent is represented as a boundary condition on the module ceiling, which switches between a closed (impermeable) state and an open (discharge to atmosphere) state according to the value of the loop pressure. Below the opening threshold, the boundary behaves as a wall, with no flow normal to it; above the threshold, it behaves as an outlet to atmospheric pressure. This boundary-condition representation lends itself naturally to a parametric study on the geometry of the vent, which is a key objective of the work.

An alternative approach, in which the vent is modelled as an explicit porous-medium sub-domain representing an ePTFE dual-stage membrane, was considered and discarded; the reasons are discussed in Appendix A.1.4. Two specific implementations of the boundary-condition approach were developed during model construction: a velocity-BC formulation, which prescribes the efflux velocity, and a pressure-BC formulation, which prescribes the boundary pressure. Only the pressure-BC formulation is used in the runs of this thesis and is described in the next subsection; the velocity-BC formulation is documented in Appendix A.1.5 for completeness.

3.4.2 Pressure-BC formulation

In the boundary-condition formulation adopted in this work, the vent boundary is represented as a prescribed pressure that switches between the closed-loop pressure and the atmospheric pressure according to the valve state variable s:

$$p_{vent} = (p_0 + K \cdot V_{excess} + R_{vent} \cdot u \cdot n \cdot (1 - s)^2) \cdot (1 - s) + p_{atm} \cdot s$$

Here $s \in [0, 1]$ is the valve state variable governed by the hysteretic ODE described in the next subsection, R_{vent} is the closed-state hydraulic resistance, $u \cdot n$ is the outward-normal velocity component on the boundary, and p_{atm} is the atmospheric pressure to which the vent discharges. The $(1 - s)^2$ weighting placed inside the parenthesis on the $R_{vent} \cdot u \cdot n$ term, combined with the additional $(1 - s)$ factor multiplying the whole parenthesis, produces an effective cubic weighting $(1 - s)^3$ on the closed-state hydraulic-resistance penalty. This cubic weighting is introduced to suppress the spurious efflux resistance during the transient opening of the valve, when s takes intermediate values between 0 and 1: with a linear $(1 - s)$ weighting the high value of R_{vent} would keep contributing a residual penalty of the order of tens of bar for the typical efflux velocities (at $s = 0.7$, $R_{vent} \cdot u \cdot n \cdot 0.3 \approx 300$ mbar at $u \cdot n = 1$ m/s), producing a near-step opening behaviour inconsistent with the smooth nature of the s state variable. The cubic weighting reduces the residual resistance by two additional factors of $(1 - s)$, so at $s = 0.7$ it falls to $R_{vent} \cdot u \cdot n \cdot 0.027 \approx 27$ mbar at $u \cdot n = 1$ m/s, while the closed state $s = 0$ retains the full penalty needed to keep the boundary effectively impervious. When the valve is fully open ($s = 1$) both terms in the first parenthesis vanish, and the prescribed pressure is the atmospheric value p_{atm}.

When the valve is closed ($s = 0$) the boundary pressure equals the internal pressure of the loop, which, combined with the wall behaviour of the laminar-flow interface elsewhere,

produces zero net flux. When the valve is open (s = 1) the boundary pressure is fixed at atmospheric, and the efflux velocity emerges from the solution of the local pressure field. This formulation avoids the need for an arbitrary numerical cap on the efflux velocity and is the one used for the 3D reference run of Section 5.1 and for the parametric study of Section 5.2.

3.4.3 Hysteretic valve model

As discussed in Chapter 2, a realistic representation of a mechanical safety valve includes a hysteresis loop, with an opening pressure p_open higher than the closing pressure p_close (blowdown). This behaviour is implemented in the present model through an additional global ODE on a state variable s ∈ [0, 1], which represents the open/closed state of the valve:

$$ds/dt = [flc2hs(p - p\_open, \varepsilon) \cdot (1 - s) - flc2hs(p\_close - p, \varepsilon) \cdot s] / \tau\_vent$$

The first term drives s towards 1 when the pressure exceeds p_open, the second drives s towards 0 when the pressure falls below p_close, and τ_vent sets the response time of the transition. In the range between p_close and p_open both terms are essentially zero and s retains its previous value: this is the memory effect that produces the hysteresis loop. The reference values adopted in the present work for the run discussed in Section 5.1 are p_open = 1.8 bar (absolute), p_close = 1.5 bar (absolute), ε = 5000 Pa for the regularisation width and τ_vent = 0.1 s for the response time. The choice p_close > p_0 = 1.2 bar ensures that, once the event is over, the valve reseats at a pressure slightly above the nominal operating point of the loop, avoiding numerical chatter around the nominal value. With this formulation, the state variable s enters directly the pressure boundary condition of Section 3.4.2 and the source term of the V_gas ODE.

The rationale for the specific numerical values of p_open, p_close and τ_vent is discussed below. The set pressure p_open must lie sufficiently above the nominal operating pressure p_0 = 1.2 bar to prevent spurious openings of the safety valve under normal automotive transients. Instantaneous pressure peaks of the order of a few tenths of a bar can be produced by pothole shocks up to 4 g, by pump pulsations during DC fast charging, and by thermal expansion of the oil during vehicle pre-conditioning. A margin of 0.6 bar between p_open and p_0 is sufficient to accommodate these transients without loss of coolant. The blowdown p_open − p_close = 0.3 bar corresponds to approximately 17 % of the set pressure in absolute terms; this value is at the upper end of the range recommended by API Standard 520 and API Standard 521 for spring-loaded pressure-relief valves in compressible-fluid service, which are typically 7–10 % for conventional spring-loaded valves and up to 16–21 % for pilot-operated valves and fire-case scenarios (API Standard 520, 2020; API Standard 521, 2014; ASME Boiler and Pressure Vessel Code Section VIII, 2023). The deliberately wide blowdown reflects a specific trade-off of the venting event modelled in this thesis: the safety valve must remain open long enough to expel a substantial fraction of the injected gas, but not so long that an excessive volume of dielectric oil is co-expelled together with the gas once the initial gas surge has been ejected. A narrower blowdown would close the valve prematurely, reducing the gas-removal efficiency and leaving a larger residual gas volume V_gas trapped in the loop (Section

5.1.6); a wider blowdown, on the contrary, would keep the valve open beyond the useful interval and progressively expel oil in place of gas, without a proportional reduction in the residual V_gas, which is ultimately set by buoyancy-driven accumulation in the dead zones of the recirculating oil flow. The residual gas volume that persists after the event is what sustains the final loop pressure at approximately p_close, slightly above the nominal p_0. The response time τ_vent = 0.1 s corresponds to the upper end of the physically realistic range for automotive-scale PRVs, whose actuation time — including the spring dynamics and the seal reseating time — is typically in the range 10–100 ms, as characterised by the classical theory of spring-loaded valves (Merritt, 1967; Rabie, 2009). This is a deliberately conservative choice from the safety-sizing standpoint: as discussed in Section 5.1.4, the pressure overshoot of the closed loop increases monotonically with τ_vent, so adopting the slow end of the range corresponds to a worst-case estimate of the peak loop pressure. The value is also convenient numerically, being short enough compared to the characteristic time scale of the venting transient (order of seconds, Section 5.1) yet long enough to avoid excessive stiffness of the ODE on the valve state variable s; adopting a faster response at the lower end of the range would require a proportional reduction of the transient solver time step without qualitatively changing the venting dynamics.

3.5 Heat source model and calibration from the literature

The cell driven into thermal runaway in the reference simulation is the central row cell of the inlet-side column of the 3×3 grid (see Section 4.2 for the geometry). This location combines two features: it lies on the transverse symmetry plane of the module, and it receives the incoming oil flow directly from the inlet, so it experiences the strongest convective cooling among the possible TR-cell positions in the module. Alternative TR-cell locations, chosen to span a range from the best-cooled to the worst-cooled configuration, are explored in the parametric sub-study of Section 5.2.4. Heat generation inside this cell is imposed as a volumetric heat source applied to the cell domain, expressed as a power density in W/m³. The choice of a volumetric formulation rather than a lumped total power is deliberate: the resulting value is an intensive quantity, independent of the geometric reduction adopted for the model, and can therefore be transferred unchanged to the three-dimensional model of the next phase. It is also the form in which peak heat-release rates are most often reported in the experimental literature on cylindrical cells in thermal runaway, which makes the calibration against published data straightforward.

The temporal profile of the heat source is governed by the latch variable phase_done, formally introduced in Section 4.5, which switches irreversibly to 1 when the cumulative gas volume V_produced reaches the empirical cap V_max. The heat source has therefore a two-state profile:

- Active TR (phase_done = 0): the heat source is applied at the peak rate q_peak until the cell temperature approaches the empirical cap T_peak = 600 K, beyond which the source is smoothly turned off through the gate flc2hs((T_peak − T_cell)/1[K],

T_smooth). This single-stage formulation is consistent with a nail-penetration TR trigger, in which the cell enters thermal runaway from the ambient temperature directly, without a slow external pre-heating ramp: the heat source is applied from t = 0 with no distinct pre-TR warming phase. The temperature cap T_peak = 600 K is directly anchored to the experimental characterisation of Arcand et al. (2025), who measured the peak cell surface temperatures during thermal runaway of two commercial 21700 NMC cells in a critical venting scenario and reported average peaks of 341.3 °C (614.5 K) for the high-power Molicel P42a and 320.2 °C (593.4 K) for the high-energy Samsung 50e; the mean of the two peaks, 603.9 K, is rounded to 600 K and adopted as the cap in the present model. This value acts as an apparent thermal saturation of the cell without resolving the chemistry of the decomposition reactions explicitly.

- Post-TR (phase_done = 1): Q_Cell = 0. The reactive material has been exhausted and no further heat is generated inside the cell. The latched flag prevents the source from reigniting if the temperature condition is met again.

The active-TR and post-TR regimes are combined in a single closed-form expression for the heat source that produces no discontinuities, suitable for direct use as the input field of the COMSOL "general heat source" node:

$$Q\_Cell(t) = q\_peak \cdot flc2hs((T\_peak - T\_cell)/1[K], T\_smooth) \cdot (1 - phase\_done)$$

The peak heat-release rate q_peak is calibrated against the experimental data of Arcand et al. (2025), who measured the venting dynamics of two commercial 21700 NMC cells (Molicel P42a and Samsung 50e) in a critical venting scenario, in which the positive vent of the cell is fully severed and the ejection of gas and solid material is maximised. The calibration is carried out via a lumped thermal balance on the cell, treated as a homogeneous mass during the rapid heating phase that accompanies the venting event:

$$q\_peak = \rho\_cell \cdot c\_p,cell \cdot (dT/dt)\_TR$$

Using the cell-averaged density obtained from the cell mass and outer volume ($\rho\_cell \approx 2890$ kg/m³ for a 70 g 21700 NMC cell), the standard reference value of the specific heat capacity ($c\_p,cell \approx 1000$ J/(kg·K) for cylindrical Li-ion cells), and the peak temperature rate observed by Arcand et al. for the high-power Molicel P42a cell ($\Delta T \approx 128$ K rise from the TR trigger temperature of 213.3 °C to the peak temperature of 341.3 °C in 376 ms, i.e. $(dT/dt)\_TR \approx 340$ °C/s), the resulting peak heat-release rate is $q\_peak \approx 1 \times 10^9$ W/m³. The same calculation applied to the high-energy Samsung 50e cell yields a value very close ($\approx 1.01 \times 10^9$ W/m³), which indicates that the result is robust within the family of commercial 21700 NMC cells. This value is therefore adopted as the baseline of the model.

An earlier version of the model used a lower peak value $q\_peak \approx 1.2 \times 10^8$ W/m³, calibrated on the broad experimental range reported by Chen al. (2020) for cylindrical Li-ion cells (peak HRR 1.1–11.8 kW per cell). That value was retained for the 2D preliminary verification of the multiphysics coupling documented in Section 4.1, but has been superseded in the final 3D reference setup by the Arcand-based value $q\_peak = 1 \times 10^9$ W/m³, which is directly anchored

to a measurement on the same chemistry and format of the cell in the model. The 2D verification runs are documented with the earlier parameter set: as discussed in Section 4.1, the 2D model is a verification-of-coupling exercise rather than a quantitative simulation, and the absolute values of cell temperature and loop pressure obtained there should not be read as physically representative of the final 3D reference case.

3.6 Cell vent opening trigger

The cell vent (the boundary on top of the NMC 21700 cell that is driven into TR) is modelled with a prescribed effective injection velocity, active only when the average cell temperature exceeds a threshold (in the COMSOL implementation this velocity is imposed as a Mass Flow Rate boundary condition rather than as a pointwise velocity, for reasons of numerical consistency with the no-slip cylindrical wall of the cell — see Section 4.7.1; the two formulations are equivalent in the integral sense at u_gas_peak):

u_inj(t) = u_gas_peak · flc2hs((T_cell − T_open_cell)/1[K], T_smooth) · (1 − phase_done)

The temperature gate flc2hs((T_cell − T_open_cell)/1[K], T_smooth) is what activates the gas injection when the cell temperature crosses the venting threshold T_open_cell. The latch variable phase_TR (introduced in Section 4.5) is not part of this boundary condition: it is a diagnostic marker that switches irreversibly from 0 to 1 the first time the same temperature gate opens, and its function is to record the instant at which the venting event begins for the post-processing of the transient (Section 4.6).

The average cell temperature T_cell(t) is defined as the area-averaged temperature of the TR cell domain, computed by an aveop2 operator in COMSOL. The threshold T_open_cell = 480 K (207 °C) is adopted as the temperature at which gas starts to be released from the cell into the module. The value is calibrated on the venting characteristics reported by Arcand et al. (2025) for 21700-format NMC cells: for the high-energy 50e cell, whose format and chemistry are closest to the LG INR21700 M50 used as reference in the present work, gas emission is observed to start in the range 202–208 °C (475–481 K), just below the thermal-runaway trigger measured at 217 °C (490 K). It should be stressed that the scenario modelled in this work is a nail-penetration TR event, in which the internal short circuit drives the cell into TR abruptly, without a slow external pre-heating ramp. The distinction between a slow pre-venting stage and a subsequent chemical trigger, observed in the external-heating experiments of Arcand et al. under controlled temperature rise, collapses in the nail-penetration case into a single fast venting event, in which the safety vent of the cell opens as soon as the internal pressure and the surface temperature reach the venting condition. The choice T_open_cell = 480 K identifies this single opening moment; the numerical smoothing of the flc2hs gate on a width of ±T_smooth·1[K] = 10 K only serves to regularise the numerical activation of the gas injection, and does not represent a physical two-stage venting process. The injection velocity is held constant at u_gas_peak during the active venting phase, gated only by the temperature flc2hs and by the (1 − phase_done) hard cut-off, without any depletion factor. In the reference simulation the sequence of events is therefore: at t = 0 the volumetric

heat source Q_Cell (Section 3.5) is activated on the TR cell and the cell temperature T_cell starts to rise from the initial value of 293.15 K; when T_cell crosses the venting threshold T_open_cell = 480 K, the temperature gate flc2hs on the cell-vent boundary opens and gas injection at the mean velocity u_gas_peak begins; venting proceeds at constant mean velocity u_gas_peak until the cumulative injected volume V_produced reaches V_max, at which point phase_done latches to 1 and the gas injection is permanently switched off. The heat source Q_Cell has its own independent cutoff at the thermal cap T_peak = 600 K (Section 3.5): as T_cell approaches T_peak, Q_Cell is smoothly reduced by the thermal gate, and the cell starts to cool by conduction to the surrounding oil while the venting continues. In the reference run the thermal cap trips a fraction of a second before the phase_done latch, so that Q_Cell has already dropped to near zero when phase_done finally forces both the heat source and the gas injection to zero. The exponential depletion factor exp(−V_produced/V_scale), used in earlier versions of the model to represent the progressive consumption of the reactive material of the cell, has been removed from both the heat-source expression (Section 3.5) and the cell-vent injection velocity in the final nail-penetration formulation adopted here. The reasoning is that in the nail-penetration scenario the TR is triggered abruptly by the internal short circuit rather than by a slow depletion of the reactants, and the regulation of the heat source through the thermal cap at T_peak is sufficient to reproduce the thermal-saturation behaviour observed in the immersion-cooling experiments of Wang et al. (2026). The baseline value u_gas_peak = 50 m/s is derived from the cumulative gas inventory of Abbott et al. (2022) and the venting duration observed by Arcand et al. (2025); its numerical value and the parametric sweep range adopted in the 3D study are documented in Section 4.5.1, and the implications for the Reynolds regime at the cell vent (transitional/turbulent at this velocity) are discussed in Section 6.2.

The use of a temperature trigger replaces the earlier choice of a fixed time window for gas injection, which made the model independent of the thermal field. Tying the cell-vent opening to T_cell makes the model self-consistent: the thermal physics drives both the cell temperature and, through the threshold, the gas injection.

4. COMSOL implementation

4.1 Preliminary 2D verification

4.1.1 Rationale and scope

Before the three-dimensional model was constructed, a simplified two-dimensional model was developed and exercised, with the sole purpose of verifying the multiphysics coupling, the valve formulation and the lumped-loop closure. The 2D model is not part of the parametric study, and the absolute values it yields are not directly comparable to experimental data. Two structural features of the 2D setup make this explicit: the 2D laminar-flow interface in COMSOL Multiphysics treats the out-of-plane dimension as unit depth (1 m), so that all volumetric integrals are scaled to a fictitious one-metre-deep cell; and the parameter set used in the 2D runs is an earlier version of the calibration, anchored to LFP literature available at the start of the work, with $V_{max} = 4.7 \times 10^{-3}$ m², $q_{peak} = 1.2 \times 10^{8}$ W/m³, $T_{open\_cell} = 402.65$ K and $\rho_{cell} = 2300$ kg/m³.

The runs documented in this section are presented as a methodological step on the structure of the model, not as a prediction of the physical system. The 3D model that constitutes the deliverable of the thesis (Sections 4.2 onward) uses an updated, NMC-coherent parameter set and is exercised with the full three-dimensional geometry of the module.

4.1.2 2D geometry and parameter set

The computational domain is a 2D rectangle of 150 × 90 mm, representing the cross-section of a three-cell immersion-cooled module. Three NMC 21700 cells are represented as rectangles of 21 × 70 mm at three x positions, with 5 mm inter-cell spacing and four plastic supports at mid-height in the inter-cell gaps. Two side plenums, a triangular one on the left and a semicircular one on the right, funnel the oil from the inlet, around the cells and to the outlet. The module vent is a 5 mm segment in the middle of the top edge, centred above the central cell.

The 2D mesh is a free triangular mesh of approximately 10 800 elements generated with the physics-controlled COMSOL setting, with refinement at solid–fluid interfaces and at inlet, outlet and vent boundaries, and a boundary layer of two elements on the wetted walls of the cells and of the module on the fluid side. The transient solver uses BDF with an initial step of 0.001 s and a maximum step of 0.05 s; a 20-simulated-second run takes about 30 minutes on the available workstation.

4.1.3 Verification of valve hysteresis

The first verification compared the loop pressure transient in three configurations. Without a module vent, the closed-loop pressure rises monotonically as gas accumulates and reaches approximately 9.5 bar at the end of the venting transient (with $K = 1 \times 10^{8}$ Pa/m²); the value is incompatible with the design limits of typical battery enclosures and quantitatively confirms the need for a vent. A non-hysteretic threshold valve contains the pressure rise around 2 bar but

exhibits clearly identifiable chattering during the transient, with rapid open/close oscillations driven by the lack of dead band. The hysteretic valve with p_open = 1.8 bar, p_close = 1.3 bar and τ_vent = 0.1 s contains the pressure rise to comparable peak values without chattering and brings the loop to a steady state at approximately 1.31 bar after the cell vent ceases, with a residual gas volume per unit depth around 108 mm² at t = 100 s.

4.1.4 V_excess vs V_gas formulation

A more subtle point clarified in the 2D phase concerns the choice of the variable that drives the lumped pressure of the closed loop. In an earlier version of the model the closed-loop pressure was defined as $p_{loop} = p_0 + K \cdot V_{gas}$, on the implicit assumption that the gas was the only fluid leaving the module and that any oil expelled at the module vent could be neglected. In practice, however, the pressure-boundary formulation of the module vent does not distinguish phases at the boundary level: when the valve opens, whatever is locally near the vent is allowed to leave, and a non-negligible fraction of the efflux can be oil rather than gas. Linking the lumped pressure to V_gas alone would therefore leave the system unaware of any oil mass leaving the loop, with two undesirable consequences: the lumped pressure would not relax when oil is expelled, and would keep building up as more gas is injected, even though the closed system has already lost a substantial amount of fluid; and the imbalance between gas injected, gas expelled and oil expelled would not appear anywhere in the model. Replacing V_gas with V_excess as the pressure driver cures both issues: any fluid leaving the module, irrespective of its nature, reduces V_excess and therefore the lumped pressure. V_gas is retained but only as a diagnostic of the gas inventory. This choice has been carried over to the 3D model.

4.1.5 The trapped-gas phenomenon

A robust feature of the 2D runs, observed across configurations, is the persistence of a residual gas pocket under the ceiling of the module at the end of the venting event. The pocket is the consequence of the density contrast between oil and gas, of the buoyancy-driven upward migration of the gas plume, and of the position of the module vent relative to the inlet/outlet flow paths. In the 2D geometry, where the module vent occupies a 5 mm segment in the middle of a 150 mm-wide ceiling, between approximately 50 % and 80 % of the gas produced by the cell remains trapped under the ceiling in the dead zones, depending on the value of K and on the inlet velocity. The phenomenon is qualitatively consistent with the experimental observations of Liu et al. (2025) on prismatic LFP cells in FS49: even under aggressive immersion cooling, a residual fraction of the vent gas remains in the module. Beyond confirming the qualitative behaviour, the 2D observation directly motivates one of the parametric sweeps of the 3D study (Section 5.2): the lateral position of the module vent on the ceiling.

4.1.6 Numerical parameters carried over to the 3D model

Several numerical choices established during the 2D phase have been retained unchanged in the 3D model: the hysteretic valve formulation with the smoothed flc2hs gates and the

relaxation time τ_{vent} = 0.1 s; the choice of V_{excess} as the variable driving the lumped pressure; the closed-state regularisation of the module-vent boundary with the $R_{vent} \cdot u \cdot n$ penalty term; and the general structure of the global ODEs (V_{gas}, V_{excess}, $V_{produced}$, s, phase_TR, phase_done), with the only modification of the dimensional units (m^3 in 3D, m^2 in the 2D unit-depth reduction). The Level Set reinitialisation parameters γ and ε_{ls} have been re-calibrated for the 3D mesh and the higher peak velocities encountered in 3D, as discussed in Section 4.8.

4.2 Three-dimensional geometry

The three-dimensional model represents a single immersion-cooled module containing nine cylindrical 21700 NMC cells, arranged in a 3×3 grid as shown in Figure 4.1. The cells are held in position by four plastic supports made of a generic thermoplastic material, and the module is enclosed in a sealed case that defines the oil-filled cavity. The cooling loop is closed: the oil is supplied through a circular inlet on one of the side walls and recovered through a circular outlet on the opposite side wall, both with the same diameter.

Figure 4.1: Three-dimensional geometry of the module used for the reference run. The nine 21700 cells are arranged in a 3×3 grid inside a rectangular case; the oil inlet is on the left side, the oil outlet on the right side, and the module vent is on the top wall in position A (aligned with the outlet-side column of cells). The four plastic supports at the corners of the module are also visible.

The cell that undergoes thermal runaway in the reference simulation is the central row cell of the inlet-side column of the 3×3 grid. Alternative TR-cell locations are examined in the parametric sub-study of Section 5.2.4. The cell-vent face is a circular boundary of 1 cm diameter ($A_{inj} = 7.854 \times 10^{-5}$ m^2) on the top of the TR cell. This value is consistent with the experimental observation of Shelke et al. (2022), who report that in severe thermal runaway of 21700 NMC cells the entire vent cap becomes completely dislodged, leaving a central opening of approximately 10 mm in diameter. In less severe TR events the effective vent area

is smaller (three ~1.5 mm openings before dislodgement of the cap); the choice of the dislodged-cap geometry represents the worst-case scenario in terms of gas mass flow rate, consistent with the critical venting-case scenario adopted by Arcand et al. (2025) as the reference condition of their transient measurements. In the present formulation, no initial gas sub-domain is placed above the cell vent: gas is introduced into the module exclusively through the cell-vent boundary condition (see Section 4.7.2 for the rationale of the unseeded Level Set initialisation). The module vent, which is the boundary on the case where the safety valve is implemented (Section 4.7.1), is a circular face of 10 mm diameter (5 mm radius) on the top wall of the case in the baseline configuration, positioned laterally relative to the venting cell at a distance that is one of the parametric study variables (Section 5.2). The diameter of the module vent is itself the object of a parametric sweep ($D_{module} \in \{5, 10, 15\}$ mm) documented in Section 5.2.2.

4.3 Materials

The material properties assigned to the sub-domains of the model are summarised below. The cell material is a homogeneous effective medium representative of a 21700 NMC cell, with $\rho_{cell} = 2890$ kg/m³, $c_{p,cell} = 1000$ J/(kg·K) and $k_{cell} = 0.85$ W/(m·K); the anisotropic conductivity of the wound jelly roll has been collapsed to the through-plane value as a conservative choice for the heat transfer to the surrounding oil. The dielectric coolant is AmpCool AC-110 (Engineered Fluids, USA), a synthetic hydrocarbon single-phase liquid coolant designed specifically for immersion cooling of battery modules (Engineered Fluids, 2026). The properties at 40 °C, taken from the manufacturer's technical data sheet, are: $\rho_{oil} = 820$ kg/m³ (density at 15.6 °C, treated as constant in the model), $c_{p,oil} = 2212$ J/(kg·K), $k_{oil} = 0.1359$ W/(m·K), $\mu_{oil} = 6.65 \times 10^{-3}$ Pa·s (dynamic viscosity, derived from the datasheet kinematic viscosity $\nu = 8.11$ cSt via $\mu = \nu \cdot \rho$). Additional properties relevant to the automotive application are flash point 193 °C, pour point −57 °C, dielectric strength > 60 kV, and boiling point: none (thermally stable, decomposes without boiling). AmpCool AC-110 is non-halogenated, non-toxic, biodegradable, and ISO 9001 certified. The vent gas is a custom material with a temperature-dependent density of ideal-gas form, $\rho_{gas} = p_{ref} \cdot M_{gas} / (R \cdot T_{inlet,smooth})$, with $M_{gas} = 25$ g/mol representative of an H_2 / CO / CO_2 mixture and $p_{ref} = (p_{open} + p_{close})/2 = 1.65 \times 10^5$ Pa. Two aspects of this expression deserve comment. First, the reference pressure p_{ref} is a constant, not the instantaneous loop pressure $p_{loop}(t)$. This choice is a consequence of the incompressible-flow assumption adopted for both phases (Section 3.1): a fluid treated as incompressible has a density independent of pressure by construction. A physically more accurate formulation would replace p_{ref} with $p_{loop}(t)$, so that ρ_{gas} responds dynamically to the loop pressurisation. This was attempted in an earlier version of the model using the built-in COMSOL Air material, whose density depends on the local pressure field; the resulting setup produced solver failures at $t = 0$, traced to the absence of a well-defined initial pressure field on the gas sub-domain before the venting event begins (the gas domain is initialised as an infinitesimally thin layer above the cell vent, where the pressure is not yet coupled to the loop ODE on V_{excess}). The

pragmatic solution adopted here is to freeze the reference pressure to the constant value p_ref = (p_open + p_close)/2 = 1.65 bar, which corresponds to the mean operating pressure of the safety valve across its open–closed cycle: during normal cycling p_loop oscillates in a narrow band around p_ref, and the resulting error on ρ_gas is bounded within ±10 % at the extremes of the cycle. Other gas properties are c_p,gas = 1100 J/(kg·K), k_gas = 0.025 W/(m·K), μ_gas = 2 × 10^{-5} Pa·s. The plastic supports are modelled as a generic engineering thermoplastic, with properties in the typical range reported in the literature for such materials: ρ_plastic = 1140 kg/m³, c_p,plastic = 1670 J/(kg·K), k_plastic = 0.25 W/(m·K). The specific choice of polymer has a second-order effect on the results, since the supports are thermally passive during the transient and do not participate in the heat-source model of Section 3.5. On the fluid domain (the oil cavity), the oil and gas materials are combined through a Multiphase Material node (labelled fluido_misto in the implementation), which exposes effective material properties to both the Laminar Flow and the Heat Transfer in Fluids interfaces. The effective density and dynamic viscosity are computed as volume averages weighted by the Level Set volume fractions (ρ_eff = (1 − Vf2)·ρ_oil + Vf2·ρ_gas, μ_eff = (1 − Vf2)·μ_oil + Vf2·μ_gas), the effective thermal conductivity is also a volume average, and the effective specific heat capacity is a mass average (c_p,eff = (m_oil·c_p,oil + m_gas·c_p,gas)/m_tot), in line with the COMSOL recommendation for two-phase systems with significant density contrast. The use of a Multiphase Material is essential for the present problem, where the local thermal conductivity changes by approximately a factor of six across the diffuse gas–oil interface (k_oil = 0.16 vs k_gas = 0.025 W/(m·K)) and would otherwise be assigned the single value of the material registered on the fluid domain, with a non-physical homogenisation of the heat conduction in the gas regions of the module.

4.4 Active physics interfaces

The COMSOL component for the 3D model, labelled comp2 in the implementation to distinguish it from the 2D verification component comp1, hosts the following active physics interfaces:

- Laminar Flow (spf2), which solves the incompressible Navier–Stokes equations on the fluid domains (oil and gas). The interface includes the 'Include gravity' option, with gravity at the standard value g = 9.81 m/s² in the −z direction; the option 'Compensate for hydrostatic pressure' is kept active.
- Heat Transfer in Solids and Fluids (ht2), which solves the energy equation on all sub-domains, with solid constitutive laws for the cells and the supports and fluid constitutive laws for the oil–gas mixture.
- Level Set (ls2), which solves the conservative Level Set equation on the fluid domains, tracking the interface between oil and gas. The reinitialisation parameters are γ = 5·u_gas_peak = 250 m/s (set to five times the mean gas injection velocity at the cell vent, above the peak local velocities observed in the reference transient — approximately 85 m/s at the centre of the cell vent and 160 m/s at the module vent

during valve opening; consistent with the COMSOL Knowledge Base recommendation $\gamma \geq v_{max}$ for two-phase flow, Olsson and Kreiss, 2005) and $\varepsilon_{ls} = 3 \times 10^{-3}$ m (approximately three times the local mesh size in the refined regions around the cell-vent and module-vent boundaries, so that the diffuse interface is resolved by approximately six mesh elements across its thickness). The rationale for the larger-than-default value is discussed in Section 4.5.1.

- Global ODEs and DAEs (ge2), which implements the six global ODEs of Section 4.6.
- Two multiphysics couplings: Two-Phase Flow, Level Set (tpf2), which couples Laminar Flow and Level Set and applies the surface tension force; and Non-Isothermal Flow (nitf2), which couples Laminar Flow and Heat Transfer in the fluid domains. The Wetted Wall multiphysics coupling sets the oil–gas contact angle on the solid walls.

4.5 Global definitions

4.5.1 Parameters

All numerical parameters of the 3D reference simulation are collected in the Parameters 1 node under Global Definitions and are listed below. The specific combination of cell format (21700 cylindrical), cell chemistry (NMC) and cooling strategy (dielectric-oil single-phase immersion) considered in this thesis is not directly covered by any single experimental study in the open literature: the available measurements refer either to cylindrical 21700 NMC cells in air (Abbott et al., 2022; Arcand et al., 2025; Shelke et al., 2022), or to immersion-cooled cells of different format and chemistry (Wang et al., 2026 on 26700 LFP; Liu et al., 2025 on 280 Ah LFP with FS49 coolant). Several of the parameters listed below have therefore been derived by the author by combining measurements from these sources, with an explicit statement of the assumptions and the associated uncertainty in each case. This derivation strategy is documented parameter-by-parameter below, and, where the input data are cross-chemistry or cross-coolant, the corresponding caveat is stated explicitly in Section 6.2.

- p_0 = 1.2 bar — nominal operating pressure of the closed loop, chosen above atmospheric to prevent cavitation in the pump and air ingress through any leak path.
- p_{atm} = 1.0 bar — atmospheric pressure outside the loop.
- p_{open} = 1.8 bar — design opening pressure of the valve. A value above ≈ p_{atm} + 0.5 bar prevents spurious openings due to non-TR transients.
- p_{close} = 1.5 bar — design closing pressure of the valve, giving a blow-down of 0.3 bar.
- τ_{vent} = 0.1 s — relaxation time of the valve state ODE.
- $R_{vent} = 1 \times 10^{9}$ Pa·s/m — hydraulic-resistance penalty included in the module-vent boundary condition when the valve is closed. The value has been calibrated during the 3D model debugging: initial choices at 1×10^{7} Pa·s/m produced non-physical pressure oscillations on the module-vent boundary at closed valve state, driven by a numerical feedback loop between the small residual oil leakage through the boundary and the

pressure imposed by the boundary condition itself. Increasing R_vent by two orders of magnitude to 1 × 10^9 Pa·s/m makes the closed-state boundary effectively impervious to any efflux (residual leakage below 10^{-10} m^3/s), eliminates the oscillations and stabilises the transient solver.

- K = 1 × 10^8 Pa/m^3 — hydraulic stiffness of the internal loop. The choice and the physical justification of this value are discussed below; the loop opens the safety valve when the cumulative excess volume in the closed loop reaches V_excess_crit = (p_open − p_0) / K ≈ 0.6 L.
- u_oil = 0.14 m/s — oil inlet velocity, calibrated on the reference physical model.
- u_gas_peak = 50 m/s — peak mean velocity of the gas injection at the cell vent, applied on a circular face of diameter 1 cm (A_inj = π · (5 mm)² ≈ 7.85 × 10^{-5} m^2; see below and Section 4.2 for the justification of A_inj). The corresponding peak volumetric flow rate is Q_peak = A_inj · u_gas_peak = 3.93 × 10^{-3} m^3/s ≈ 3.93 L/s at the injection temperature T_inlet_gas = 1000 K. Direct measurements of the local peak velocity at the safety vent of a 21700 NMC cell in immersion cooling are not available in the open literature. The value has been derived by combining three literature elements. First, Arcand et al. (2025) measured average volumetric flow rates of 57.7 L/s for the high-power Molicel P42a and 97.0 L/s for the high-energy Samsung 50e, adjusted to the direct ejection temperature, over acute venting durations of 376 ms and 291 ms respectively — but these values refer to the cell venting in air, without any immersion cooling. Second, Wang et al. (2026) measured that immersion cooling reduces the total volume of gas released by an average of 82 % across the main gas species (H_2, CO, CO_2, C_2H_4, CH_4) with respect to the same event in air, because the shorter exposure to high temperatures partially suppresses the gas-producing reactions of the cell. Third, the full venting event of a 21700 NMC cell extends over a longer time interval than the acute ejection phase measured by Arcand: Shelke et al. (2022) report that the jet-flame phase of a 21700 NMC cell in air, following the acute gas ejection, lasts approximately 4 seconds in radial nail-penetration tests and 18 seconds in axial nail-penetration tests, the latter being the configuration that most closely mimics the top-safety-vent release considered in the present work. In immersion cooling, the jet-flame phase is expected to be either absent or strongly suppressed by the dielectric oil surrounding the cell, but the corresponding duration under dielectric-oil immersion has not been directly measured for a 21700 NMC cell. Combining these three effects gives an expected volumetric flow rate in immersion cooling of the order of a few L/s, of which the adopted value 3.93 L/s represents an order-of-magnitude estimate. The choice of u_gas_peak (rather than the volumetric flow rate) as the primary parameter is dictated by the boundary condition adopted in the COMSOL implementation: the injection is imposed as a Mass Flow Rate on the cell-vent face (Section 4.7.1), and u_gas_peak is the equivalent mean velocity through A_inj. A consequence of the adopted value is that the modelled venting duration t_vent ≈ V_max / (A_inj ·

u_gas_peak) ≈ 1.65 s is stretched with respect to the 291–376 ms acute ejection measured by Arcand in air; the total gas volume balance is preserved by the choice V_max above, but the fine temporal structure of the ejection is not resolved. The rationale for the stretched-time modelling choice and its consequences on the interpretation of the transient are discussed in Section 6.2.

- V_max = 6.5 × 10^{-3} m^3 — empirical upper bound on the cumulative gas volume produced by the TR cell, evaluated at the injection temperature T_inlet_gas = 1000 K. The value is derived by combining three literature sources, since no direct measurement is available for a 21700 NMC cell under dielectric-oil immersion cooling. First, Abbott et al. (2022) measured a total gas volume of approximately 6.5 L at standard temperature and pressure (STP: 288.15 K, 1 atm) for the LG INR21700 M50 NMC-811 cell at 100 % SoC in inert atmosphere by mass spectrometry. Applying the ideal-gas law to correct this value to the injection temperature T_inlet_gas = 1000 K gives V(1000 K) ≈ 6.5 · (1000/288) ≈ 22.6 L. This correction is consistent with the volumes at direct ejection temperature reported by Arcand et al. (2025), of 20.4 L for the high-power P42a and 26.2 L for the high-energy 50e, whose average of 23.3 L further validates the range. Second, Wang et al. (2026) measured that immersion cooling of 26700 LFP cells in a dielectric fluid reduces the total gas released with respect to the same event in air by 87.7 % for H_2, 95.8 % for CO, 76.2 % for CO_2, 72.2 % for C_2H_4 and 79.1 % for CH_4, giving an average reduction of approximately 82 % and a minimum reduction (i.e. maximum residual, for C_2H_4) of approximately 72 %. Third, applying the most conservative reduction (72 %) to the 23.3 L average from Arcand gives V_max ≈ 23.3 · 0.28 ≈ 6.5 L. This value therefore represents the conservative worst-case gas inventory that combines air-vented gas volumes measured on the same cell format and chemistry as the reference (Abbott, Arcand) with the least favourable species-specific LIC reduction factor extrapolated from a different cell format and chemistry (Wang). The associated uncertainty is dominated by the cross-chemistry extrapolation of the reduction factor and by the assumption of a uniform gas temperature at the injection boundary.
- q_peak = 1 × 10^9 W/m^3 — peak volumetric heat-release rate inside the cell during the active TR stage.
- T_open_cell = 480 K (207 °C) — cell-vent opening temperature threshold. The value is directly calibrated on the pre-venting temperature range measured by Arcand et al. (2025) for the Samsung 50e high-energy 21700-NMC cell, which the same authors identify as the closest commercial cell to the LG INR21700 M50 reference of this work: pre-venting was observed between 202 and 208 °C (475–481 K), and the value 207 °C = 480 K adopted here corresponds to the upper end of that range, i.e. just below the thermal-runaway trigger measured at 217 °C (490 K), and adopted here as the effective venting onset in the nail-penetration scenario (see Section 3.6). See Section 3.6 for the rationale of the choice in the nail-penetration scenario.

- $\tau_{phase} = 0.1$ s — relaxation time of the latch variables phase_TR and phase_done.
- $V_{scale,phase} = 8 \times 10^{-5}$ m³ — smoothing width of the phase_done latch transition around V_produced = V_max.
- T_smooth = 10 — dimensionless smoothing parameter of the temperature gate flc2hs((T_cell − T_open_cell)/1[K], T_smooth) used in the cell-vent injection and in the phase_TR latch. With T_smooth = 10, the gate transitions smoothly from 0 to 1 as T_cell moves from T_open_cell − T_smooth·1[K] = 470 K to T_open_cell + T_smooth·1[K] = 490 K, i.e. on a total interval of 2·T_smooth·1[K] = 20 K around the venting threshold.
- $\sigma = 0.025$ N/m — surface tension coefficient for the AmpCool AC-110 / vent gas interface. This value is slightly below the range 27–30 mN/m reported by Coelho de Sousa Marques et al. (2021) for four polyalphaolefins (PAO6, PAO20, PAO32, PAO40) at the oil-air interface between 293 K and 323 K, and reflects the reduction of surface tension expected at the higher temperatures reached at the oil-gas interface during venting (gas ejected at 1000 K). The Bond number $Bo = (\rho_{oil} - \rho_{gas}) \cdot g \cdot L^2 / \sigma \approx 129$ (for L = 0.02 m, $\Delta\rho \approx 820$ kg/m³) and the Weber number at the cell vent $We = \rho_{gas} \cdot u_{gas,peak}^2 \cdot L / \sigma \approx 1000$ both indicate that capillary forces are dominated by buoyancy and by inertia respectively. Consequently the integrated quantities of interest (loop pressure dynamics, valve actuation, gas volume balance) are essentially independent of the precise value of σ within this regime.
- $\gamma = 5 \cdot u_{gas,peak} = 250$ m/s — Level Set reinitialisation velocity. The value has been set as five times the mean gas injection velocity at the cell vent ($u_{gas,peak} = 50$ m/s) as a conservative choice, sufficient to remain above the maximum velocity encountered anywhere in the domain during the venting transient. Two regions can reach velocities appreciably above $u_{gas,peak}$. At the cell vent, the Mass Flow Rate boundary condition (Section 4.7.1) leaves the solver free to redistribute the profile across the injection face: the resulting profile is peaked at the centre, with local values of order 85 m/s observed in the reference simulation (a factor of order 1.7 above the mean). At the module vent, once the safety valve opens, the efflux is driven by the pressure difference between the loop and the atmosphere and the peak local velocity reaches approximately 160 m/s. The chosen $\gamma = 250$ m/s remains above both peaks and satisfies the COMSOL Multiphysics Reference Manual recommendation $\gamma \geq v_{max}$ for a stable reinitialisation step (Olsson and Kreiss, 2005).
- $\varepsilon_{ls} = 3 \times 10^{-3}$ m — Level Set diffuse interface thickness. The value has been chosen to be approximately three times the local mesh size in the refined regions around the cell-vent and module-vent boundaries, so that the diffuse interface is resolved by approximately six mesh elements across its thickness. A larger value than the default $\varepsilon_{ls} = h_c/2$, where h_c is the local mesh size in the interface region, was found to be necessary to reduce the numerical stiffness of the multiphase flow on the 3D mesh,

especially in conjunction with the high gradient of material properties (viscosity, thermal conductivity) imposed by the Multiphase Material on the diffuse interface (Section 4.3).

- T_peak = 600 K (327 °C) — empirical peak cell temperature beyond which the heat source is sharply reduced (Section 3.5). The value is the arithmetic mean of the peak cell temperatures measured by Arcand et al. (2025) after the thermal runaway trigger for the two 21700-NMC cell types tested: 341.3 °C for the high-power Molicel P42a and 320.2 °C for the high-energy Samsung 50e, giving a mean of 330.75 °C = 603.9 K, rounded to 600 K in the model. Both cell types are 21700-format NMC cells with chemistry similar to the LG INR21700 M50 reference of this work. This value replaces earlier attempts to derive T_peak from a range between the immersion-cooled peak temperature reported by Wang et al. (2026) for a 26700 LFP cell (~495 K) and the unconstrained adiabatic value for 21700 NMC in inert atmosphere (~800 K, Abbott et al., 2022), by anchoring the parameter directly to a measurement on the same cell format and chemistry.
- T_inlet_gas = 1000 K (727 °C) — temperature of the gas injected at the cell vent during venting. The value is calibrated on the direct ejected-gas temperature measurements of Arcand et al. (2025), obtained using a fast-response thermocouple positioned 2 inches from the cell vent: the average ejection gas temperature is 719.0 °C for the Molicel P42a and 780.3 °C for the Samsung 50e, giving a mean of 749.6 °C ≈ 1023 K, rounded down to 1000 K in the model as a slightly conservative underestimate of the thermal load imposed by the gas plume on the surrounding oil. The individual test peak values reported by Arcand range from 640 °C to 930 °C (913–1203 K), so the adopted value sits well within the measured range. In the boundary condition of the cell vent (Section 4.7.3), the injected gas temperature is smoothed in time as T_inlet_smooth = T_initial + (T_inlet_gas − T_initial) · flc2hs((T_cell − T_open_cell)/1[K], T_smooth) with T_initial = 293.15 K, to avoid a step discontinuity in the thermal boundary condition at the moment of vent opening.

The choice of K = 1 × 10⁸ Pa/m³ as the baseline value deserves explicit justification, because no single measured datum exists for this parameter in the literature and the dynamics of the safety valve is directly controlled by it. The hydraulic stiffness K of the closed loop has been derived from a simple physical model of the external cooling circuit, treated as a fixed-volume container of dielectric oil to which a small excess volume V_excess is added during the venting event. Two internal volumes must be distinguished: the volume of liquid coolant V_oil filling the tubes, the heat exchanger and the pump body, and the volume V_air of the gas cushion in the expansion vessel (either a dedicated diaphragm accumulator or a residual entrapped-air pocket). For a compact automotive test rig of representative size, values of V_oil ≈ 2.5 L and V_air ≈ 1 L are assumed as the reference configuration.

The two volumes contribute to the loop stiffness through two independent compliance terms, which combine in series:

$$1 / K = V\_oil \cdot \beta\_oil + V\_air / (\gamma \cdot p\_0)$$

The first term is the compressibility of the liquid coolant, where $\beta_{oil} \approx 6 \times 10^{-10}$ Pa^{-1} is the isothermal compressibility of the synthetic hydrocarbon oil (bulk modulus ≈ 1.6 GPa, in line with the range 1.5–2.0 GPa reported for mineral and PAO-based hydraulic fluids). The second term is the isothermal compliance of the entrapped gas cushion at nominal loop pressure p_0 = 1.2 bar, with γ = 1.4 the ratio of specific heats of the cushion air. Numerically, for the reference configuration V_{oil} = 2.5 L and V_{air} = 1 L, the compressibility term of the liquid coolant is $V_{oil} \cdot \beta_{oil} \approx 1.5 \times 10^{-12}$ $Pa^{-1} \cdot m^3$, and the compliance term of the gas cushion is $V_{air} / (\gamma \cdot p_0) \approx 6 \times 10^{-9}$ $Pa^{-1} \cdot m^3$. The gas cushion compliance dominates by approximately four orders of magnitude, so the resulting stiffness is $K \approx (\gamma \cdot p_0) / V_{air} \approx 1 \times 10^8$ Pa/m^3. This value has been adopted as the baseline for the reference simulation. Two additional values are considered in the parametric sub-study of Section 5.2.3 to bracket the range of realistic BTMS designs: $K = 5 \times 10^7$ Pa/m^3, which corresponds to a larger cushion $V_{air} \approx 2$ L (softer loop), and $K = 1 \times 10^9$ Pa/m^3, which corresponds to a smaller cushion $V_{air} \approx 100$ mL (stiffer loop). It follows from the analysis above that the stiffness of the loop is essentially independent of the volume of liquid coolant within the ranges of interest: doubling or halving V_{oil} while keeping V_{air} constant modifies K by a fraction of a percent, because the cushion compliance dominates. The reference value V_{oil} = 2.5 L can therefore be regarded as a design parameter of the physical rig, but not as a driver of the loop stiffness modelled here.

At the baseline $K = 1 \times 10^8$ Pa/m^3, the critical excess volume that triggers the safety valve is $V_{excess,crit} = (p_{open} - p_0)/K = (1.8 - 1.2) \times 10^5 / 10^8 \approx 0.6$ L. This is approximately 9 % of the empirical gas inventory V_{max} of a single 21700 NMC cell derived above, which means that under the reference configuration the safety valve is expected to trip after a small fraction of the venting event has been completed, keeping the loop pressure well controlled throughout the transient.

4.5.2 Functions

The analytic function used in the model is the standard smoothed Heaviside flc2hs (transitioning smoothly from 0 to 1 over a width specified by its second argument). No other user-defined functions are needed; all the physical and numerical relations are expressed inline in the boundary-condition and ODE expressions.

4.5.3 Integration and averaging operators

Three component-level operators are defined on comp2. The operator intop3 is a surface integration on the cell-vent boundary, used in the source term Φ_{in} of the volume balances. The operator intop4 is a surface integration on the module-vent boundary, used in the efflux terms of the V_{gas} and V_{excess} balances. The operator aveop2 is a volume average over the TR cell domain, used to compute the cell-mean temperature T_{cell} that gates the cell-vent opening and that drives the heat source. The numbering 3, 4, 2 of these operators reflects the existence of the corresponding 2D operators intop1, intop2, aveop1 in the verification component comp1; the 3D operators are independent COMSOL entities and act only on the 3D geometry.

4.6 Global ODEs

The model relies on six coupled global ordinary differential equations, declared in the Global ODEs and DAEs interface of COMSOL. They track quantities that are scalar by construction (volumes integrated over the entire fluid domain, the state of the safety valve, the two phase latches that govern the temporal staging of the heat source). The six variables and their physical meaning are listed below.

- V_gas [m³] — volume of gas currently held inside the module and the lumped external loop. Used as a diagnostic of the gas content of the closed system.
- V_excess [m³] — net volume in excess in the closed loop, used as the driver of the lumped pressure p_loop = p_0 + K · V_excess.
- V_produced [m³] — cumulative volume of gas injected by the cell vent into the module, monotonically non-decreasing.
- s ∈ [0, 1] — state variable of the module-level safety valve.
- phase_TR ∈ {0, 1} — diagnostic latch that switches irreversibly to 1 the first time the cell temperature crosses the venting threshold T_open_cell, marking the TR-start instant of the transient. The variable is governed by the ODE d(phase_TR)/dt = flc2hs((T_cell − T_open_cell)/1[K], T_smooth) · (1 − phase_TR) / τ_phase, in the same form as the phase_done latch defined below, and reaches saturation on a time scale of order τ_phase = 0.1 s after the trigger. In the nail-penetration TR formulation adopted in Section 3.5, phase_TR does not enter the heat-source expression, and is retained only as a temporal marker that identifies the start of the venting event in the plots and in the post-processing of the transient (Section 5.1).
- phase_done ∈ {0, 1} — latch for the exhausted-cell stage.

The three volume-like variables share the same source term at the cell vent, defined as the surface integral of the prescribed gas injection velocity:

$$\Phi_{in}(t) = \mathrm{intop3}[\, u_{gas,peak} \cdot \mathrm{flc2hs}((T_{cell} - T_{open,cell})/1[K], T_{smooth}) \,]$$

V_produced is non-decreasing by definition, V_gas counts only the gas leaving the module, and V_excess counts every fluid element crossing the vent regardless of phase:

$$dV_{produced}/dt = (1 - phase_{done}) \cdot \Phi_{in}(t)$$

$$dV_{gas}/dt = (1 - phase_{done}) \cdot \Phi_{in}(t) - \mathrm{intop4}[\, v_n \cdot (1 - ls2.Vf1) \,]$$

$$dV_{excess}/dt = (1 - phase_{done}) \cdot \Phi_{in}(t) - \mathrm{intop4}[\, v_n \,]$$

Here v_n = u · n is the outward-normal velocity component on the module-vent boundary, evaluated by the integration operator intop4. The expression (1 − ls2.Vf1) selects the gas fraction of the efflux (it is equivalent to the local Level Set gas volume fraction ls2.Vf2). The valve gating is not applied explicitly in these balances: the valve state s acts directly on the pressure boundary condition of Section 3.4.2 (Eq. (3.7)), where the closed-state hydraulic resistance R_vent · u · n · (1 − s)² suppresses any spurious efflux when the valve is closed.

When the valve is closed the efflux velocity v_n at the module-vent boundary is therefore negligible, and the corresponding integrals intop4[v_n · (1 − ls2.Vf1)] and intop4[v_n] vanish naturally. All three volume balances are integrated from zero initial conditions.

The valve state is governed by the hysteretic ODE introduced in Section 3.4.3, with the loop pressure based on V_excess:

$$ds/dt = [\ flc2hs((p\_0 + K \cdot V\_excess - p\_open)/1[Pa], 5000) \cdot (1 - s) - flc2hs((p\_close - p\_0 - K \cdot V\_excess)/1[Pa], 5000) \cdot s\ ] / \tau\_vent$$

The remaining two ODEs of the system are the two latches phase_done and phase_TR. Only phase_done governs the model dynamics: through the (1 − phase_done) factors of Section 3.5 and Section 3.6, it forces both the heat source and the cell-vent injection to zero once the cumulative gas inventory reaches V_max. phase_TR is instead a purely diagnostic marker of the venting-onset instant (Section 3.6), and does not enter any other equation of the model. Both latches are governed by ODEs of the same functional form:

$$d(phase\_TR)/dt = flc2hs((T\_cell - T\_open\_cell)/1[K], T\_smooth) \cdot (1 - phase\_TR) / \tau\_phase$$

$$d(phase\_done)/dt = flc2hs((V\_produced - V\_max)/V\_scale\_phase, 5) \cdot (1 - phase\_done) / \tau\_phase$$

Both latches start at zero. phase_TR rises smoothly to one when the cell temperature crosses T_open_cell, on a time scale of order τ_phase, and then remains saturated: its function is to record the venting-onset instant for the post-processing (Section 5.1) without affecting the physical dynamics of the model. Analogously, phase_done switches to one when V_produced exceeds V_max and, through the (1 − phase_done) factors of Section 3.5 and Section 3.6, forces both the heat source and the gas injection to zero in the post-TR stage.

4.7 Boundary conditions

4.7.1 Laminar Flow

- Oil inlet — prescribed normal velocity u_oil, applied to the circular face on the side wall of the case.
- Oil outlet — prescribed pressure p_0 + K · V_excess on the opposite side wall, with the 'Suppress backflow' option active. The driving variable is V_excess so that the outlet sees the same lumped loop pressure as the module-vent boundary, regardless of whether the efflux at the module vent is gas or oil.
- Cell vent — prescribed mass flow rate $\dot{m}$ = A_inj · u_gas_peak · flc2hs((T_cell − T_open_cell)/1[K], T_smooth) · ρ_gas_inj · (1 − phase_done), applied to the circular face of area A_inj on the top of the central cell. The flc2hs gate activates the injection only when the cell temperature crosses the venting threshold, with a smoothing width of 2·T_smooth·1[K] = 20 K; the (1 − phase_done) factor ensures a clean cut-off when the empirical limit V_max is reached. The local gas density on the boundary is

computed as ρ_gas_inj = (p_open + p_close)/2 · M_gas / (R · max(T_inlet_smooth, 200 K)) consistent with the ideal-gas density definition adopted for the custom gas material, and evaluated at the boundary temperature T_inlet_smooth prescribed by the corresponding node in the Heat Transfer interface (Section 4.7.3). The use of T_inlet_smooth in the density expression, rather than the cell temperature T_cell, ensures that the effective normal velocity of the injected gas on the cell-vent face equals u_gas_peak when the gate flc2hs saturates to unity: had T_cell been used instead, the density mismatch between the mass-flow prescription (evaluated at T_cell ≈ 600 K at the thermal cap) and the actual boundary temperature (T_inlet_smooth = 1000 K) would have introduced an inflation factor of approximately 5/3 on the injected volumetric flow rate. The choice of a mass flow rate boundary condition, in place of the velocity boundary condition originally implemented, was motivated by a numerical issue: a prescribed velocity on the cell-vent face produces, at the boundary edge shared with the no-slip cylindrical wall of the cell, an inconsistent pointwise constraint (the velocity must be simultaneously u = u_gas_peak from the cell-vent BC and u = 0 from the no-slip BC). The mass flow rate boundary condition imposes the constraint as an integral over the face, leaving the solver free to redistribute the velocity profile across the face (peak velocity at the centre, vanishing at the edge), which removes the constraint conflict at the shared edge and improves the convergence of the transient solver.

- Module vent — pressure-BC formulation with the hysteretic state s, a closed-state hydraulic-resistance penalty and a cubic weighting of that penalty on (1 − s): p_vent = (p_0 + K · V_excess + R_vent · u · n · (1 − s)²) · (1 − s) + p_atm · s. When the valve is closed (s = 0) the term R_vent · u · n suppresses any spurious efflux with its full penalty; when the valve is fully open (s = 1) the prescribed pressure is the atmospheric value; at intermediate valve states the cubic weighting (produced by the combination of the inner (1 − s)² factor and the outer (1 − s) factor multiplying the whole parenthesis) prevents the high value of R_vent from dominating the BC and producing a near-step opening behaviour, as discussed in Section 3.4.2.
- All other boundaries — default no-slip wall, with the Wetted Wall multiphysics coupling applied to set the oil–gas contact angle on the solid walls.

4.7.2 Level Set

- Initial condition — Fluid 1 (oil) everywhere. The earlier formulation of the model used a small cylindrical sub-domain of approximately 2.5 mm height above the TR cell, initialised as Fluid 2 (gas), as a seed for the Level Set interface (an approach analogous to the one used in the COMSOL Conference paper by Rivera-Salinas & Cruz-Ramírez, 2015). This formulation was subsequently abandoned in favour of a fully unseeded Level Set initialisation, in which the gas is introduced into the module exclusively through the cell-vent boundary condition (Mass Flow Rate for the Laminar Flow interface, Vf2 = 1 inlet for the Level Set interface). The unseeded formulation

removes the non-physical initial gas pocket above the cell (no gas is present in a sealed immersion-cooled module before the TR trigger) and makes the activation of the cell vent the sole source of gas in the model. The numerical issues associated with the dynamic generation of an interface from a single-phase initial state, which had motivated the original seed-bubble approach, are mitigated in the present setup by the relatively large diffuse-interface thickness $\varepsilon_{ls} = 3 \times 10^{-3}$ m (Section 4.5.1), which provides sufficient regularisation for the Level Set scalar to nucleate from the cell-vent boundary without spurious oscillations.

- Cell vent (LS) — Fluid 2 (gas) injection, consistent with the mass-flow-rate injection at the same boundary in the Laminar Flow interface (Section 4.7.1).
- Oil outlet and module vent (LS) — Outlet nodes, allowing the interface scalar to leave the domain when the corresponding boundary is active.

4.7.3 Heat transfer

- Oil inlet — prescribed temperature 293.15 K.
- Oil outlet — outflow boundary.
- Cell vent — prescribed temperature T_inlet_smooth = T_initial + (T_inlet_gas − T_initial) · flc2hs((T_cell − T_open_cell)/1[K], T_smooth), with T_initial = 293.15 K and T_inlet_gas = 1000 K. This smoothing of the gas inlet temperature avoids the step discontinuity that would result from imposing T = T_inlet_gas directly at the moment of vent opening, which would create a 707 K gradient across the cell-vent face on a single time step and stall the time-dependent solver. With the smoothing, the inlet temperature ramps smoothly from 293 K to 1000 K as T_cell crosses the venting threshold.
- External walls — thermal insulation (default).
- TR cell volume — heat source Q_Cell(t) = q_peak · flc2hs((T_peak − T_cell)/1[K], T_smooth) · (1 − phase_done) applied as Heat Source in the cell domain, with the form derived in Section 3.5. The gate flc2hs((T_peak − T_cell)/1[K], T_smooth) holds the source at the peak value q_peak when T_cell is below T_peak, and switches it smoothly to zero when T_cell exceeds T_peak, over a transition width of 2·T_smooth·1[K] = 20 K. Once V_produced reaches V_max, phase_done latches to 1 and the source is permanently switched off.
- Initial condition — 293.15 K everywhere.

The Boussinesq approximation has been deactivated: oil density does not depend on temperature in the present model, consistent with the small temperature range of the oil in the runs analysed in Chapter 5.

4.8 Mesh

The mesh of the 3D model is shown in Figure 4.2. It is a free tetrahedral mesh, with refinement at the boundaries of the cells and supports, at the cell-vent and module-vent faces,

and at the interface between the gas sub-domain and the surrounding oil. A boundary-layer mesh of two prismatic elements is generated on the wetted walls of the cells and of the case on the fluid side, with a stretch factor of 1.1 and a thickness adjustment factor of 3.

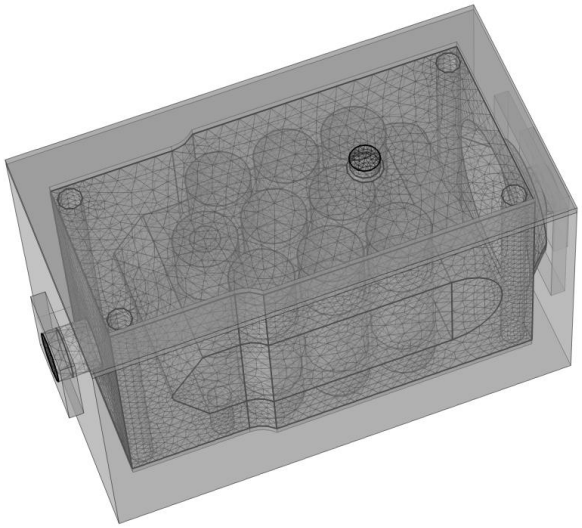

Figure 4.2: Final mesh of the 3D module used in the reference simulation and the parametric study. The mesh consists of approximately 120 000 tetrahedral elements with boundary layers on the cylindrical cell walls and localised refinement on the module vent and cell vent (visible on the ceiling). Minimum element quality on the residual slivers is approximately 0.155 (Appendix A.2).

The mesh has been the subject of a dedicated calibration step. Three competing requirements have been balanced. First, the in-core constraint on the PARDISO direct solver, which on the available 32 GB workstation sets a practical upper bound of approximately 250 000 fluid elements (the memory required by the factorisation scales approximately as N^1.5 to N²). Second, the requirement to resolve the Level Set interface across at least two elements in the regions of physical interest, which sets a local lower bound on the element size of approximately $\varepsilon_{ls}/2 \approx 5 \times 10^{-4}$ m. Third, the requirement to maintain an acceptable minimum element quality on the fluid domains, which is sensitive to the geometric features generated by the Boolean operations of the geometry sequence.

The final mesh of the reference run consists of approximately 120 000 tetrahedral elements (120 207 in the reference run; variations between the individual parametric runs are of the order of a few hundred elements and are not significant), with the bulk of the oil meshed with the 'Coarser' COMSOL preset ($h_{max} \approx 8$ mm) and the gas sub-domain meshed with the 'Finer' preset ($h_{max} \approx 3.34$ mm), plus local Size constraints with refinement factor 0.4 at the cell-vent and module-vent edges. The minimum element quality is approximately 0.155, concentrated on a small number of sliver elements at the intersection between the cylindrical cells and the planar walls of the case; the average element quality is approximately 0.65. The memory footprint of the PARDISO factorisation is approximately 13–16 GB, well within the in-core regime.

4.9 Solver setup

The study consists of two sequential steps:

- Step 1 — Phase Initialisation. Solves for the Level Set, the two-phase flow multiphysics and the global ODEs only, in order to establish the initial phase distribution. The laminar-flow interface and the heat-transfer interface are not solved in this step. With the unseeded initialisation adopted in the present formulation (Section 4.7.2), the step simply confirms ls2.Vf1 = 1 (oil) on the entire fluid domain and resets all the global ODE variables to zero. The step is retained for consistency with the standard COMSOL two-phase flow workflow.
- Step 2 — Time Dependent. Solves the fully coupled problem from t = 0 to t_end = 10 s. The output time list is range(0, 0.01, 10), i.e. one output every 10 ms across the entire simulation window, sufficient to resolve the sub-second venting transient and the subsequent long-term relaxation with the same temporal resolution.

The transient solver uses BDF with maximum order 2 and minimum order 1. The maximum time step is 1×10^{-2} s and the initial time step is 5×10^{-4} s. The fraction of the initial step used for the consistent initialisation with Backward Euler is 0.5, and the safety factor for the Backward Euler step is 30. The Fully Coupled Newton solver is used, with PARDISO as the linear solver in direct mode. The initial damping factor of Newton is 1, the maximum number of Newton iterations per step is 15, and the tolerance factor is 1. The relative tolerance of the time stepper is 0.005. These solver settings have been the subject of an extended trial-and-error calibration phase, which is documented in Appendix A.2.

5. Three-dimensional results

This chapter documents the three-dimensional simulation campaign of the venting system. Section 5.1 presents the 3D reference run at the baseline parameter set, used as the anchor case of the work. Section 5.2 presents the parametric sweeps on the three design variables retained in the study (D_module, vent position and K), each exercised on the 3D model with all other parameters held at their baseline values. Section 5.3 synthesises the outputs of the parametric study into a set of design implications for the module-level venting system. The 2D model documented in Section 4.1 was used only as a preliminary verification of the multiphysics coupling and of the lumped-loop closure prior to the 3D implementation, and is not part of the results reported in this chapter.

5.1 Reference three-dimensional simulation

5.1.1 Setup and parameters

The reference 3D run uses the parameter set of Section 4.5.1, with $K = 1 \times 10^8$ Pa/m³, u_gas_peak = 50 m/s, D_module = 10 mm (5 mm radius), $\sigma = 0.025$ N/m and $\varepsilon_{ls} = 3 \times 10^{-3}$ m. The mesh is the optimised one described in Section 4.8 (approximately 120 000 elements with minimum quality of approximately 0.155, obtained after several iterations on the geometry virtual operations and the mesh size constraints, see Appendix A.2). The time integration is set to extend to t_end = 10 s, which is sufficient to cover the cell heating phase, the cell-vent activation, the rise of the loop pressure, the opening of the module vent, the blow-down of the gas inventory and the post-event relaxation of the system. The actual time horizon reached by the simulation is reported with the results.

5.1.2 Cell heating and TR trigger dynamics

Figure 5.1 shows the time evolution of the volume-averaged temperature T_cell of the cell undergoing thermal runaway. The temperature rises from the initial condition T_initial = 293.15 K following the volumetric heat source q_peak (Section 3.5) applied from t = 0. During the first 0.6 s the heating is essentially linear, at a rate approximately consistent with $(dT/dt) = q_{peak} / (\rho_{cell} \cdot c_{p,cell}) \approx 340$ K/s prescribed by the calibration of Section 3.5, with a mild reduction due to convective heat loss to the surrounding oil. At $t \approx 0.6$ s the cell reaches T_open_cell = 480 K, at which point the Mass Flow Rate boundary condition at the cell vent is smoothly activated by the flc2hs gate on $(T_{cell} - T_{open\_cell})$. The switch of the latch variable phase_TR from 0 to 1 marks the onset of the venting event.

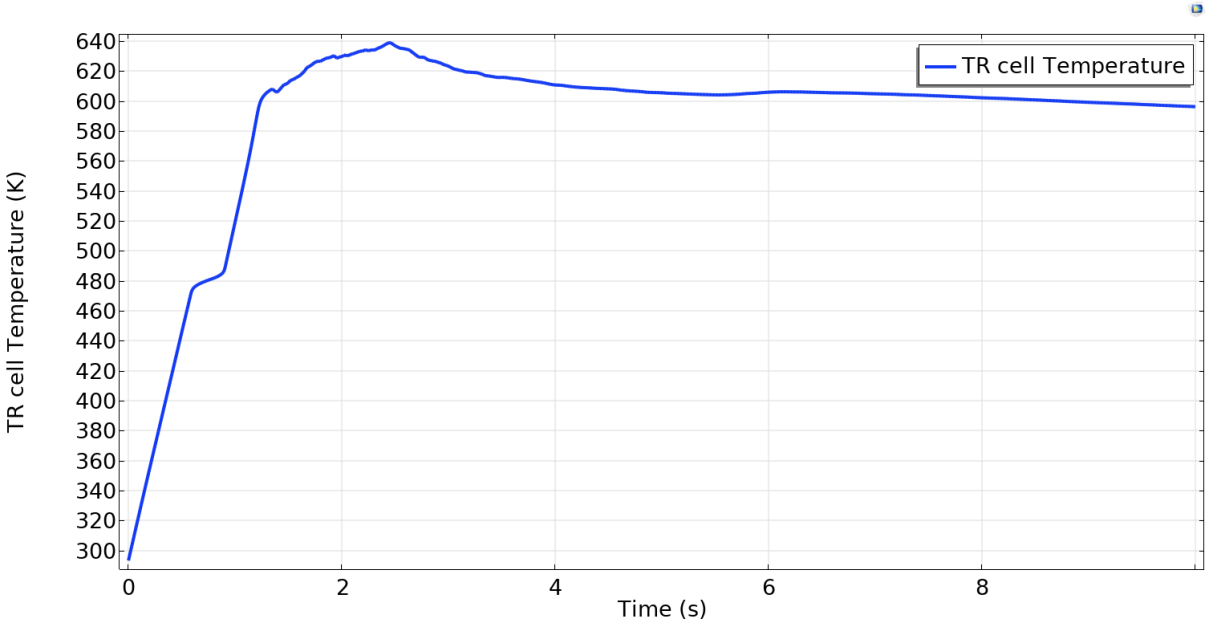

Figure 5.1: Time evolution of the volume-averaged temperature T_cell of the TR cell during the reference run. The initial rapid heating from T_initial = 293.15 K to T_open_cell = 480 K takes approximately 0.6 s; after the cell-vent activation the heating slows down, T_cell reaches an apparent peak of 638 K at t ≈ 2.4 s and then decreases smoothly to approximately 596 K at t = 10 s.

After the cell-vent activation the slope of T_cell(t) is reduced, reflecting the balance between the residual heat source, the convective heat loss to the surrounding oil and the enthalpy absorbed by the gas leaving the cell at the vent. T_cell continues to rise, reaches an apparent peak of approximately 638 K at t ≈ 2.4 s and then decreases smoothly. The mild overshoot of about 38 K above the empirical cap T_peak = 600 K (Section 3.5) is a consequence of the finite width of the flc2hs smoothing (of the order of T_smooth·1[K] = 10 K on either side of T_peak) combined with the thermal inertia of the cell during the transition of the gate: the source is not sharply switched off at T_peak but only smoothly reduced over an interval of ±10 K, and the residual heating during the transition drives T_cell above T_peak before the cap fully suppresses q_peak. The overshoot is quantitatively small (approximately 6 % of T_peak) and does not affect the physical interpretation of the results.

The latch variable phase_done switches from 0 to 1 at t ≈ 2.6 s, indicating that the cumulative gas volume V_produced has reached the empirical upper bound V_max = 6.5 × 10^{-3} m^3 = 6.5 L (Section 4.5.1). From that point onward both the heat source and the cell-vent mass flow rate are permanently switched off, and the cell enters a slow cooling phase governed by conduction to the surrounding oil. At t = 10 s the cell temperature is approximately 596 K and continues to decrease slowly, with a characteristic time significantly longer than the simulation window.

This evolution of T_cell is essentially identical across all the parametric sweeps of Section 5.2, because the temperature field inside the cells is set by the internal heat generation q_peak and by the direct dielectric-fluid contact on the outer surfaces, both of which are held constant

across the sweeps. For this reason, T_cell is not reported separately in the parametric studies that follow.

5.1.3 Gas venting and bubble evolution

Figures 5.3a–5.3f present six snapshots of the gas volume fraction field Vf2 (Level Set output) in the 3D module during the venting transient of the reference run, at t = 0.6 s, 0.88 s, 2.36 s, 3.2 s, 6 s and 10 s. The red regions correspond to gas (Vf2 ≈ 1), the blue regions to oil (Vf2 ≈ 0), and the intermediate colours to the diffuse interface of thickness $\varepsilon_{ls} = 3 \times 10^{-3}$ m. In all figures the oil inlet is on the left and the oil outlet is on the right, with the module vent located near the outlet on the top wall in the baseline configuration of position A (Section 5.2.4). The TR cell is the one in the central row of the inlet-side column, i.e. the cell directly under the far end of the ceiling opposite to the module vent.

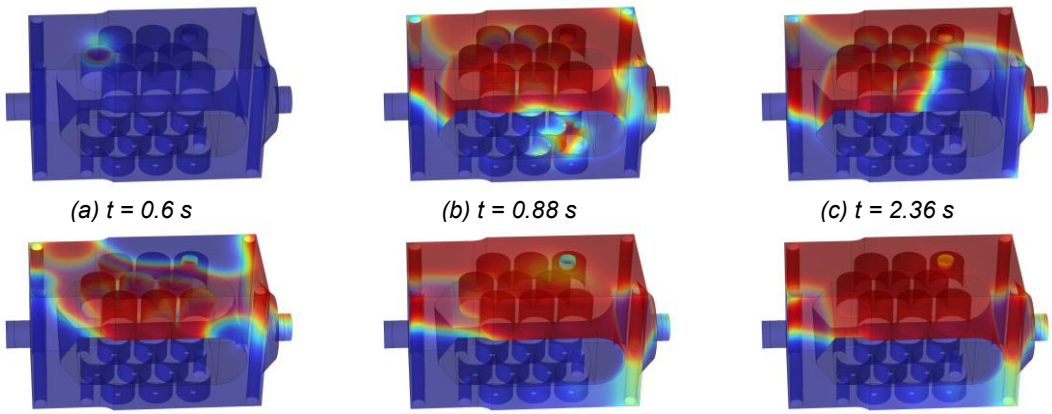

(a) t = 0.6 s *(b) t = 0.88 s* *(c) t = 2.36 s*

(d) t = 3.2 s *(e) t = 6 s* *(f) t = 10 s*

Figure 5.2: Gas volume fraction Vf2 in the 3D module at six representative timestamps of the reference run: (a) t = 0.6 s, cell-vent activation; (b) t = 0.88 s, module-vent opening; (c) t = 2.36 s, onset of phase_done; (d) t = 3.2 s, end of cell-vent injection; (e) t = 6 s, module vent almost fully closed (s ≈ 0.026); (f) t = 10 s, end of the simulation. Red = gas (Vf2 ≈ 1), blue = oil (Vf2 ≈ 0), intermediate colours = diffuse interface. The residual gas region under the ceiling visible in (e) and (f) corresponds to only a fraction of the total V_gas ≈ 0.9 L reported in Section 5.1.6; the remainder has been transported into the external portion of the closed loop through the oil outlet.

At t = 0.6 s (Figure 5.2a) the cell vent has just been activated and a small localised gas region has formed on the top of the TR cell, initiating the plume. At t = 0.88 s (Figure 5.2b), 0.28 s later, the plume has already stratified along the entire ceiling of the module: the upper half is essentially gas (red) and the lower half is essentially oil (blue), separated by a nearly horizontal interface of transition. This very rapid stratification is a direct manifestation of the dominance of buoyancy over the other body forces at the scale of the module: with a Bond number $Bo = (\rho_{oil} - \rho_{gas}) \cdot g \cdot L^2 / \sigma \approx 129$ (Section 4.5.1), capillary forces are negligible and the gas plume rises to the ceiling on the buoyancy timescale $L / \sqrt{(g \cdot L)} \approx 0.05$ s, well below the injection timescale.

Between t = 0.88 s and t ≈ 3.2 s the gas ejection at the cell vent is still active and the module vent is fully open (Section 5.1.5): the transient is not a simple monotonic filling of the upper half but an oscillatory redistribution between gas and oil. Local high-pressure regions

produced by the gas jet at the cell vent transiently push oil into upper areas of the module, from which the oil then falls back by gravity as the local pressure relaxes, and the gas moves in to replace the oil. This cycle is visible in the intermediate snapshots (Figure 5.2c and 5.5, at t = 2.36 s and 3.2 s): the interface is not flat but exhibits regions where oil has re-entered the upper half and other regions where the gas has penetrated into the lower half. The oscillation is damped by the viscous stresses in the oil and by the dissipation of pressure energy at the module vent, and repeats two or three times before the system stabilises into a stratified configuration.

At t = 3.2 s (Figure 5.2d) the cell-vent injection has just been switched off by the phase_done latch (Section 5.1.2) and the module vent is still open: the module is now in a pure blowdown phase in which the gas trapped under the ceiling is progressively expelled through the vent, and oil re-enters the module. By t = 6 s (Figure 5.2e) the module vent has almost fully closed (s ≈ 0.026, Section 5.1.5) and the system is essentially in its final stratified configuration; only minor readjustments of the interface are still visible between t = 6 s and t = 10 s (Figure 5.2f).

The final configuration shows a persistent gas region under the ceiling of the module, occupying approximately the upper half of the volume between the top of the cells and the top wall of the case. This is the three-dimensional manifestation of the trapped-gas phenomenon identified in the 2D preliminary verification (Section 4.1.5): once the module vent has closed, the gas that has migrated to the ceiling and is not adjacent to the vent opening cannot be expelled, and remains locked in place by buoyancy. The residual gas visible in the 3D domain at t = 10 s corresponds to only a fraction of the value V_gas ≈ 0.9 L reported in the volume balance of Section 5.1.6: the remainder of V_gas has been transported out of the 3D module through the oil outlet by the recirculating oil flow during the venting transient, and now resides in the external portion of the closed loop (external piping, expansion vessel and air cushion) that is represented in a lumped fashion by the ODE on V_excess. In the notation of the ODE formulation of Section 4.6, V_gas accounts for all the gas that has entered the closed loop from the cell vent minus the gas that has been ejected through the module vent to the atmosphere; the gas that has left the 3D module through the oil outlet is still part of the loop and is therefore still counted in V_gas. This distinction is important for the interpretation of the 3D snapshots: the residual gas visible under the ceiling of the module is the trapped fraction locked by buoyancy, while the invisible fraction is distributed along the external loop and might either be recirculated back into the module through the oil inlet or accumulated in the air cushion of the expansion tank in the long time limit; the present model does not resolve the distribution of gas in the external loop and does not distinguish between these two long-term fates.

5.1.4 Pressure dynamics in the closed loop

Figure 5.3 shows the time evolution of the loop pressure p_loop(t) = p_0 + K · V_excess(t) during the reference run. From t = 0 to t ≈ 0.6 s the pressure remains constant at the nominal operating value p_0 = 1.2 bar, since no gas has yet been produced (V_excess = 0). Starting

from t ≈ 0.6 s (cell-vent activation) the pressure rises rapidly as gas is injected at the cell vent, crosses the valve opening threshold p_open = 1.8 bar at t ≈ 0.9 s and continues to rise while the valve state s transitions from 0 to 1 over its characteristic time τ_vent = 0.1 s. During this transient the volumetric injection rate at the cell vent (Q_inj = A_inj · u_gas_peak ≈ 3.93 L/s) exceeds the ejection rate through the partially open module vent, and V_excess continues to accumulate.

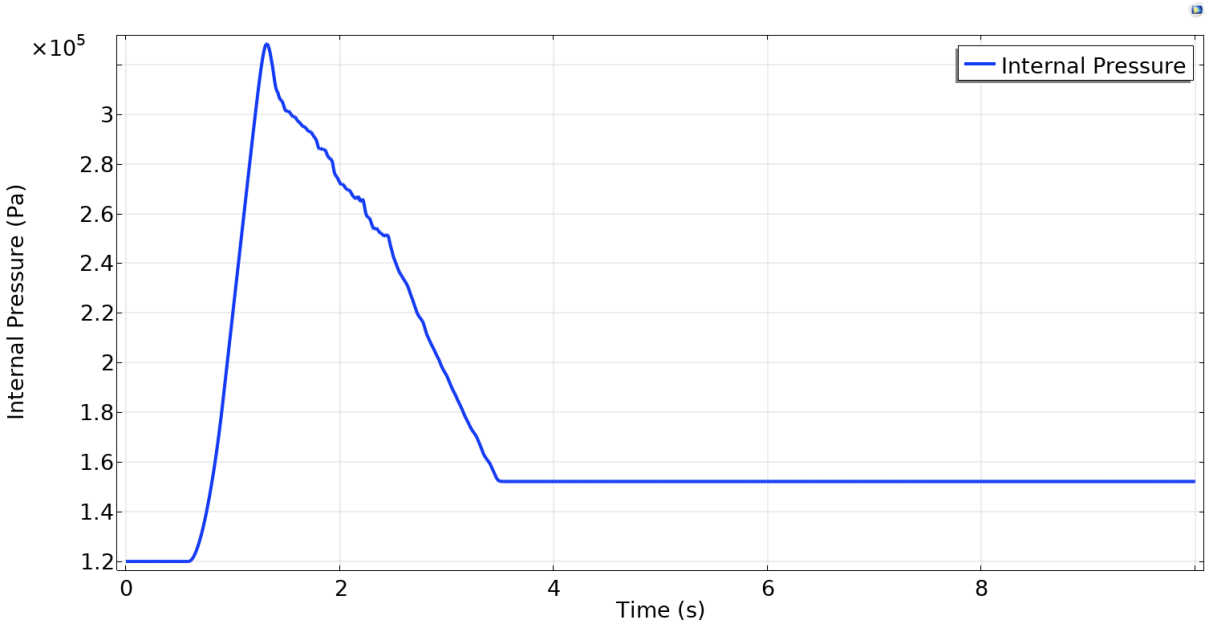

Figure 5.3: Time evolution of the internal loop pressure p_loop = p_0 + K · V_excess during the reference run. The pressure remains at the nominal p_0 = 1.2 bar until t ≈ 0.6 s, then rises rapidly to a peak of approximately 3.28 bar at t ≈ 1.32 s, blows down between t ≈ 1.3 s and t ≈ 3.5 s, and stabilises at a residual value of approximately 1.52 bar, approximately 20 mbar above the design closing threshold p_close = 1.5 bar.

A peak pressure of p_peak ≈ 3.28 × 10⁵ Pa ≈ 3.28 bar is reached at t ≈ 1.32 s, corresponding to an overshoot of approximately 83 % above the opening threshold p_open. The origin of this overshoot is the finite opening time τ_vent = 0.1 s of the module valve (Section 3.4.3): when p_loop crosses p_open at t ≈ 0.9 s the valve begins to open smoothly through the flc2hs gate, but the ejection rate through the module vent is initially zero and ramps up over an interval of the order of τ_vent. During this ramp, the gas injection at the cell vent continues at the nominal rate Q_inj ≈ 3.93 L/s and V_excess keeps accumulating until the ejection rate matches the injection rate; only at this point does dV_excess/dt vanish and p_loop reach its peak. The overshoot is therefore a direct consequence of the mismatch between the fixed valve time constant and the rate of gas injection during the opening ramp: a faster-acting valve (smaller τ_vent) would produce a lower overshoot, while a larger τ_vent would produce a higher one. The specific value of the peak p_loop depends on both the module-vent size (through the ejection rate for a given pressure difference) and on the closed-loop stiffness K (through the pressure–volume relation p_loop = p_0 + K · V_excess); the sensitivity of the peak to these two parameters is investigated in Section 5.2.2 and Section

5.2.3 respectively. From t ≈ 1.3 s onward the ejection rate at the module vent exceeds the injection rate at the cell vent, and V_excess (and therefore p_loop) begins to decrease monotonically. The blowdown phase extends from t ≈ 1.3 s to t ≈ 3.5 s, during which the pressure decreases from p_peak to values close to p_close = 1.5 bar. The spatial distribution of p_loop at six representative timestamps of this evolution is shown in Figure 5.4.

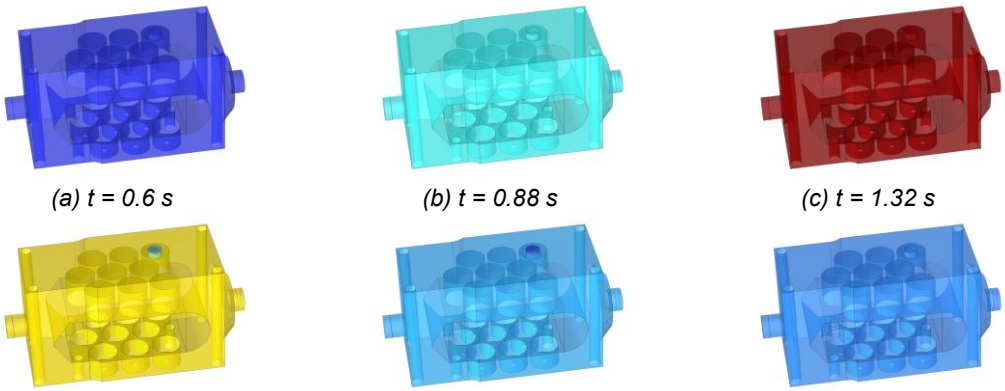

(a) t = 0.6 s *(b) t = 0.88 s* *(c) t = 1.32 s*

(d) t = 2.36 s *(e) t = 3.2 s* *(f) t = 10 s*

Figure 5.4: Loop pressure p_loop in the 3D module at six representative timestamps of the reference run: (a) t = 0.6 s, before cell-vent activation; (b) t = 0.88 s, module-vent about to open; (c) t = 1.32 s, peak loop pressure; (d) t = 2.36 s, blowdown phase with module vent fully open; (e) t = 3.2 s, module vent closing; (f) t = 10 s, end of the simulation. The colour scale is common to all six snapshots and extends from 1 × 10^5 Pa (dark blue) to 3.5 × 10^5 Pa (dark red). The essentially uniform colour field of each snapshot reflects the lumped-compliance assumption of Section 3.3, whereby p_loop is treated as spatially uniform; the localised low-pressure spot near the module vent visible in (d) and (e) marks the pressure drop across the open valve during the blowdown phase.

At t ≈ 3.5 s the pressure has decreased sufficiently that the valve begins to close, and the loop pressure stabilises at a residual value p_steady ≈ 1.52 × 10^5 Pa ≈ 1.52 bar, approximately 20 mbar above the design closing threshold p_close = 1.5 bar. The residual pressure corresponds to a residual excess volume V_excess,final = (p_steady − p_0) / K ≈ 0.3 L trapped in the loop, consistent with the blowdown behaviour expected from a hysteretic PRV that resets at p_close > p_atm (API Standard 520, 2020): the design choice p_close > p_atm keeps the module partially pressurised even after the venting event, avoiding a full depressurisation of the coolant loop. The additional 20 mbar offset above the nominal p_close is a numerical consequence of the smooth regularisation of the valve-state ODE (ε = 5 kPa in the flc2hs gate): as p_loop approaches p_close from above, the closing gate G_close = flc2hs((p_close − p_loop)/1[Pa], 5000) is only partially active (G_close ≈ 0.15 at p_loop = p_close + 2 kPa), and the self-consistent equilibrium between the residual gas efflux and the closure of s stabilises at this offset.

5.1.5 Valve dynamics

Figure 5.5 shows the time evolution of the valve state variable s(t), which represents the fraction of the module vent that is effectively open, from s = 0 (fully closed) to s = 1 (fully open). From t = 0 to t ≈ 0.9 s the valve remains closed because the loop pressure is below p_open. At t ≈ 0.9 s the loop pressure crosses p_open and the valve opens rapidly: s transitions from

0 to values close to 1 in approximately 0.4 s, driven by the opening gate G_open = flc2hs((p_loop − p_open)/1[Pa], 5000) that saturates to unity as soon as p_loop rises well above p_open (Section 3.4.3). The intrinsic characteristic time of the opening ODE is τ_vent = 0.1 s, consistent with the actuation time of automotive-grade PRVs discussed in Section 4.5.1.

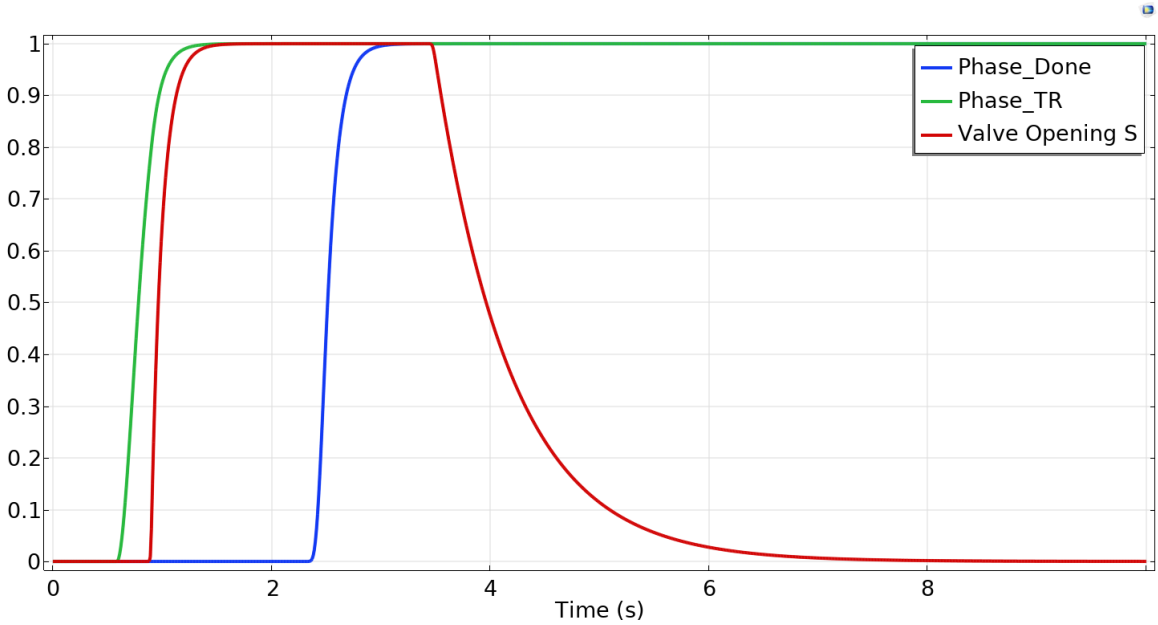

Figure 5.5: Time evolution of the valve state variable s (red), superimposed on the latch variables phase_TR (green) and phase_done (blue). The valve opens at t ≈ 0.9 s (s rises from 0 to 1 over approximately 0.4 s), remains fully open during the blowdown phase from t ≈ 1.3 s to t ≈ 3.5 s, then closes slowly with an effective time constant τ_eff ≈ 0.7 s that is significantly longer than the nominal τ_vent = 0.1 s. The phase_TR latch marks the activation of the cell vent at T_cell = T_open_cell (t ≈ 0.6 s); the phase_done latch marks the switch-off of the cell-vent mass flow rate at V_produced = V_max (t ≈ 2.5 s). The delay between the activation of the cell vent and the opening of the module vent (approximately 0.3 s), and between the end of the cell-vent injection and the beginning of the module-vent closure (approximately 1 s), are clearly visible.

The valve remains fully open (s ≈ 1) from t ≈ 1.3 s to t ≈ 3.5 s, spanning the blowdown transient of the loop pressure. During this interval the loop pressure stays within the hysteresis dead band (p_close, p_open) or above, so neither the opening gate nor the closing gate is significantly active and the value s = 1 is preserved by the ODE. No chattering is observed at any time during the open plateau, confirming that the hysteretic formulation of Section 3.4.3 correctly suppresses the numerical oscillations that a non-hysteretic threshold valve would produce.

At t ≈ 3.5 s the loop pressure has decreased to approximately p_close and the closing gate G_close = flc2hs((p_close − p_loop)/1[Pa], 5000) begins to activate. The subsequent decay of s is exponential but with an effective time constant τ_eff ≈ 0.7 s, significantly longer than the nominal τ_vent = 0.1 s. The reason has been discussed in Section 5.1.4: the loop pressure stabilises at approximately p_close + 20 mbar rather than exactly at p_close, so the closing gate evaluates to approximately G_close ≈ 0.15 rather than to its saturated value of unity

below p_close. The rate of closure ds/dt = − G_close · s / τ_vent is therefore of the order of − 1.5 · s, corresponding to a characteristic decay time τ_eff = τ_vent / G_close ≈ 0.7 s. This slow, self-consistent closure is not a chattering artefact of the model: the smooth regularisation of the flc2hs gate and the finite compliance of the loop conspire to leave the valve slightly open for several seconds before the transient extinguishes, and s → 0 asymptotically as t → 10 s. In a physical PRV with a mechanical spring the closure below p_close would be sharper; the slow-closure behaviour is a small consequence of the smoothed regularisation adopted here, discussed under the 'Lumped valve dynamics' bullet of Section 6.2.

5.1.6 Volume balance at the module vent

Figure 5.6 shows the time evolution of the three volumetric integrals of the closed-loop model: V_produced (cumulative gas injected at the cell vent), V_gas (residual gas volume inside the loop, i.e. V_produced minus the gas expelled through the module vent) and V_excess (net excess volume in the loop, equal to V_gas minus the volume of oil expelled through the module vent). The three quantities are governed by the coupled global ODEs of Section 4.6.

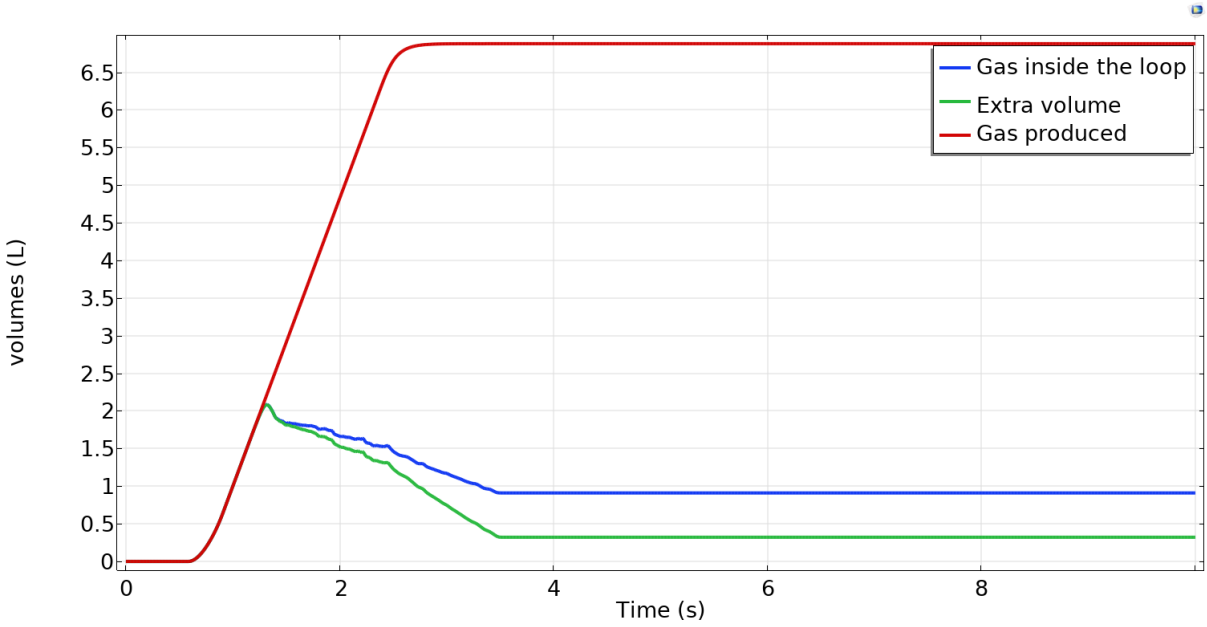

Figure 5.6: Time evolution of the volumetric integrals V_produced (red, cumulative gas injected at the cell vent), V_gas (blue, residual gas volume inside the closed loop) and V_excess (green, net excess volume in the loop). V_produced saturates at approximately 6.85 L when phase_done latches at t ≈ 2.6 s; V_gas reaches a peak of approximately 2.1 L at t ≈ 1.32 s and stabilises at approximately 0.9 L residual in the loop; V_excess stabilises at approximately 0.3 L residual. The difference V_gas − V_excess ≈ 0.6 L represents the volume of dielectric oil that has been expelled through the module vent together with the gas.

The cumulative injected volume V_produced rises linearly from t ≈ 0.6 s at the constant injection rate Q_inj ≈ 3.93 L/s prescribed at the cell vent during the active phase (Section 4.5.1), and saturates at V_produced,final ≈ 6.85 L at t ≈ 2.6 s, corresponding to the switch of phase_done from 0 to 1 through the smoothed flc2hs latch centred on V_max = 6.5 L: the final

value is approximately 5 % higher than V_max because of the finite width V_scale_phase = 8 × 10^{-5} m^3 of the latch transition, which allows the gas injection to continue smoothly for a short interval after V_produced has crossed V_max.

The residual gas volume inside the loop V_gas rises together with V_produced until t ≈ 1.3 s, at which point it reaches a maximum of approximately 2.1 L (the same instant at which the loop pressure reaches its peak, Section 5.1.4). From that time onward the ejection rate at the module vent exceeds the injection rate at the cell vent, and V_gas decreases quasi-linearly during the blowdown phase from t ≈ 1.3 s to t ≈ 3.5 s. When the valve begins to close, V_gas stabilises at a residual value of approximately 0.9 L still present in the closed loop, part of which remains visible under the ceiling of the 3D module (Section 5.1.3) and part of which has migrated into the external loop through the oil outlet.

The net excess volume V_excess follows a similar profile but with a lower final residual of approximately 0.3 L. The difference V_gas − V_excess ≈ 0.6 L represents the volume of dielectric oil that has been carried out of the module together with the gas during the open-valve interval. This is a substantial loss on the scale of the loop, corresponding to approximately 24 % of the total oil inventory of the closed loop (V_oil ≈ 2.5 L, Section 4.5.1). In a practical BTMS this oil would need to be replaced from the external reservoir to restore the operating configuration after the venting event.

The volume balance of the reference run can be summarised as follows: total gas produced by the cell V_produced,final = 6.85 L (at the injection temperature T_inlet_gas = 1000 K); gas expelled through the module vent V_produced − V_gas ≈ 5.95 L, corresponding to a gas removal efficiency of approximately 87 %; gas still inside the closed loop at t = 10 s V_gas ≈ 0.9 L (approximately 13 % of V_produced), of which only a fraction remains inside the 3D module domain (trapped under the ceiling by buoyancy, see Section 5.1.3) while the remainder has been transported into the external portion of the closed loop through the oil outlet; oil expelled with the gas V_gas − V_excess ≈ 0.6 L. The gas removal efficiency of 87 % obtained in the reference configuration is one of the primary output metrics of the module design, and its sensitivity to the vent diameter, position and closed-loop stiffness is investigated in the parametric study of Section 5.2.

5.2 Parametric study

5.2.1 Sweep design

The parametric study is organised as a sequential exploration of the main design variables of the venting system, in which each parameter is varied in turn while the others are held at their baseline values listed in Section 4.5.1. The order of the studies has been chosen so that each step first exercises the parameters with the strongest influence on the pressure dynamics of the closed loop, and then addresses the geometric layout of the vent on the ceiling: the diameter of the module-level vent is studied first, the closed-loop hydraulic stiffness K is studied second, and the position of the vent on the ceiling combined with the location of the TR cell within the array is quantified last. In each sweep the other design variables are held at

their baseline values (Section 5.1.1). This sequential approach is preferred to a full factorial design because the latter would require an excessive number of simulations given the computational cost of a single 3D run. Each sweep is exercised on the 3D model documented in Sections 4.2–4.9, at u_gas_peak = 50 m/s and with the other parameters held constant at the values of Section 4.5.1. The three sweeps and the corresponding ranges are listed below.

- D_module ∈ {5, 10, 15} mm — three values that span a factor of three in vent diameter and a factor of nine in vent area. The intermediate value D_module = 10 mm corresponds to the baseline configuration of the reference simulation (Section 5.1); the smaller and larger values explore the engineering trade-off between a faster pressure relief (larger vent → lower peak loop pressure, larger gas mass flow rate at opening) and a larger oil loss during the blow-down phase (larger vent → larger fraction of oil expelled with the gas before the valve reseats), and between a slower pressure relief (smaller vent → higher peak loop pressure but reduced oil losses).
- Closed-loop stiffness $K \in \{5 \times 10^7, 1 \times 10^8, 1 \times 10^9\}$ Pa/m³ — spans about two decades and corresponds to critical excess volumes V_excess_crit = (p_open − p_0)/K of approximately 1.2 L, 0.6 L and 0.06 L respectively. The range explores configurations in which the valve opens after a substantial fraction of V_max has been injected (low K, soft loop) and configurations in which the valve trips almost immediately (high K, stiff loop).
- Vent position on the module ceiling and TR-cell location — the effect of the vent position was investigated for a range of failure scenarios. Two vent positions were considered (A: baseline; B: displaced), and the TR-cell locations spanned a range from best-cooled to worst-cooled. For position A, three sub-studies were performed with three different TR-cell locations; for position B, only the two extreme TR-cell locations were considered, as the intermediate case is expected to fall between the two extremes and does not add information relevant to the design decision on the vent position. The sub-study therefore consists of five runs in total (three at position A, two at position B). This sub-study addresses the trapped-gas phenomenon observed in the 2D verification (Section 4.1.5), in which a substantial fraction of the gas remains under the ceiling in the dead zones of the recirculating oil flow.

The cell-level parameters (u_gas_peak, V_max, q_peak, T_open_cell) are held fixed at the baseline values of Section 4.5.1 throughout the parametric study, because they govern the integrated gas inventory and the timing of the TR trigger rather than the venting dynamics that are the focus of the present analysis.

For each simulation of each study, the same four output quantities are extracted: the peak loop pressure p_loop reached during the venting transient, the open time of the safety valve, the gas removal efficiency $\eta = (V_{produced} - V_{gas,final}) / V_{produced}$ and the residual gas distribution in the module at the end of the simulation. The aggregate of these results feeds the design discussion in Section 5.3.

5.2.2 Sensitivity to module-vent diameter

The parametric study on the module-vent diameter comprises three runs with D_module ∈ {5, 10, 15} mm; all other parameters are held at their reference values (Section 5.1.1). The three runs cover a factor of three in diameter and a factor of nine in vent area, chosen to bracket the range of realistic BTMS designs.

The loop pressure shows the clearest sensitivity to D_module (Figure 5.7). The reference run (Figure 5.7b, D_module = 10 mm) reaches a peak of 3.28 bar at t ≈ 1.32 s, after which the blowdown reduces the pressure back to p_close within approximately 2.2 s. With D_module = 5 mm (Figure 5.7a) the vent is undersized: the ejection rate through the vent is comparable to the injection rate at the cell vent, so the internal pressure keeps rising even after the valve opens, reaches a delayed peak of 3.68 bar at t ≈ 2.40 s (a 12 % increase over the reference), and takes approximately 5.5 s to blow down to p_close (reached at t ≈ 7.9 s, Figure 5.7a). With D_module = 15 mm (Figure 5.7c) the vent is oversized, and the peak of 3.19 bar (only 3 % below the reference) is reached at t ≈ 1.29 s, essentially simultaneously with the reference case; the blowdown is significantly faster, with p_loop reaching p_close in approximately 1.6 s. The physical mechanism of these differences follows the general overshoot description of Section 5.1.4. During the finite opening time τ_vent, the gas injection at the cell vent continues at the nominal rate Q_inj ≈ 3.93 L/s and V_excess keeps accumulating; once the valve is fully open, the ejection rate through the module vent for a given driving pressure difference scales with the vent cross-section area, and therefore with D_module². With D_module = 5 mm the effective vent area is a quarter of the baseline, so the ejection rate remains below the injection rate for a much longer interval, V_excess keeps accumulating and the peak is both delayed and much higher; with D_module = 15 mm the vent area is 2.25 times the baseline and the ejection rate matches the injection rate very quickly, so the peak is only marginally lower than the reference. The scaling of the peak with D_module is therefore not linear but saturates: increasing the vent area to 2.25 times the baseline produces only a marginal reduction of the peak (from 3.28 to 3.19 bar), while reducing it to a quarter of the baseline produces a substantial increase (from 3.28 to 3.68 bar) and a very slow relaxation. The residual pressure at t = 10 s is very close but not identical across the three runs: p_steady = 1.526 bar for D_module = 5 mm, 1.522 bar for D_module = 10 mm and 1.532 bar for D_module = 15 mm. All three values lie approximately 20–30 mbar above the design threshold p_close = 1.5 bar, consistent with the self-consistent equilibrium of the closure ODE discussed in Section 5.1.4. The small differences among the three runs — approximately 10 mbar between the minimum and the maximum — are not numerical noise but a real consequence of the flc2hs regularisation of the valve closure: the offset above p_close depends on the rate at which p_loop crosses p_close during the blowdown, which is set by the ejection rate through the vent and therefore by D_module. The 15 mm case, in which the pressure crosses p_close very rapidly (dp/dt ≈ −1.5 bar/s), leaves a slightly larger residual V_excess through the self-consistent feedback discussed below; the 10 mm case, with an intermediate rate, sits at the minimum of the offset. The small residual is inherited by V_excess (0.326, 0.322 and 0.332 L

respectively for D_module = 5, 10 and 15 mm), which would ideally equal V_excess,steady = (p_close − p_0) / K = 0.3 L in the absence of the flc2hs smoothing.

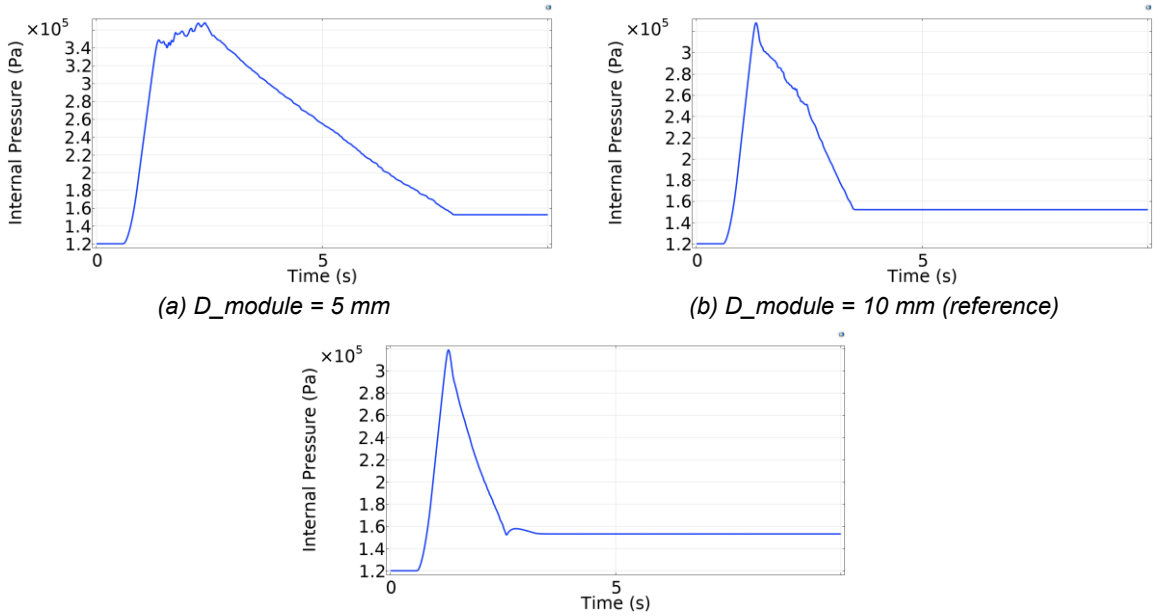

(a) D_module = 5 mm

(b) D_module = 10 mm (reference)

(c) D_module = 15 mm

Figure 5.7: Time evolution of the internal loop pressure p_loop for the three sweep runs on D_module. (a) D_module = 5 mm: undersized vent, delayed peak at approximately 3.68 bar at t ≈ 2.40 s, blowdown extending to t ≈ 7.9 s. (b) D_module = 10 mm: baseline reference, peak at 3.28 bar at t ≈ 1.32 s, blowdown to p_close by t ≈ 3.5 s. (c) D_module = 15 mm: oversized vent, peak at 3.19 bar at t ≈ 1.29 s, fast blowdown to p_close by t ≈ 2.6 s.

The valve state variable shows a striking spread in the open-vent time (Figure 5.8). All three runs open the valve at approximately the same instant (t ≈ 0.97 s), because the opening trigger depends on the accumulation of V_excess which is dominated by the cell-vent injection common to all three runs. The closing dynamics however diverge: with D_module = 10 mm (Figure 5.8b) the valve closes by t ≈ 4 s; with D_module = 15 mm (Figure 5.8c) by t ≈ 4.7 s; with D_module = 5 mm (Figure 5.8a) by t ≈ 8.7 s. The vent open interval (defined as the range during which s ≥ 0.5) is therefore approximately 3.0, 3.7 and 7.7 s respectively. The 5 mm case is not just quantitatively different but qualitatively pathological: the valve remains open for more than seven seconds, keeping the module continuously connected to the atmosphere and dissipating pressure inefficiently. A notable feature of the D_module = 15 mm run (Figure 5.8c) is a small oscillation of the valve state around t ≈ 2.6 s (visible as a brief plateau near s ≈ 0.96). This is a self-consistent feedback of the closure ODE when the pressure crosses p_close rapidly: the larger vent produces a very fast pressure decrease (dp/dt ≈ −1.5 bar/s at the crossing), which triggers a sharp activation of the closing gate; the resulting quick reduction of s cuts the ejection rate, causing V_excess and p_loop to overshoot p_close by a small amount, which halts the closure at s ≈ 0.96 for a fraction of a second before the transient dies out. The oscillation is small in amplitude and does not affect the steady state, but it is a warning that oversized vents can produce marginally underdamped valve dynamics.

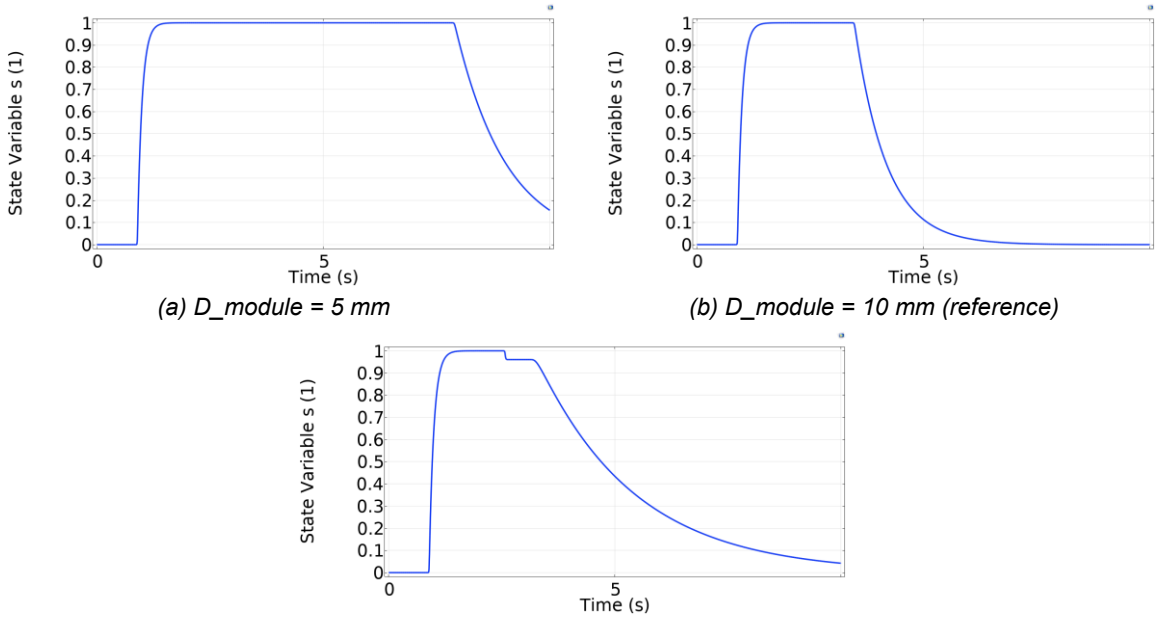

(a) D_module = 5 mm

(b) D_module = 10 mm (reference)

(c) D_module = 15 mm

Figure 5.8: Time evolution of the valve state variable s for the three sweep runs on D_module. All three runs open the valve at approximately the same instant (t ≈ 0.97 s); the closing dynamics diverge strongly. (a) D_module = 5 mm: valve fully open for approximately 7.7 s, still not fully closed at t = 10 s. (b) D_module = 10 mm: baseline reference, valve open for approximately 3.0 s (s ≥ 0.5), slow self-consistent closure with τ_eff ≈ 0.7 s. (c) D_module = 15 mm: valve open for approximately 3.7 s with a small underdamped oscillation visible as a brief plateau at s ≈ 0.96 around t ≈ 2.6 s, discussed in the text.

The volumetric integrals V_gas and V_excess (Figure 5.9) complete the picture. The cumulative gas produced V_produced is essentially identical across the three runs at V_produced,final ≈ 6.88 L, as expected because V_max is a property of the TR cell and not of the vent geometry, and is therefore not shown. The residual gas volume in the loop at t = 10 s is V_gas ≈ 0.94 L for D_module = 5 mm (Figure 5.9a), 0.91 L for the reference D_module = 10 mm (Figure 5.9b) and 0.99 L for D_module = 15 mm (Figure 5.9c), corresponding to a gas removal efficiency of approximately 86 %, 87 % and 85 %. The oil expelled with the gas is V_gas − V_excess ≈ 0.62, 0.59 and 0.66 L for the three runs. Overall, the reference D_module = 10 mm sits at the sweet spot of the trade-off: a smaller vent penalises the pressure dynamics (higher peak, much longer relaxation) without improving the volume balance; a larger vent marginally reduces the peak pressure but increases the oil loss and can introduce underdamped valve dynamics. The parametric sub-study therefore confirms the choice of D_module = 10 mm as a suitable baseline for the module.

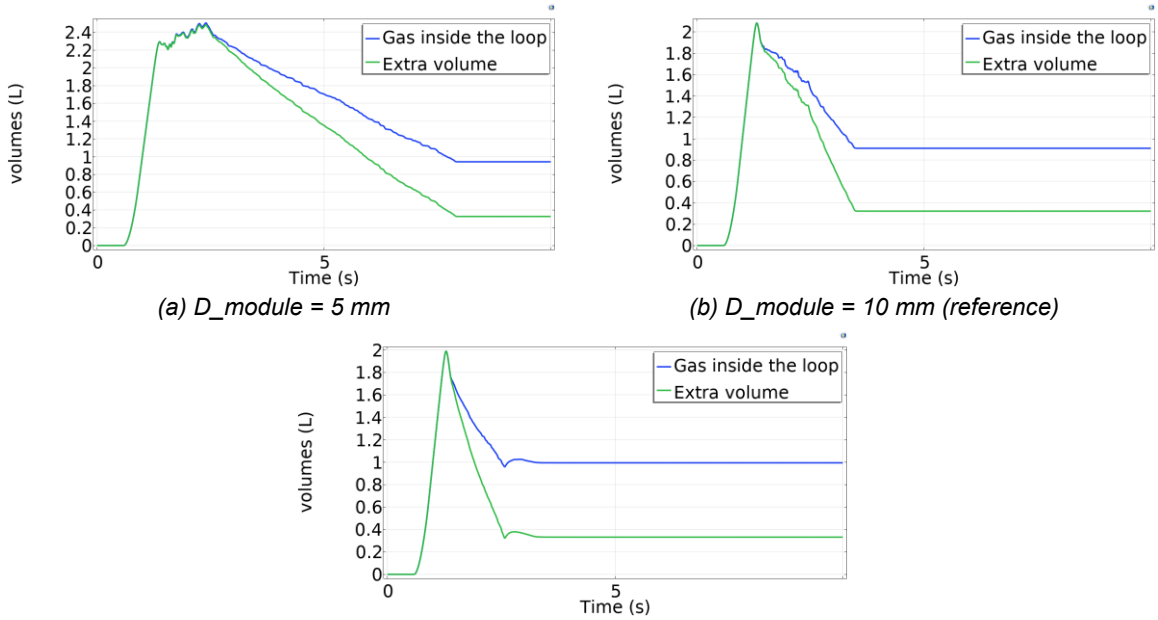

(a) D_module = 5 mm

(b) D_module = 10 mm (reference)

(c) D_module = 15 mm

Figure 5.9: Time evolution of the volumetric integrals V_gas (blue) and V_excess (green) for the three sweep runs on D_module. The residual gas in the loop at t = 10 s is V_gas ≈ 0.94 L for (a) D_module = 5 mm, 0.91 L for (b) D_module = 10 mm and 0.99 L for (c) D_module = 15 mm; the residual V_excess is 0.326, 0.322 and 0.332 L respectively. The oil expelled through the module vent (V_gas − V_excess) is approximately 0.62, 0.59 and 0.66 L for the three runs.

5.2.3 Sensitivity to closed-loop stiffness K

The lumped hydraulic stiffness K represents the inverse of the hydraulic capacitance C_h of the closed loop external to the module, following the classical hydraulic analogy that models fluid compressibility, plumbing expansion and entrapped air as equivalent capacitors in the pressure–flow domain [Merritt, 1967; Rabie, 2009]. For a closed loop with fluid volume V_fluid and effective bulk modulus β_eff, the hydraulic capacitance is C_h = V_fluid / β_eff and the stiffness is K = 1 / C_h; the same lumped-parameter formulation is adopted, for example, by Falcone et al. (2021) for the closed cooling loop of a hybrid-train battery thermal management system, on which the model of Chapter 3 is conceptually based. The effective bulk modulus of the system combines the bulk modulus of the fluid and the compliance of the plumbing and of any expansion accumulator present in the loop. Four sources of hydraulic capacitance are recognised in the literature: fluid compressibility, plumbing or component expansion under pressure, air entrapped in the hydraulic fluid, and non-rigidity of the machine structure [Rabie, 2009]. For synthetic hydrocarbon dielectric fluids at ambient temperature, the bulk modulus is in the range $\beta_{oil} \approx 1.5–2.0 \times 10^9$ Pa (typical values for polyalphaolefins and refined mineral hydraulic oils), similar to that of water-based hydraulic fluids ($\beta \approx 2.1 \times 10^9$ Pa) [Kambic and Kalb, 2014] and higher than that of silicone oils.

A closed loop with a fluid volume of a few litres of AmpCool AC-110 synthetic hydrocarbon oil and no expansion accumulator would therefore have $K = 1/(V_{oil} \cdot \beta_{oil}) \approx 6 \times 10^{11}$ Pa/m³

for V_{oil} = 2.5 L, corresponding to a very stiff response of the pressure to any accumulated gas volume. In real cooling loops, however, a small volume of trapped air or a dedicated expansion accumulator is invariably present, and dominates the effective bulk modulus of the system: the equivalent bulk modulus of a small entrapped gas pocket at pressure p_0 is $\gamma \cdot p_0$ (with $\gamma = 1.4$ for air), which for $p_0 = 1.2$ bar is only 1.7×10^5 Pa, six orders of magnitude below the fluid bulk modulus. With an equivalent air volume V_{air} in the range of 100 mL to a few litres, which is representative of an automotive battery thermal management loop with either a dedicated diaphragm accumulator or residual entrapped air, the resulting K falls in the range 10^8–10^9 Pa/m³. The baseline value $K = 10^8$ Pa/m³ adopted in the reference case corresponds to an equivalent air volume of approximately 1 L, in line with typical expansion volumes of BTMS accumulators.

The sensitivity sweep on K spans a factor of twenty around the baseline, from $K = 5 \times 10^7$ Pa/m³ (loop with a very compliant expansion vessel of the order of 2 L of equivalent air) to $K = 10^9$ Pa/m³ (compact loop with only ~100 mL of residual air, at the limit of a stiff-loop configuration). The corresponding critical excess volumes $V_{excess,crit} = (p_{open} - p_0) / K$ required to trigger the module vent are 1.2 L for $K = 5 \times 10^7$ Pa/m³, 0.6 L for the baseline $K = 10^8$ Pa/m³, and 60 mL for $K = 10^9$ Pa/m³, which cover the design space between a loop that opens the valve only after a substantial fraction of the cell gas inventory has been injected (compliant loop) and a loop that trips the valve almost immediately after the cell vent activation (stiff loop). Values above 10^9 Pa/m³ approach the fully rigid loop and are not explored, since the resulting valve dynamics collapse to instantaneous opening.

The parametric study on the closed-loop stiffness K comprises three runs with $K \in \{5 \times 10^7, 10^8, 10^9\}$ Pa/m³, spanning a factor of twenty between the softest and the stiffest loop configurations. All other parameters are held at their reference values (Section 5.1.1).

The loop pressure shows a very large spread across the three runs (Figure 5.10). The soft loop (Figure 5.10a, $K = 5 \times 10^7$ Pa/m³) reaches a peak $p_{peak} \approx 2.58$ bar at $t \approx 1.5$ s, followed by a slow blowdown to p_{close} by $t \approx 4.2$ s: the peak is approximately 21 % lower than the reference run (Figure 5.10b, baseline $p_{peak} = 3.28$ bar). The stiff loop (Figure 5.10c, $K = 10^9$ Pa/m³) is dramatically different: p_{loop} reaches a peak of approximately 11.7 bar at $t \approx 1.05$ s, more than three times the baseline value, and the blowdown to p_{close} is extremely fast (approximately 1.4 s) and includes visible pressure oscillations between $t \approx 2$ s and $t \approx 2.5$ s. The physical mechanism of the K-dependence of the peak, referring to the general overshoot description of Section 5.1.4, is the following. When p_{loop} crosses p_{open}, the valve requires the same fixed opening time $\tau_{vent} = 0.1$ s across the three runs to reach $s \approx 1$, during which the gas injection at the cell vent continues at the nominal rate $Q_{inj} \approx 3.93$ L/s. The additional excess volume accumulated during this opening ramp is approximately $Q_{inj} \cdot \tau_{vent} \approx 0.4$ L and is essentially the same across the three runs, because τ_{vent} does not depend on K. The pressure rise $\Delta p = K \cdot \Delta V_{excess}$ produced by this residual accumulation, however, scales linearly with K: it is approximately 0.2 bar for $K = 5 \times 10^7$ Pa/m³, 0.4 bar for the baseline, and 4 bar for $K = 10^9$ Pa/m³. In the stiff configuration this is the dominant contribution to the peak:

the loop cannot open the valve fast enough to prevent a large pressure spike, and the resulting peak p_peak ≈ 11.7 bar is more than three times the baseline. Once the valve is fully open, the excess pressure of the loop over the atmospheric side drives the gas out through the module vent at a rate that increases with the pressure difference, so the very high peak pressure of the stiff configuration also produces a very fast blowdown, as visible in Figure 5.10c. This is a fundamental design consideration for the closed-loop cooling system: a loop with excessive stiffness produces internal pressures that can easily exceed the burst pressure of automotive-grade plumbing (typically 5–10 bar for cooling loops), potentially causing catastrophic failure of the coolant circuit during a thermal runaway event. A stiff cooling loop therefore requires either a faster-acting pressure-relief valve (smaller τ_vent) or a lower opening threshold p_open, in order to keep the peak pressure within the safe operating range of the coolant plumbing.

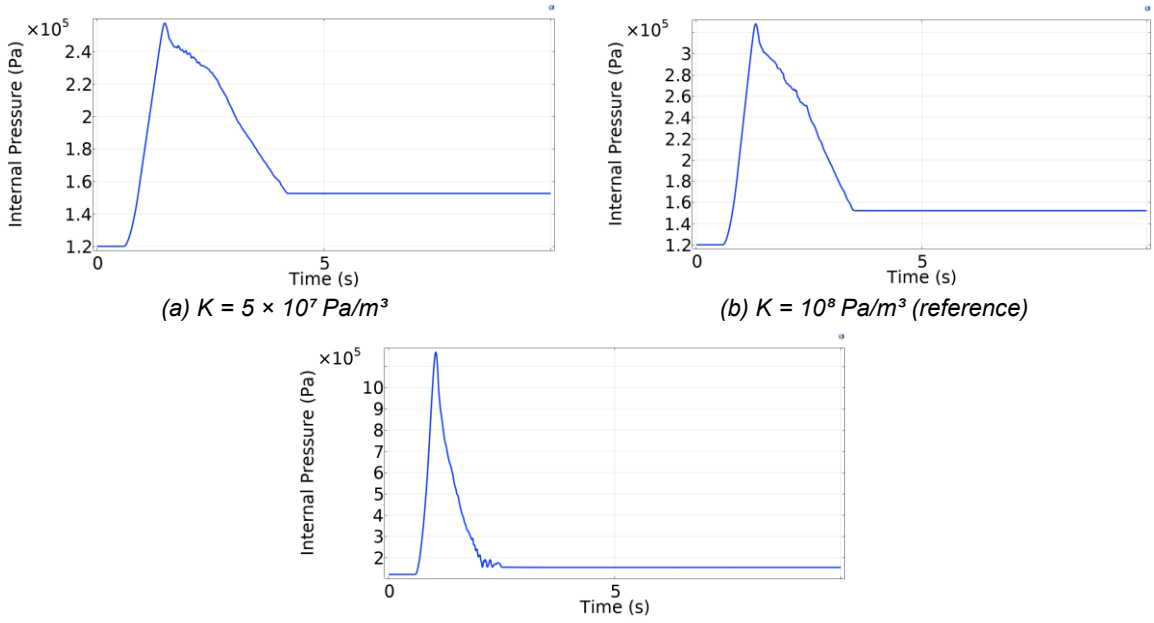

(a) K = 5 × 10⁷ Pa/m³ *(b) K = 10⁸ Pa/m³ (reference)*

(c) K = 10⁹ Pa/m³

Figure 5.10: Time evolution of the internal loop pressure p_loop for the three sweep runs on K. (a) K = 5 × 10⁷ Pa/m³: soft loop, moderate peak of approximately 2.58 bar at t ≈ 1.5 s, slow blowdown to p_close by t ≈ 4.2 s. (b) K = 10⁸ Pa/m³: baseline reference, peak at 3.28 bar at t ≈ 1.32 s. (c) K = 10⁹ Pa/m³: stiff loop, dramatic peak of approximately 11.7 bar at t ≈ 1.05 s (more than three times the baseline), very fast blowdown with visible pressure oscillations between t ≈ 2 s and t ≈ 2.5 s.

The valve state variable shows a spread both in the opening instant and in the closing dynamics (Figure 5.11). The stiff loop (Figure 5.11c, K = 10⁹ Pa/m³) opens the valve earliest (t ≈ 0.7 s), because the pressure crosses p_open with only V_excess ≈ 60 mL of accumulated gas — a fraction of what is needed in the softer configurations. The reference case (Figure 5.11b) opens at t ≈ 0.9 s and the soft loop (Figure 5.11a) at t ≈ 1.0 s. The closing dynamics are somewhat counter-intuitive: the soft loop keeps the valve fully open for the longest interval (approximately 3.2 s of plateau at s = 1), because p_loop varies slowly with V_excess and takes longer to cross p_close from above; the stiff loop, despite the very high peak pressure, terminates its s = 1 plateau within approximately 2.5 s because the small V_excess reservoir

empties quickly through the vent, causing a rapid pressure drop through the hysteresis dead band. A small underdamped oscillation of s is visible in the stiff configuration between t ≈ 2 s and t ≈ 2.5 s (Figure 5.11c), similar to the one observed in the D_module = 15 mm case of Section 5.2.2 and generated by the same mechanism: a rapid pressure crossing of p_close triggers a sharp activation of the closing gate that leaves a small oscillatory transient before the valve fully closes.

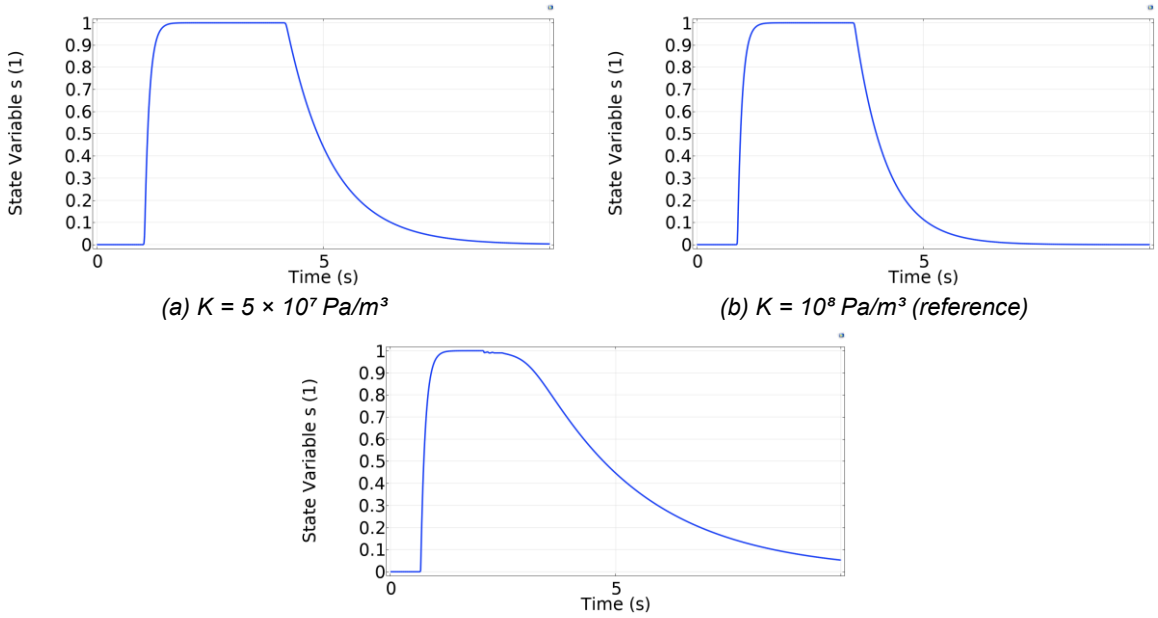

(a) K = 5 × 10⁷ Pa/m³ *(b) K = 10⁸ Pa/m³ (reference)*

(c) K = 10⁹ Pa/m³

Figure 5.11: Time evolution of the valve state variable s for the three sweep runs on K. (a) K = 5 × 10⁷ Pa/m³: latest valve opening (t ≈ 1.0 s) but longest open plateau. (b) K = 10⁸ Pa/m³: baseline reference, opening at t ≈ 0.9 s, plateau from t ≈ 1.3 s to t ≈ 3.5 s. (c) K = 10⁹ Pa/m³: earliest valve opening (t ≈ 0.7 s), small underdamped oscillation between t ≈ 2 s and t ≈ 2.5 s, slow closure extending to t = 10 s.

The volumetric integrals V_gas and V_excess (Figure 5.12) reveal the trade-off between the peak pressure and the gas removal efficiency. The residual gas volume in the loop at t = 10 s is V_gas ≈ 1.42 L for (a) K = 5 × 10⁷ Pa/m³, 0.91 L for (b) K = 10⁸ Pa/m³ and 0.44 L for (c) K = 10⁹ Pa/m³, corresponding to gas removal efficiencies of approximately 79 %, 87 % and 94 % respectively. A higher stiffness therefore expels more of the produced gas but at the cost of dramatically higher peak pressures. The residual V_excess is 0.65 L, 0.32 L and 0.03 L respectively; these values are close to the design expectation V_excess,steady = (p_close − p_0) / K = 1.2 L, 0.3 L and 0.06 L, with the soft loop showing the largest deviation because its self-consistent equilibrium is reached asymptotically and 10 s is not sufficient to fully drain the loop. The oil expelled with the gas (V_gas − V_excess) is approximately 0.77, 0.59 and 0.41 L for the three runs, decreasing with increasing K: a stiffer loop expels less oil in absolute terms because the total V_excess to be blown down is smaller. The K = 10⁹ Pa/m³ case therefore offers the best mass balance and the best gas removal efficiency, but the peak loop pressure of 11.7 bar makes it unsuitable for automotive applications; the K = 5 × 10⁷ Pa/m³ case is safe from the pressure point of view but leaves a substantial residual gas volume in

the loop. The reference K = 10⁸ Pa/m³ is a reasonable compromise, and the sub-study of this section highlights the need for a design methodology that couples the sizing of the cooling-loop capacitance (expansion vessel, air cushion) with the sizing of the pressure-relief valve, in order to prevent both catastrophic pressure spikes and inefficient gas removal in the event of a thermal runaway.

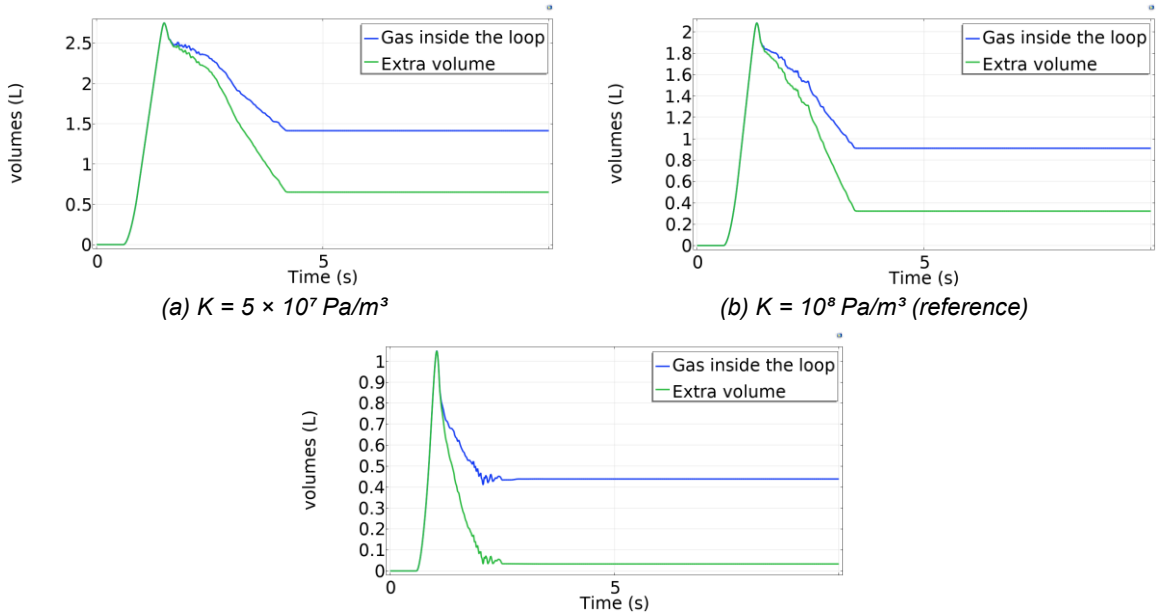

(a) K = 5 × 10⁷ Pa/m³

(b) K = 10⁸ Pa/m³ (reference)

(c) K = 10⁹ Pa/m³

Figure 5.12: Time evolution of the volumetric integrals V_gas (blue) and V_excess (green) for the three sweep runs on K. The residual gas in the loop at t = 10 s is V_gas ≈ 1.42 L for (a) K = 5 × 10⁷ Pa/m³, 0.91 L for (b) K = 10⁸ Pa/m³ and 0.44 L for (c) K = 10⁹ Pa/m³, corresponding to gas removal efficiencies of approximately 79 %, 87 % and 94 %. The oil expelled through the module vent (V_gas − V_excess) is approximately 0.77, 0.59 and 0.41 L for the three runs.

5.2.4 Effect of the vent position on the module ceiling for different locations of the TR cell

The effect of the vent position on the module ceiling was investigated for a set of different possible TR-cell locations inside the 3×3 grid. Two vent positions were considered: position A, corresponding to the baseline configuration of Section 5.1, and position B, displaced to a second location on the ceiling. The sub-study was designed to bracket the range of possible failure scenarios rather than to cover the full 2 × 3 factorial: for position A, three sub-studies were performed with three different cells driven into thermal runaway, so as to span the range from the best-cooled to the worst-cooled TR-cell location; for position B, only the two extreme TR-cell locations (the best-cooled and the worst-cooled) were considered, because the intermediate case (central TR cell) is expected to fall between the two extremes and does not add information relevant to the design decision on the vent position. The resulting simulation matrix consists of five runs in total: three at position A and two at position B. The run at position A with the reference TR cell coincides with the reference simulation of Section 5.1 and is not repeated, so the sub-study reported in this section requires four additional simulations.

The cells of the 3×3 array are numbered from [1] to [9] according to the top view of Figure 5.13: cells [1], [2] and [3] form the outlet-side column (nearest to the module outlet and aligned with the module vent in the baseline configuration), cells [4], [5] and [6] form the central column, and cells [7], [8] and [9] form the inlet-side column (nearest to the module inlet). The three horizontal rows are [7, 4, 1] (top), [8, 5, 2] (central) and [9, 6, 3] (bottom). Module-vent position A places the vent on the ceiling directly above cell [2] — the central row of the outlet-side column — corresponding to the baseline configuration of Section 5.1 (Figure 5.13a). Module-vent position B is displaced by 2 cm further towards the outlet along the longitudinal axis of the module (Figure 5.13b). The three TR-cell locations selected for the sub-study are cell [8] (central row of the inlet-side column), corresponding to the reference run of Section 5.1 and representing the worst-case scenario in terms of distance from the module vent; cell [3] (outlet-side column, bottom row), directly adjacent to vent-A position on the ceiling and representing the best-case scenario in terms of vent proximity; and cell [4] (central column, top row), an asymmetric intermediate case. Each run is identified by the pair (vent position, TR-cell), so that the reference run of Section 5.1 corresponds to (A, [8]) and the four additional runs of the sub-study are (A, [4]), (A, [3]), (B, [8]) and (B, [3]).

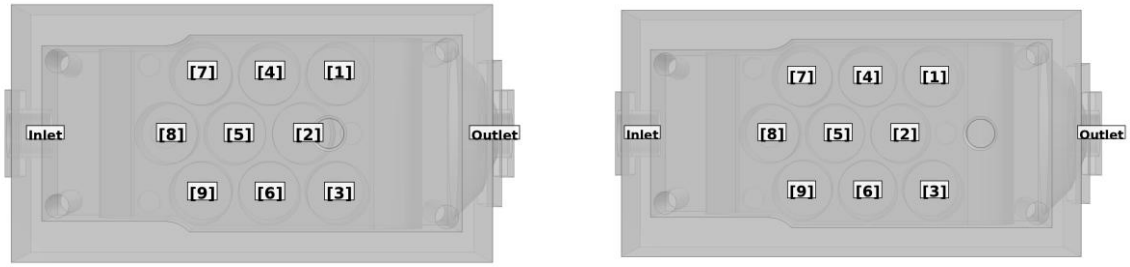

(a) Position A

(b) Position B

Figure 5.13: Top view of the module used for the vent-position and TR-cell-location sub-study, with the numbering of the nine cells and the module-vent position highlighted for the two configurations considered. (a) Position A: the module vent is placed on the ceiling directly above cell [2] (central row of the outlet-side column), corresponding to the baseline configuration of Section 5.1. (b) Position B: the module vent is displaced by 2 cm further towards the outlet along the longitudinal axis of the module.

The five runs of the sub-study are analysed in terms of the internal loop pressure p_loop, the valve state variable s and the lumped volumes V_gas and V_excess. The analysis focuses on how the buoyancy-driven migration of the injected gas interacts with the position of the module vent on the ceiling, and on how this interaction depends on the location of the TR event within the array.

The internal loop pressure exhibits only a marginal sensitivity to both design variables of this sub-study (Figure 5.14). All five runs reach essentially the same peak of approximately 3.28 bar at t ≈ 1.32 s (Figure 5.14a–e), which confirms the mechanism identified in Section 5.1.4: the overshoot is dictated by the finite valve opening time τ_vent and by the excess volume V_excess accumulated during the valve ramp — both quantities depend on the

injection rate Q_inj at the cell vent and on the closed-loop stiffness K, not on the position at which the gas enters the module nor on the location of the module vent on the ceiling. The subsequent blowdown exhibits small but systematic differences. At module-vent position A the reference run (A, [8]) closes the valve at t ≈ 3.5 s (Figure 5.14a), whereas (A, [4]) closes at t ≈ 3.3 s (Figure 5.14b) and (A, [3]) at t ≈ 3.1 s (Figure 5.14c): the closer the TR cell is to the outlet-side column, the shorter the blowdown. At module-vent position B the same trend is preserved: (B, [8]) closes at t ≈ 3.4 s (Figure 5.14d), while (B, [3]) exhibits the fastest blowdown of the sub-study at approximately 3.0 s (Figure 5.14e). Comparing pairs at fixed TR-cell location, the effect of the vent shift alone (A → B) is small — 3.5 s versus 3.4 s for cell [8], 3.1 s versus 3.0 s for cell [3] — but confirms that both moving the TR cell towards the outlet-side column and shifting the module vent towards the outlet lead to a slightly shorter blowdown. The pressure signal alone does not strongly discriminate between the five configurations; the sensitivity to the vent position and to the TR-cell location emerges more clearly in the volumetric analysis reported later in this section, where these small blowdown differences correlate directly with the residual gas fraction trapped in the loop.

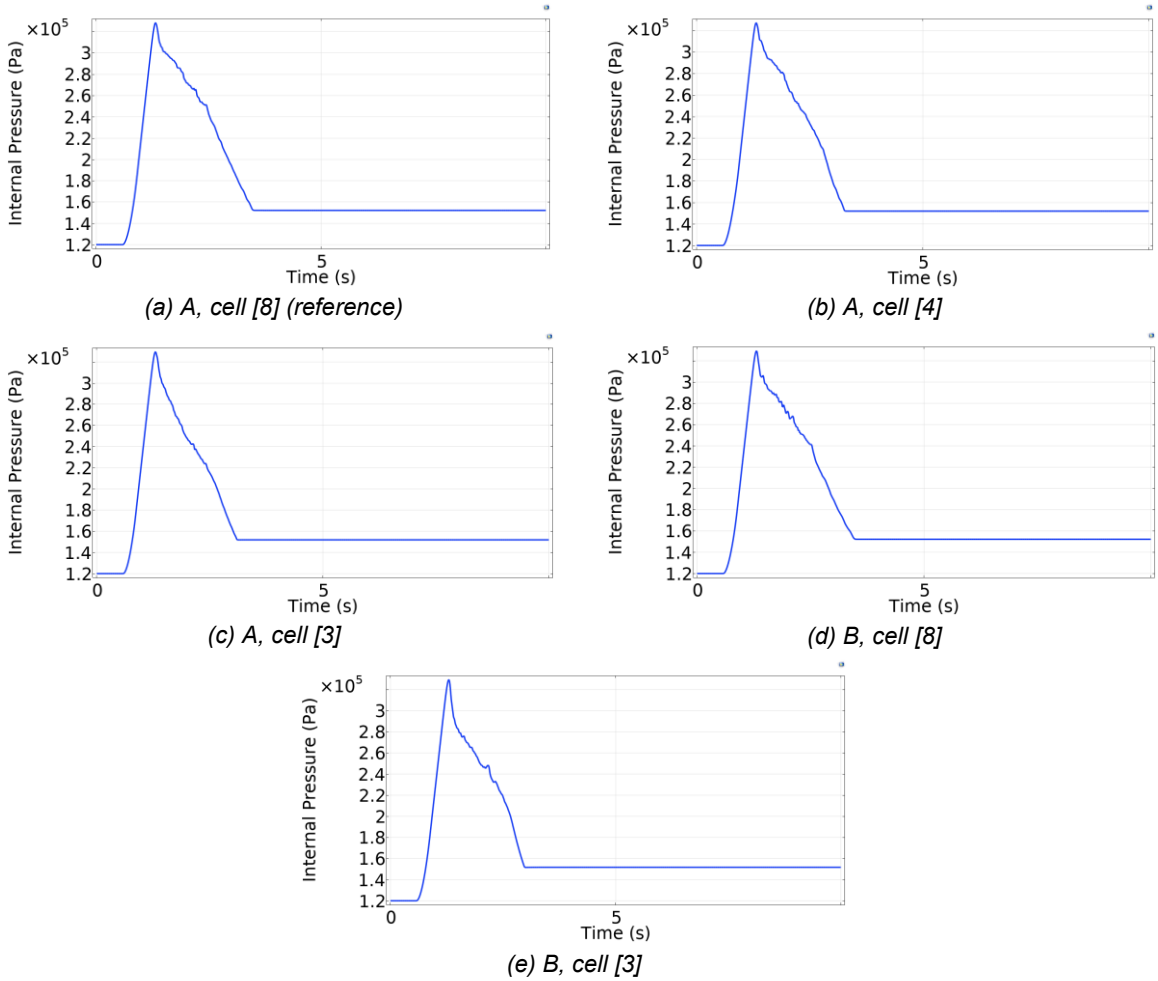

(a) A, cell [8] (reference)

(b) A, cell [4]

(c) A, cell [3]

(d) B, cell [8]

(e) B, cell [3]

Figure 5.14: Time evolution of the internal loop pressure p_loop for the five runs of the vent-position and TR-cell-location sub-study. (a) Position A, cell [8] in TR (reference): peak of approximately 3.28 bar at t ≈ 1.32 s, blowdown to p_close by t ≈ 3.5 s. (b) Position A, cell [4]: peak

3.26 bar, blowdown by t ≈ 3.3 s. (c) Position A, cell [3]: peak 3.30 bar, blowdown by t ≈ 3.1 s. (d) Position B, cell [8]: peak 3.30 bar, blowdown by t ≈ 3.4 s. (e) Position B, cell [3]: peak 3.30 bar, blowdown by t ≈ 3.0 s.

The evolution of the valve state variable s(t) is consistent with the loop pressure signal analysed above (Figure 5.15). All five runs open the valve at approximately the same instant (t ≈ 0.9 s), as expected from the p_open crossing that is essentially identical across the sub-study. The valve remains fully open (s ≈ 1) throughout the venting transient, and starts closing when p_loop crosses back through p_close: at t ≈ 3.5 s for (A, [8]) (Figure 5.15a), t ≈ 3.3 s for (A, [4]) (Figure 5.15b) and t ≈ 3.1 s for (A, [3]) (Figure 5.15c), reproducing the same trend already discussed for the blowdown of p_loop. At module-vent position B the same trend is preserved: (B, [8]) closes at t ≈ 3.4 s (Figure 5.15d), while (B, [3]) closes at t ≈ 3.0 s (Figure 5.15e), the shortest open interval of the sub-study. The subsequent decay of s follows the closing dynamics of the valve (Section 3.4.3) and is qualitatively identical in shape across the five runs: s drops to a residual value of order 10^{-2} within approximately 2 s of the closing onset in all configurations, with no signature of chattering or partial re-opening.

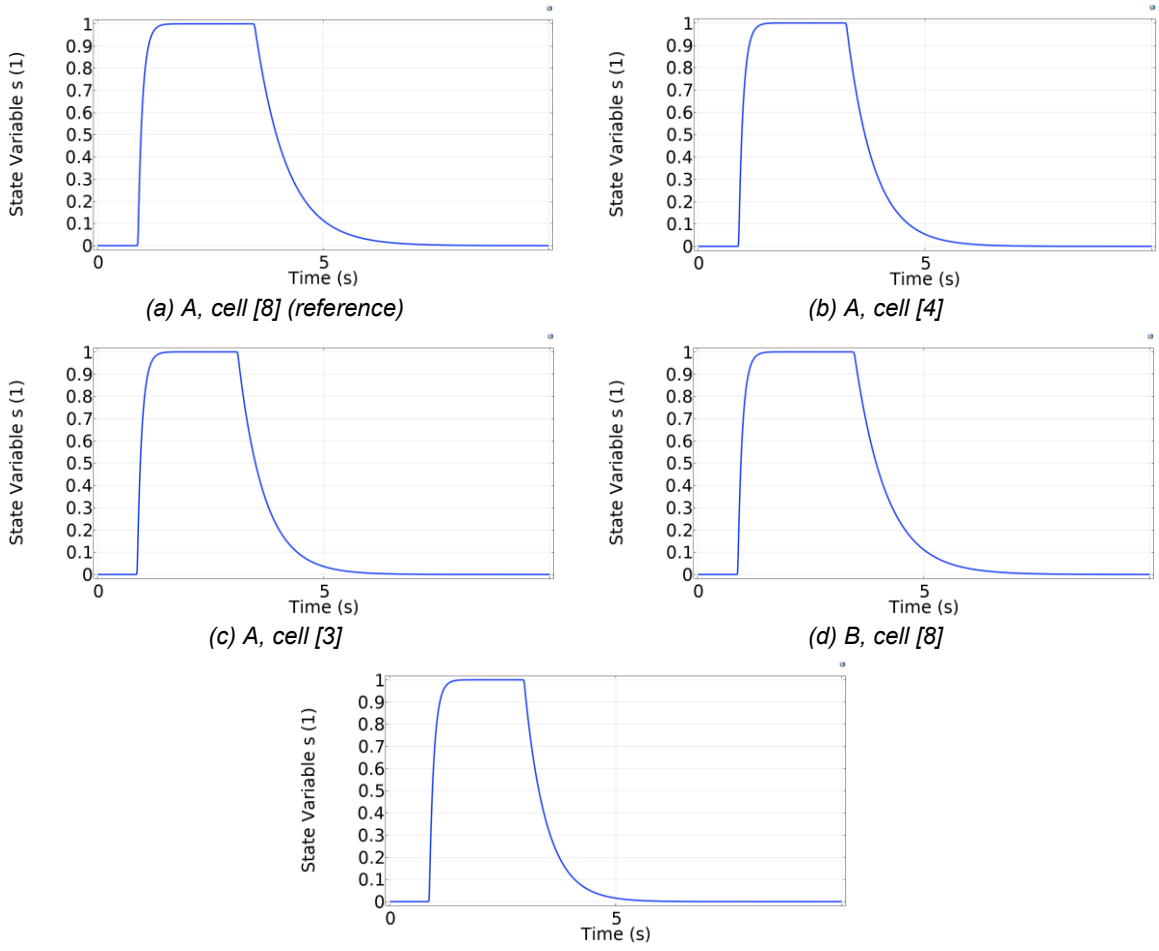

(a) A, cell [8] (reference)

(b) A, cell [4]

(c) A, cell [3]

(d) B, cell [8]

(e) B, cell [3]

Figure 5.15: Time evolution of the valve state variable s for the five runs of the vent-position and TR-cell-location sub-study. (a) Position A, cell [8] (reference): valve opens at t ≈ 0.9 s and closes at t ≈ 3.5 s. (b) Position A, cell [4]: opens at t ≈ 0.9 s, closes at t ≈ 3.3 s. (c) Position A, cell [3]: opens at t ≈

0.9 s, closes at t ≈ 3.1 s. (d) Position B, cell [8]: opens at t ≈ 0.9 s, closes at t ≈ 3.4 s. (e) Position B, cell [3]: opens at t ≈ 0.9 s, closes at t ≈ 3.0 s. All five runs share the same exponential decay of s after the closing onset, with a residual value of order 10^{-2} within approximately 2 s.

The volume balance completes the analysis of the sub-study and quantifies the buoyancy-dominated trapped-gas phenomenon anticipated in the 2D verification (Section 4.1.5) and now resolved in 3D (Figure 5.16). Because the injection conditions are identical across the five runs, the cumulative injected volume V_produced saturates at the same value V_produced,final ≈ 6.85 L in all configurations, and the residual excess volume V_excess is dictated by the pressure equilibrium V_excess,final = (p_close − p_0)/K ≈ 0.32 L, again identical across the five runs. The physically informative quantity is therefore the residual gas volume trapped inside the closed loop at the end of the transient, V_gas,final, which measures the fraction of injected gas that could not be expelled through the module vent before the valve reseated. This quantity shows a clear systematic dependence on both the TR-cell location and the module-vent position. At module-vent position A the residual gas decreases monotonically as the TR cell moves from the inlet-side column towards the outlet-side column: from V_gas,final ≈ 0.91 L for (A, [8]) (Figure 5.16a, worst case, gas removal efficiency η = (V_produced − V_gas,final)/V_produced ≈ 86.7 %) to V_gas,final ≈ 0.79 L for (A, [4]) (Figure 5.16b, η ≈ 88.5 %) to V_gas,final ≈ 0.69 L for (A, [3]) (Figure 5.16c, η ≈ 89.9 %). At module-vent position B the same trend is preserved with a further improvement for the near-vent scenario: V_gas,final ≈ 0.90 L for (B, [8]) (Figure 5.16d, η ≈ 86.9 %) and V_gas,final ≈ 0.62 L for (B, [3]) (Figure 5.16e, η ≈ 90.9 %, best case of the sub-study). Comparing pairs at fixed TR-cell location isolates the effect of the vent shift alone (A → B): for the far-vent scenario ([8]) the improvement is negligible (0.91 → 0.90 L, roughly 1 %), whereas for the near-vent scenario ([3]) the improvement is significant (0.69 → 0.62 L, roughly 10 %). This asymmetric response is consistent with two concurrent mechanisms already visible in the reference run (Section 5.1.3): the buoyancy-driven redistribution of the injected gas, which rises to the ceiling and accumulates preferentially above the TR cell; and the horizontal convection driven by the oil flow from the inlet side to the outlet side of the module, which transports the buoyant plume towards the outlet corner of the ceiling and promotes its ejection through a module vent placed on that side. The two mechanisms act in the same direction and reinforce each other: a TR cell close to the outlet-side column already sits below a vent that is aligned with the natural accumulation region of the gas, and a further shift of the vent towards the outlet (position B) intercepts the plume even more effectively. When the TR cell is far from the module vent (cell [8], inlet-side column), the buoyant plume forms on the opposite side of the ceiling and the 2 cm shift of the vent along the outlet direction is a small displacement compared to the length of the ceiling that separates the TR cell from the vent, so the improvement is negligible. When the TR cell is close to the module vent (cell [3], outlet-side column), the same 2 cm shift represents a larger fraction of the local plume width, and the vent-position improvement is significant. The oil loss follows the mirror trend: the volume of oil expelled through the module vent, V_gas,final − V_excess,final, ranges from 0.30 L in the best case (B, [3]) to 0.59 L in the

worst case (A, [8]) — a factor of two between the best and worst configurations of the sub-study.

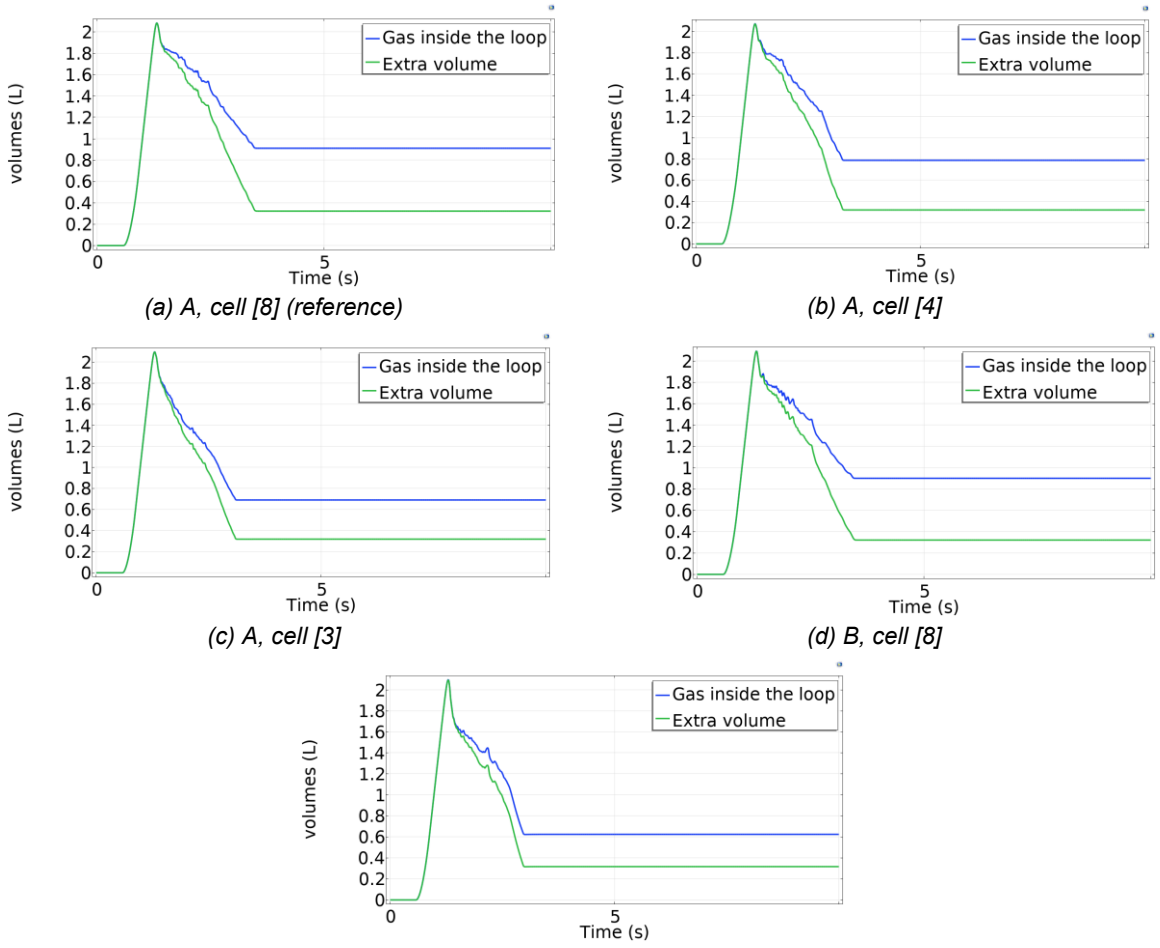

(a) A, cell [8] (reference)

(b) A, cell [4]

(c) A, cell [3]

(d) B, cell [8]

(e) B, cell [3]

Figure 5.16: Time evolution of the residual gas volume inside the closed loop V_gas (blue) and of the net excess volume V_excess (green) for the five runs of the vent-position and TR-cell-location sub-study. In all runs V_produced,final ≈ 6.85 L (not shown, saturates monotonically) and V_excess,final ≈ 0.32 L (determined by (p_close − p_0)/K). (a) Position A, cell [8] (reference): V_gas peaks at 2.08 L and stabilises at 0.91 L (gas removal efficiency η ≈ 86.7 %). (b) Position A, cell [4]: V_gas peaks at 2.05 L, final 0.79 L (η ≈ 88.5 %). (c) Position A, cell [3]: V_gas peaks at 2.10 L, final 0.69 L (η ≈ 89.9 %). (d) Position B, cell [8]: V_gas peaks at 2.10 L, final 0.90 L (η ≈ 86.9 %). (e) Position B, cell [3]: V_gas peaks at 2.10 L, final 0.62 L (η ≈ 90.9 %, best case of the sub-study).

6. Conclusions and future work

6.1 Summary of the work

Immersion cooling with dielectric oils has emerged as a candidate architecture for the next generation of battery thermal management systems, promising uniform temperature control and inherent electrical insulation of the cells. The vast majority of the published research on immersion cooling addresses the steady-state and normal-operation performance of the system; a much smaller number of works have investigated the response of the immersed configuration during a thermal runaway event, where the gas produced by the failing cell must be safely expelled from the module. In particular, the coupling between the gas released by the runaway cell, the buoyancy-driven and oil-flow-driven redistribution of the gas inside the immersion loop, and the dynamic response of a module-level pressure-relief valve on a closed hydraulic loop has, to the author's knowledge, not previously been resolved in a three-dimensional CFD model on the reference 21700 NMC cell format.

To address this gap, the present work has developed a three-dimensional CFD model of a nine-cell 21700 NMC module immersed in the synthetic hydrocarbon coolant AmpCool AC-110, integrated with a lumped-parameter model of the external closed loop and a hysteretic pressure-threshold model of the module-level relief valve. The model has been implemented in COMSOL Multiphysics 6.2 and couples three physics interfaces in a self-consistent way — laminar single-phase flow of the dielectric oil, the two-phase Level Set method for the tracking of the gas-oil interface, and heat transfer in fluids and solids for the cell temperature and the ejected-gas temperature — with six global ODEs, of which the two physically essential ones are the closed-loop compliance driven by the excess volume V_excess and the hysteretic valve state s, driven by the loop pressure with separate opening (p_open = 1.8 bar) and closing (p_close = 1.5 bar) thresholds; the remaining four (V_gas, V_produced, phase_TR, phase_done) act as diagnostic integrals and temporal latches for the heat-source staging. The three-dimensional model has been preceded by a two-dimensional preliminary study (Section 4.1) used to verify the coupling between the fluid physics and the essential ODEs (V_excess and s). The reference cell geometry and physical parameters have been anchored to the experimental characterisation of the 21700 NMC format available in the recent literature (Arcand et al., 2025; Abbott et al., 2022) and to the LIC reduction correlations for hydrocarbon immersion coolants (Wang et al., 2026).

A reference three-dimensional simulation (Section 5.1) has been performed on the nine-cell module with a single cell driven into thermal runaway. The simulation reproduces the full sequence of the venting event: the pre-vent temperature rise of the failing cell, the injection of the vent gases at 1000 K into the surrounding oil, the buoyant migration of the gas plume to the ceiling of the module, the opening of the module vent when the loop pressure crosses p_open, the blowdown of the loop during the open interval, the closing of the vent when the pressure crosses p_close and the residual gas trapped inside the closed loop at the end of the transient. Under the reference conditions, the peak loop pressure reaches 3.28 bar

(approximately 82 % above the atmospheric-side reference), the gas removal efficiency at the end of the transient is $\eta \approx 87$ %, and approximately 0.59 L of oil is expelled together with the gas. Three parametric sub-studies (Section 5.2) have exercised the sensitivity of the reference configuration to the module-vent diameter D_{module}, to the closed-loop stiffness K, and to the position of the module vent on the ceiling combined with the location of the TR cell within the array. The D_{module} sweep reveals a trade-off between peak pressure and oil loss (D = 5 mm produces peak 3.68 bar, D = 15 mm reduces the peak to 3.19 bar at the cost of more oil expelled). The K sweep reveals a more critical trade-off between safety and gas removal efficiency: a stiff loop ($K = 10^9$ Pa/m³) reaches a peak of 11.7 bar, well above the burst threshold of automotive-grade cooling plumbing (5–10 bar). The vent-position sub-study confirms that the buoyancy of the gas plume and the horizontal convection driven by the oil flow concentrate the residual gas near the outlet-side portion of the ceiling, so that a module vent positioned further towards the outlet intercepts the trapped gas more effectively. The three sub-studies interact qualitatively: a larger D_{module} partially compensates a stiffer loop on the peak pressure, but the correct design response to a stiff loop is to reduce the valve response time τ_{vent} to the order of 10 ms rather than to enlarge the vent, in order to keep the peak pressure within the burst-pressure margin of the plumbing and of the pack enclosure.

The main contribution of this work is a three-dimensional CFD model that couples, within a single self-consistent formulation, the immersion cooling of a 21700 NMC module, the closed-loop lumped compliance of the external hydraulic circuit and the hysteretic dynamics of a module-level pressure-relief valve. To the author's knowledge, no previous work in the open literature has resolved this coupling on the reference cell format under a thermal-runaway trigger. The parametric study documented in Chapter 5 identifies a viable baseline configuration for the experimental scale-up of the immersion-cooled architecture: closed-loop stiffness $K = 10^8$ Pa/m³, valve response time $\tau_{vent} = 0.1$ s, module-vent diameter D_{module} = 10 mm and module-vent position offset towards the outlet of the module.

The framework developed in this work is amenable to a number of extensions, including a dedicated experimental campaign for the calibration of the vent-gas parameters on the immersed configuration, a turbulence-enabled extension at the cell vent and at the module vent, a compressible-gas formulation that removes the constant-pressure approximation adopted for the gas density, and an extension to the pack-level architecture that accounts for the distribution of TR-cell locations across multiple modules and for the propagation of the thermal runaway to the cells surrounding the initial failure. These directions are outlined in more detail in Section 6.3.

6.2 Limitations of the model

The model has a number of limitations that the reader should keep in mind when interpreting the results.

First, the TR is represented by an externally imposed heat source and gas injection, with no coupling to the actual reactive kinetics of the cell and therefore it is unable to predict the

onset of TR from a given thermal or mechanical stimulus: it accepts the trigger as given and computes the consequences.

Also, the vent gas is modelled as a single-phase, incompressible fluid with a temperature-dependent density, while the actual vent gas is a multi-species mixture (H_2, CO, CO_2, CH_4, C_2H_4, traces of HF) with significant differences in molecular weight and thermal properties, and is generated at conditions where compressibility effects are not negligible. The single-phase incompressible representation has been adopted because the parametric study targets the global pressure dynamics of the closed loop, where the local gas dynamics in the immediate vicinity of the cell vent has a secondary role.

Moreover, the density of the vent gas is evaluated at a constant reference pressure p_ref = (p_open + p_close)/2 = 1.65 bar rather than at the instantaneous loop pressure p_loop(t) (Section 4.3). The choice is a pragmatic response to solver failures observed when the built-in COMSOL Air material — whose density depends on the local pressure — was used in place of the custom gas material. The bias on ρ_gas is bounded at ±10 % during the normal open–closed cycle of the safety valve but grows in transient regimes where p_loop departs significantly from p_ref: in particular, during the initial rise of the loop pressure from p_0 = 1.2 bar to p_open = 1.8 bar, and during the transient overshoot above p_open in the stiff-loop configuration (K = 10^9 Pa/m³ of the sensitivity sweep of Section 5.2.3), where the assumption of a constant p_ref introduces a systematic mismatch between the modelled gas density and the one that would be obtained with a fully coupled compressible-gas formulation. The bias is inherited by all quantities that depend on ρ_gas, in particular V_produced (integrated by the ODE from the mass flow rate imposed at the cell vent). A fully compressible-gas formulation is listed as future work in Section 6.3.

At the baseline mean injection velocity u_gas_peak = 50 m/s the Reynolds number at the cell vent computed from the mean velocity is approximately 8 × 10^3, above the transitional threshold for pipe flow. On this regard two additional considerations must be noted. First, because the cell-vent boundary condition is imposed as a Mass Flow Rate (Section 4.7.1), the local velocity profile is peaked at the centre of the face and the peak local velocity, observed from the reference simulation, reaches approximately 85 m/s (a factor of order 1.7 above the mean), corresponding to a local Reynolds number Re_cell,local ≈ 4 × 10^4. Second, once the safety valve opens, the module-vent boundary itself experiences peak local velocities of the order of 160 m/s (Section 5.1), driven by the pressure difference between the loop and the atmosphere, at a corresponding local Reynolds number Re_module ≈ 8 × 10^4. Both values lie well within the turbulent regime. The Navier–Stokes equations have nevertheless been solved with the laminar flow interface throughout, without a turbulence closure. This choice is motivated by three considerations: first, the two regions in which Re exceeds the laminar threshold are geometrically confined — the centre of the cell-vent injection face and the module-vent efflux region during valve opening — while the bulk of the module (oil flow around the cells, buoyancy-driven propagation of the gas plume towards the ceiling, oil recirculation between inlet and outlet) remains well within the laminar regime, with maximum Reynolds

numbers below 10^3; second, the introduction of a turbulence model in the transient two-phase Level Set setup would add significant numerical complexity (two additional transport equations for k and ε or k and ω, wall functions on the cylindrical cell walls, and a turbulence–interface coupling on the diffuse Level Set that is known to require a substantially smaller time step for stable reinitialisation) and would risk destabilising a model that only recently reached a converged reference run; third, the resulting increase in wall time per run, estimated at approximately one order of magnitude, would make the parametric campaign of Chapter 5 unfeasible within the time frame of the present work. The laminar assumption is therefore accepted as a deliberate modelling choice, and the parametric study targets integrated quantities (loop pressure, valve dynamics, volume balance) that are relatively insensitive to the detailed turbulent structure at the vents. A turbulent extension is listed as future work in Section 6.3.

In addition, the peak gas injection velocity adopted at the cell vent, u_gas_peak = 50 m/s (Section 4.5.1), corresponds to a modelled venting duration of t_vent ≈ 1.65 s. This duration is deliberately stretched with respect to the acute gas ejection phase of 291–376 ms measured by Arcand et al. (2025) for the same cell format in air, in order to keep the local Reynolds number at the cell vent within the range of validity of the laminar-flow approximation adopted in this work (Section 3.1). The chosen value is intermediate between the Arcand acute phase and the full jet-flame durations of 4–18 s reported by Shelke et al. (2022) for nail penetration of 21700 NMC cells in air, and is consistent with the qualitative expectation that immersion cooling suppresses the jet-flame phase while leaving the acute gas ejection largely unchanged. However, no direct experimental measurement of the venting duration under dielectric-oil immersion for a 21700 NMC cell is currently available in the open literature; the associated uncertainty is dominated by this measurement gap. The total gas volume balance is preserved by the choice V_max derived in Section 4.5.1, and the peak loop pressure is representative because it is governed by the integral of gas produced against the closed-loop stiffness K, but the fine temporal structure of the transient (peak instantaneous pressure gradients, valve reaction dynamics under a sharper trigger) is not resolved. A turbulent, compressible extension of the model at the cell vent, listed in Section 6.3, would remove this limitation.

Another limitation comes from the volumetric heat-release rate which is switched off above T_peak = 600 K through the gate flc2hs((T_peak − T_cell)/1[K], T_smooth) of Section 3.5. This empirical cap-off is not derived from chemical-kinetic modelling but represents the apparent thermal saturation observed in the immersion-cooling experiments of Wang et al. (2026) on NMC and similar cells. The unconstrained adiabatic temperature of the cell would be of the order of 800 K (Abbott et al., 2022 for dry inert atmosphere), but the cap-off introduced here reflects the partial cooling provided by the oil and the implicit chemical saturation. As a consequence of this formulation, the venting event terminates not through the exhaustion of the gas inventory of the cell (phase_done = 1, which would require V_produced to reach V_max), but through the thermal feedback loop in which the heat source is switched

off above T_peak, the cell starts to cool by conduction to the oil, and the cell-vent injection ceases when T_cell drops back below T_open_cell. The total amount of gas vented in the present formulation can therefore be smaller than V_max, depending on the balance between heat injection and cooling rate. Modelling the chemistry explicitly would replace this empirical cap with a multi-step Arrhenius kinetic scheme and a chemistry-driven gas source; this is left as future work.

The external cooling loop is represented by a single stiffness parameter K, without resolving the geometry of the pump, the expansion chamber and the heat exchanger. This is a deliberate generalisation, because the actual loop architecture is installation-specific; the K parameter aggregates all the relevant compliances into a single variable that is then explored parametrically.

Lastly, the safety valve is represented by a state variable $s \in [0, 1]$ governed by a relaxation ODE, with no explicit modelling of the mechanical inertia of the valve element, of its seal dynamics, or of the burst transient of a rupture disc. The model is appropriate for a generic pressure-threshold valve and is parametrised by p_open, p_close and τ_vent; mapping these to specific valve technologies is left to the engineering design phase.

6.3 Future work

The present work might be extended through several new additions.

A side-by-side comparison with the Phase Field formulation of the same problem would clarify the relative merits and limitations of the two diffuse-interface methods on a 3D closed-loop configuration. The Phase Field method has advantages in mass conservation and in interface sharpness control, at the cost of an additional mobility parameter; the present work has stopped short of completing this comparison.

Also, the parametric study keeps u_gas_peak fixed at the baseline value 50 m/s and explores the three geometric/hydraulic variables D_module, vent position and K. A dedicated sensitivity sweep on u_gas_peak would provide a more complete picture of the response of the venting system to different TR severities. The upper end of the sweep would require the introduction of a turbulence model at the cell vent (see item below).

In addition, a coupling between the present pack-level model and a cell-level chemical-kinetic model of the TR (e.g. a multi-step Arrhenius reaction model of the cathode and anode decomposition) would allow the TR trigger to be predicted from the thermal/mechanical conditions, rather than imposed externally.

The study has been performed under the laminar-flow approximation for the reasons discussed in Section 6.2. A natural extension is the introduction of a Reynolds-averaged turbulence closure (typically k–ε or k–ω SST) in the two localised regions where the local Reynolds number exceeds the laminar threshold: the centre of the cell-vent injection face ($Re_{local} \approx 4 \times 10^4$ from the reference simulation) and the module-vent efflux region once the valve opens ($Re_{module} \approx 8 \times 10^4$). A turbulent-in-vents, laminar-in-bulk hybrid strategy could

be implemented via a Reynolds-based blending function on the turbulent viscosity. Such an extension would validate the laminar assumption on the integrated quantities of interest and enable a more accurate resolution of the local gas plume structure and of the atomisation at the vents. The main obstacle is the additional wall-time per run — of the order of one order of magnitude on the same hardware — which requires either a larger computational allocation or a more targeted parametric campaign than the one exercised here.

A physically more consistent extension of the present model would replace the constant reference pressure p_ref adopted for the vent gas (Section 4.3) with the instantaneous loop pressure p_loop(t), so that ρ_gas responds dynamically to the loop pressurisation. This was attempted using the built-in COMSOL Air material, whose density depends on the local pressure field, but the resulting setup produced solver failures at t = 0, traced to the absence of a well-defined initial pressure field on the gas sub-domain before the venting event begins. A viable path forward is to switch the flow physics from the current incompressible laminar interface to a compressible-flow (weakly compressible or full compressible Navier–Stokes) interface, in which the initial pressure field is part of the state and the solver initialises consistently; the gas density can then be computed from p_loop directly without the need for the p_ref workaround, and the ±10 % bias documented in Section 6.2 is removed. The change of physics interface has non-trivial implications on convergence and time-step size, and was therefore deferred to future work.

The study has been performed on a single nine-cell module. A natural extension is the coupling of the present module-level model to a pack-level architecture in which several modules share a common enclosure and a common coolant circuit. Two aspects require dedicated attention at this scale. First, the distance between the TR cell and the nearest module-level vent, which controls the buoyancy-driven and oil-flow-driven transport of the gas plume across the enclosure, sets a natural bound on the size of the region drained by a single module vent; beyond this bound, a distributed vent architecture with one vent per module is likely to outperform a single centralised vent per pack. Second, the vapour-space head at the pack scale, of the order of a few tens of litres for a compact automotive enclosure, relaxes the tight safety-volume ratio of the single-module configuration and modifies the sizing criteria for D_module and K established at the module scale.

The gas ejected by the failing cell reaches the neighbouring cells at temperatures of the order of 1000 K and may trigger a cascade of secondary runaways if the immersion-cooling architecture cannot arrest the propagation on time scales comparable to the venting duration itself. The present model treats the TR of a single cell as an isolated event and does not resolve the thermal loading of the neighbouring cells. A natural extension is the coupling of the vent-gas transport to a heat-transfer criterion on the neighbouring cell surfaces, with a temperature threshold that triggers the venting of the next cell when exceeded for a duration compatible with the empirical propagation kinetics. This extension would allow the parametric study to be repeated on multi-cell-TR scenarios and to quantify the propagation-arrest capability of the immersion-cooling architecture as a function of the design variables.

Lastly, the natural extension of the work is an experimental campaign on a single-cell test rig representative of the immersion-cooled module configuration, with measurements of the loop pressure transient, of the cell temperature and of the gas inventory at the vent.

References

Peer-reviewed literature

Abbott, K.C., Buston, J.E.H., Gill, J., Goddard, S.L., Howard, D., Howard, G., Read, E., Williams, R.C.E. (2022). Comprehensive gas analysis of a 21700 Li($Ni_{0.8}Co_{0.1}Mn_{0.1}$)O_2 cell using mass spectrometry. Journal of Power Sources, 539, 231585. https://doi.org/10.1016/j.jpowsour.2022.231585

Arcand, X., Mooney, W., Francoeur, D., Bouchard, M., Rancourt, D., Azarkish, H. (2025). Evaluation of Thermal Runaway Venting Characteristics for High-Energy Density and High-Power Density 21700-format Li-ion NMC Cells. In AIAA AVIATION Forum, AIAA 2025-3064. https://doi.org/10.2514/6.2025-3064

Bernardi, D., Pawlikowski, E., Newman, J. (1985). A general energy balance for battery systems. Journal of the Electrochemical Society, 132(1), 5–12. https://doi.org/10.1149/1.2113792

Coelho de Sousa Marques, M.A., Guimarey, M.J.G., Domínguez-Arca, V., Amigo, A., Fernández, J. (2021). Heat capacity, density, surface tension, and contact angle for polyalphaolefins and ester lubricants. Thermochimica Acta, 703, 178994. https://doi.org/10.1016/j.tca.2021.178994

Falcone, M., Palka Bayard De Volo, E., Hellany, A., Rossi, C., Pulvirenti, B. (2021). Lithium-Ion Battery Thermal Management Systems: A Survey and New CFD Results. Batteries, 7(4), 86. https://doi.org/10.3390/batteries7040086

Feng, X., Ouyang, M., Liu, X., Lu, L., Xia, Y., He, X. (2018). Thermal runaway mechanism of lithium ion battery for electric vehicles: A review. Energy Storage Materials, 10, 246–267. https://doi.org/10.1016/j.ensm.2017.05.013

Feng, X., Ren, D., He, X., Ouyang, M. (2020). Mitigating thermal runaway of lithium-ion batteries. Joule, 4(4), 743–770. https://doi.org/10.1016/j.joule.2020.02.010

Kambic, M., Kalb, R., Tasner, T., Lovrec, D. (2014). High Bulk Modulus of Ionic Liquid and Effects on Performance of Hydraulic System. The Scientific World Journal, 2014, 504762. https://doi.org/10.1155/2014/504762

Kim, J., Mallarapu, A., Finegan, D. P., Santhanagopalan, S. (2021). Modeling cell venting and gas-phase reactions in 18650 lithium ion batteries during thermal runaway. Journal of Power Sources, 489, 229496. https://doi.org/10.1016/j.jpowsour.2021.229496

Kong, D., Wang, G., Ping, P., Wen, J. (2022). A coupled conjugate heat transfer and CFD model for the thermal runaway evolution and jet fire of 18650 lithium-ion battery under thermal abuse. eTransportation, 12, 100157. https://doi.org/10.1016/j.etran.2022.100157

Li, W., León Quiroga, V., Crompton, K. R., Ostanek, J. K. (2021). High Resolution 3-D Simulations of Venting in 18650 Lithium-Ion Cells. Frontiers in Energy Research, 9, 788239. https://doi.org/10.3389/fenrg.2021.788239

Liu, X., Li, Y., Gao, S., Lv, X., Zhou, Z., Wu, W.-T., Wei, L., Lyu, J., Li, Y., Song, Y. (2025). Suppression thermal runaway propagation of $LiFePO_4$ batteries during nail penetration test based on liquid immersion cooling. Energy, 330, 136973. https://doi.org/10.1016/j.energy.2025.136973

Chen, H., Buston, J.E.H., Gill, J., Howard, D., Williams, R.C.E., Rao Vendra, C.M., Shelke, A., Wen, J.X. (2020). An experimental study on thermal runaway characteristics of lithium-ion batteries with high specific energy and prediction of heat release rate. Journal of Power Sources, 472, 228585. https://doi.org/10.1016/j.jpowsour.2020.228585 Merritt, H. E. (1967). Hydraulic Control Systems. John Wiley & Sons, New York.

Olsson, E., Kreiss, G. (2005). A conservative level set method for two phase flow. Journal of Computational Physics, 210(1), 225–246. https://doi.org/10.1016/j.jcp.2005.04.007

Ostanek, J. K., Li, W., Mukherjee, P. P., Crompton, K. R., Hacker, C. (2020). Simulating onset and evolution of thermal runaway in Li-ion cells using a coupled thermal and venting model. Applied Energy, 268, 114972. https://doi.org/10.1016/j.apenergy.2020.114972

Rabie, M. G. (2009). Fluid Power Engineering. McGraw-Hill, New York.

Roe, C., Feng, X., White, G., Li, R., Wang, H., Rui, X., Li, C., Zhang, F., Null, V., Parkes, M., Patel, Y., Wang, Y., Wang, H., Ouyang, M., Offer, G., Wu, B. (2022). Immersion cooling for lithium-ion batteries – A review. Journal of Power Sources, 525, 231094. https://doi.org/10.1016/j.jpowsour.2022.231094

Shelke, A.V., Buston, J.E.H., Gill, J., Howard, D., Abbott, K.C., Goddard, S.L., Read, E., Howard, G.E., Abaza, A., Cooper, B., Wen, J.X. (2022). Characterizing and predicting 21700 NMC lithium-ion battery thermal runaway induced by nail penetration. Applied Thermal Engineering, 209, 118278. https://doi.org/10.1016/j.applthermaleng.2022.118278

Wahab, A., Najmi, A.-U.-H., Senobar, H., Amjady, N., Kemper, H., Khayyam, H. (2025). Immersion cooling innovations and critical hurdles in Li-ion battery cooling for future electric vehicles. Renewable and Sustainable Energy Reviews, 211, 115268. https://doi.org/10.1016/j.rser.2024.115268

Wang, G., Kong, D., Ping, P., He, X., Lv, H., Zhao, H., Hong, W. (2023). Modeling venting behavior of lithium-ion batteries during thermal runaway propagation by coupling CFD and thermal resistance network. Applied Energy, 334, 120660. https://doi.org/10.1016/j.apenergy.2023.120660

Wang, Y., Hao, T., Qi, Y., Fu, Y. (2026). Experimental investigation of immersion cooling in removing heat and gas from 26700 lithium-ion batteries during thermal runaway. Journal of Energy Storage, 153, 121044. https://doi.org/10.1016/j.est.2026.121044

Ye, Y., Mao, Y., Zhao, L., Chen, Y., Chen, M. (2024). Experimental investigation of thermal runaway behavior and propagation inhibition of lithium-ion battery by immersion cooling. Applied Thermal Engineering, 256, 124093. https://doi.org/10.1016/j.applthermaleng.2024.124093

Zhao, L., Song, G., Yang, X., Pan, R. (2025). Experiments on the effects of liquid immersion cooling on the thermal runaway (TR) behaviors of 280 Ah lithium-ion batteries subjected to different TR-triggering conditions. Journal of Energy Storage, 121, 116358. https://doi.org/10.1016/j.est.2025.116358

Standards and technical documentation

API Standard 520 (2020). Sizing, Selection, and Installation of Pressure-relieving Devices in Refineries — Part I: Sizing and Selection, 10th ed. American Petroleum Institute, Washington, DC.

API Standard 521 (2014). Pressure-relieving and Depressuring Systems, 6th ed. American Petroleum Institute, Washington, DC.

ASME Boiler and Pressure Vessel Code Section VIII (2023). Rules for Construction of Pressure Vessels, Division 1. American Society of Mechanical Engineers, New York.

Donaldson Company (2023). Dual-Stage Vent and Tetratex® ePTFE Membrane Technology — Technical Documentation. Retrieved from https://www.donaldson.com/en/products/venting/battery-pack/

Engineered Fluids (2026). AmpCool AC-100 Series Technical Data Sheet (DOC.ID AC1XX-ENG-TDS-20260505). Engineered Fluids, Portland, OR, USA.

Gore W. L. & Associates (2023). PolyVent ePTFE protective vents — Technical documentation. Retrieved from https://www.gore.com/products/vents-automotive-battery-components

Lotus Cars (2024). Emeya – Technical Specifications. Retrieved from https://www.lotuscars.com/en/emeya/ (accessed July 2026).

Porsche AG (2024). Range and charging – The new Taycan. Porsche Newsroom. Retrieved from https://newsroom.porsche.com/en/press-kits/taycan/ (accessed July 2026).

Rivera-Salinas, J. E., Cruz-Ramírez, A. (2015). Simulation of gas injection with the level set method. Proceedings of the 2015 COMSOL Conference, Grenoble. https://www.comsol.com/paper/simulation-of-gas-injection-with-the-level-set-method-30731

Appendix A — Methodology notes

A.1 Discarded modelling alternatives

A.1.1 Phase Field vs Level Set

The Phase Field method was considered as an alternative to Level Set for the two-phase interface. Phase Field has two advantages: better global mass conservation and an energy-consistent formulation of the interface. The disadvantages are higher numerical complexity (the mobility parameter χ is critical and not directly tied to a physical observable) and higher computational cost. Given the time constraint of this work and the additional complexity introduced by the vent valve, Level Set has been preferred. A Phase Field comparison could be the subject of a follow-up study.

A.1.2 Lumped vs distributed external-loop model

A lumped-parameter representation of the internal loop, with a single stiffness K, has been preferred over alternatives that would explicitly resolve an expansion chamber, a 1D loop model or a hydraulic circuit network. The reason is generality: a single K parameter allows a sensitivity study without committing to a specific installation, since the actual loop architecture varies from one vehicle to another.

A.1.3 Simple-threshold vs hysteretic valve

A non-hysteretic threshold valve, in which the valve opens and closes at the same instantaneous pressure, was implemented in the earliest versions of the model as the simplest possible representation of the module vent. It was subsequently discarded in favour of the hysteretic formulation (Sections 3.4.3 and 4.6), because the non-hysteretic version led to valve chattering during the venting transient and did not keep the vent open long enough to remove the residual gas trapped under the module ceiling. The hysteretic model with a blowdown $p_{open} - p_{close} = 0.3$ bar solves both issues. A latching (rupture-disc) variant of the hysteretic model, in which the valve remains open once triggered, was also considered and discarded, because it is incompatible with the design intent of keeping the coolant loop pressurised after the venting event.

A.1.4 ePTFE porous-membrane representation of the module vent

An alternative representation of the module vent was considered, in which the vent is modelled as an explicit porous-medium sub-domain representing an ePTFE dual-stage membrane. The membrane would be resolved as a thin sub-domain with permeability dependent on the local stress (to model the burst event), described by a Darcy or Brinkman formulation coupled to the surrounding laminar flow. This representation would in principle allow the transient response of the membrane, the burst pressure distribution and the residual flow through the intact ePTFE film to be captured explicitly.

As argued in Chapter 2, however, ePTFE dual-stage membranes are physically incompatible with the operating regime of the present system: they are designed for open

packs with air on both sides and rely on passive pressure equalisation, which does not occur in a closed and pre-pressurised liquid loop. In the reference architecture of this thesis the module vent must be a mechanical pressure-relief valve with a defined opening threshold and blowdown, for which the boundary-condition representation of Section 3.4.1 is appropriate. The additional physical detail of the porous-medium approach is therefore not informative for the design questions addressed in the parametric study, while it would increase the numerical cost of every run substantially (finer mesh in the membrane sub-domain, additional Darcy or Brinkman physics, tighter time step during the burst event). The boundary-condition representation has therefore been retained as the reference formulation.

A.1.5 Velocity-BC formulation of the module vent

In a first implementation of the boundary-condition approach, the module vent was represented as a prescribed normal velocity at the boundary, modulated by a smooth opening function of the loop pressure. In compact form, $u_{n,vent} = u_{vent,max} \cdot vent_{open}(\Delta p)$, where $\Delta p = p_0 + K \cdot V_{gas} - p_{open}$ is the difference between the loop pressure and the opening threshold, $u_{vent,max}$ is a numerical cap on the efflux velocity, and $vent_{open}$ is a smooth Heaviside step centred on zero (built via the flc2hs COMSOL primitive with smoothing width 10 000 Pa). The smoothing is required for numerical stability of the transient solver.

This formulation is robust but introduces an artificial parameter $u_{vent,max}$ that does not have a clean physical interpretation: it is in practice a numerical clamp on the efflux velocity, not a physical quantity. Its value in early tests of the model was set to 5 m/s. The formulation was not used in any run of the present thesis: both the 2D preliminary verification (Section 4.1) and the 3D reference simulation (Chapter 5) adopt the pressure-BC formulation of Section 3.4.2, which does not require an arbitrary velocity cap. The velocity-BC formulation is documented here for completeness and as a possible starting point for future work on inertia-dominated venting regimes, where prescribing the efflux velocity may be preferable to prescribing the boundary pressure.

A.2 Numerical difficulties encountered

During the development of the model, several numerical issues have been encountered and resolved. They are documented here for completeness.

- The option 'Compensate for hydrostatic pressure' in the laminar-flow interface must be kept active, despite the conservative warning that COMSOL displays for its coexistence with two-phase flow. Deactivating it produces unphysical velocity profiles even in single-phase oil flow.
- The time step of the transient solver must be set manually to ensure convergence over the entire time horizon of the simulation. In the 2D preliminary model (Section 4.1) the values initial step = 0.001 s, maximum step = 0.05 s were sufficient; in the 3D reference model (Section 4.9) a tighter maximum step of 1×10^{-2} s and initial step of 5×10^{-4} s were required to control the sharper transients of the fully-resolved venting event. Automatic settings produce convergence failures at the first step.

- The opening / closing transition of the valve must be smoothed over a finite pressure interval (p_smooth ≈ 0.05 bar in the present model) to avoid numerical oscillations. A discontinuous step at the exact threshold compromises convergence.
- Phase Field convergence. The Phase Field formulation was attempted as an alternative to Level Set, but on the 3D closed-domain configuration with the present density contrast and thermal source it has been unable to accept the first time step in less than approximately 40 minutes of wall time, with both the default mobility ($\chi = 1$) and the value recommended by the COMSOL Knowledge Base for the local mesh size, surface tension and peak velocity ($\chi = 500$). The Level Set formulation, with the same mesh, the same boundary conditions and the same solver settings, accepts the first step in a few seconds and proceeds at approximately 7 ms of simulated time per minute of wall time once the start-up transient has completed.
- Spurious gas patches in the Level Set. An early version of the 3D Level Set model showed spurious patches of high ls2.Vf2 in regions far from the physical interface, appearing within the first 100 ms of the simulation. The cause was an insufficient reinitialisation velocity γ relative to the peak velocity in the domain, which in the current reference model reaches approximately 160 m/s at the module vent during the efflux phase. The final value adopted is $\gamma = 5 \cdot u\_gas\_peak = 250$ m/s, above this peak, and has eliminated the spurious patches. Note: the initial $\gamma = 25$ m/s was calibrated on a preliminary u_gas_peak = 5 m/s, subsequently updated to u_gas_peak = 50 m/s on the basis of Arcand et al. (2025).
- Interface thickness ε_ls. The value of ε_ls must be consistent with the local mesh size in the regions of physical interest. Initial values of $\varepsilon\_ls = 7 \times 10^{-4}$ m and the meshes used in early runs led to interface thicknesses below one mesh element across, with consequent loss of mass conservation. The final value $\varepsilon\_ls = 3 \times 10^{-3}$ m, combined with the refined mesh of Section 4.8, gives approximately six elements across the diffuse interface in the regions of interest — larger than the h_c/2 rule of thumb recommended by the COMSOL Knowledge Base, but justified by the strong material-property gradients on the interface (density contrast $\Delta\rho/\rho \approx 3$ orders of magnitude, thermal conductivity ratio k_oil/k_gas ≈ 6) which benefit from additional diffuse-interface regularisation for solver stability.
- PARDISO out-of-core. With an early mesh of approximately 400 000 elements, the PARDISO factorisation exceeded the 32 GB of available RAM and the solver passed to out-of-core mode, with a rate degradation of approximately a factor of 10. The mesh has been re-engineered with size constraints local to the most demanding regions (cell vent, module vent, oil–gas interface) and coarser elements in the bulk of the oil, bringing the total count down to approximately 120 000 elements and the factorisation memory to approximately 15 GB, well within the in-core regime.
- Mesh quality and sliver elements. The geometry sequence, obtained by Boolean operations between the cylindrical cells and the planar walls of the case, initially

produced a large number of short intersection lines at the interfaces between the curved cell surfaces and the flat case walls. These short lines forced the mesh generator to create disproportionately small boundary faces, which in turn seeded highly distorted tetrahedral elements (sliver elements) with minimum quality well below 0.05. The solver crashed at the first time step on these degenerate elements, and reaching a converged configuration required refining the mesh globally to approximately 400 000 elements — well beyond the in-core PARDISO limit and with an unacceptable wall time per run. The geometry has been progressively cleaned through Virtual Operations (Ignore Edge and Ignore Vertex nodes applied to the shortest intersection lines) to remove the seeds of the sliver elements. The final mesh of the reference run consists of approximately 120 000 elements with a minimum element quality of approximately 0.155 on the residual slivers, against an average of 0.65; the residual slivers are concentrated in regions where the physics is well-behaved (no gradients) and do not compromise the convergence of the solver.

- Solver settings for the first step. The first time step of the time-dependent solver, preceded by a Backward Euler consistent initialisation, has been the subject of an extended calibration. The default values of the relevant parameters (initial damping factor of Newton = 1; maximum Newton iterations per step = 4; fraction of the initial step for Backward Euler = 0.01; safety factor for Backward Euler = 20) led to a first-step time of approximately 40 minutes of wall time in early versions of the 3D model. The values adopted in the final reference run (maximum Newton iterations = 15; fraction of the initial step for BE = 0.5; safety factor for BE = 30) reduce the first-step wall time to a few seconds and lead to a stable BDF evolution thereafter.

Appendix B — Nomenclature

p	pressure [Pa]
p_0	nominal operating pressure of the closed loop [Pa]
p_atm	atmospheric pressure outside the loop [Pa]
p_open / p_close	opening and closing thresholds for the hysteretic valve [Pa]
K	hydraulic stiffness of the external closed loop [Pa/m³]
V_gas	volume of gas remaining in the closed loop [m³]
u_oil	oil inlet velocity [m/s]
u_gas_peak	peak gas injection velocity at the cell vent [m/s]
q_peak	peak volumetric heat-release rate of the TR cell [W/m³]
T_cell	volume-averaged temperature of the TR cell domain [K]
T_open_cell	cell-vent opening temperature threshold [K]
T_peak	empirical peak cell temperature above which the heat source is switched off [K]
T_inlet_gas	temperature of the gas injected at the cell vent during venting [K]
p_ref	reference pressure used in the ideal-gas density of the vent gas, = (p_open + p_close)/2 [Pa]
M_gas	molar mass of the vent gas mixture (representative H_2/CO/CO_2) [kg/mol]
ε	regularisation width of the pressure-based smoothed Heaviside gate in the valve state ODE [Pa]
ρ	density [kg/m³]
μ	dynamic viscosity [Pa·s]
k	thermal conductivity [W/(m·K)]
c_p	specific heat at constant pressure [J/(kg·K)]
σ	surface tension coefficient [N/m]
φ (ls2.Vf1)	Level Set scalar / volume fraction of fluid 1 (oil) [—]
ε_ls	Level Set diffuse interface thickness [m]
γ	Level Set reinitialisation velocity [m/s]
s	valve state variable (0 closed, 1 open) [—]
τ_vent	response time of the valve state ODE [s]
V_excess	excess fluid volume in the closed loop [m³]
V_produced	cumulative volume of gas produced by the cell vent [m³]
V_max	empirical limit of gas inventory per cell [m³]
V_scale_phase	smoothing width of the phase_done latch transition around V_max (8×10^{-5} m³) [m³]

R_vent	closed-state hydraulic resistance of the module-vent boundary [Pa·s/m]
D_module	diameter of the module-vent face [m]
A_inj	area of the cell-vent injection face [m²]
ls2.Vf2	Level Set volume fraction of fluid 2 (gas), = 1 − ls2.Vf1 [—]
phase_TR	latch variable indicating active-TR stage (0 → 1) [—]
phase_done	latch variable indicating exhausted-cell stage (0 → 1) [—]
τ_phase	response time of the phase latches [s]
Φ_in(t)	cell-vent gas injection flow rate [m³/s]
Q_vent(t)	module-vent fluid efflux flow rate [m³/s]
T_smooth	dimensionless smoothing parameter of the temperature gate in cell-vent injection and phase_TR latch (effective transition = 2·T_smooth·1[K] around T_open_cell) [—]