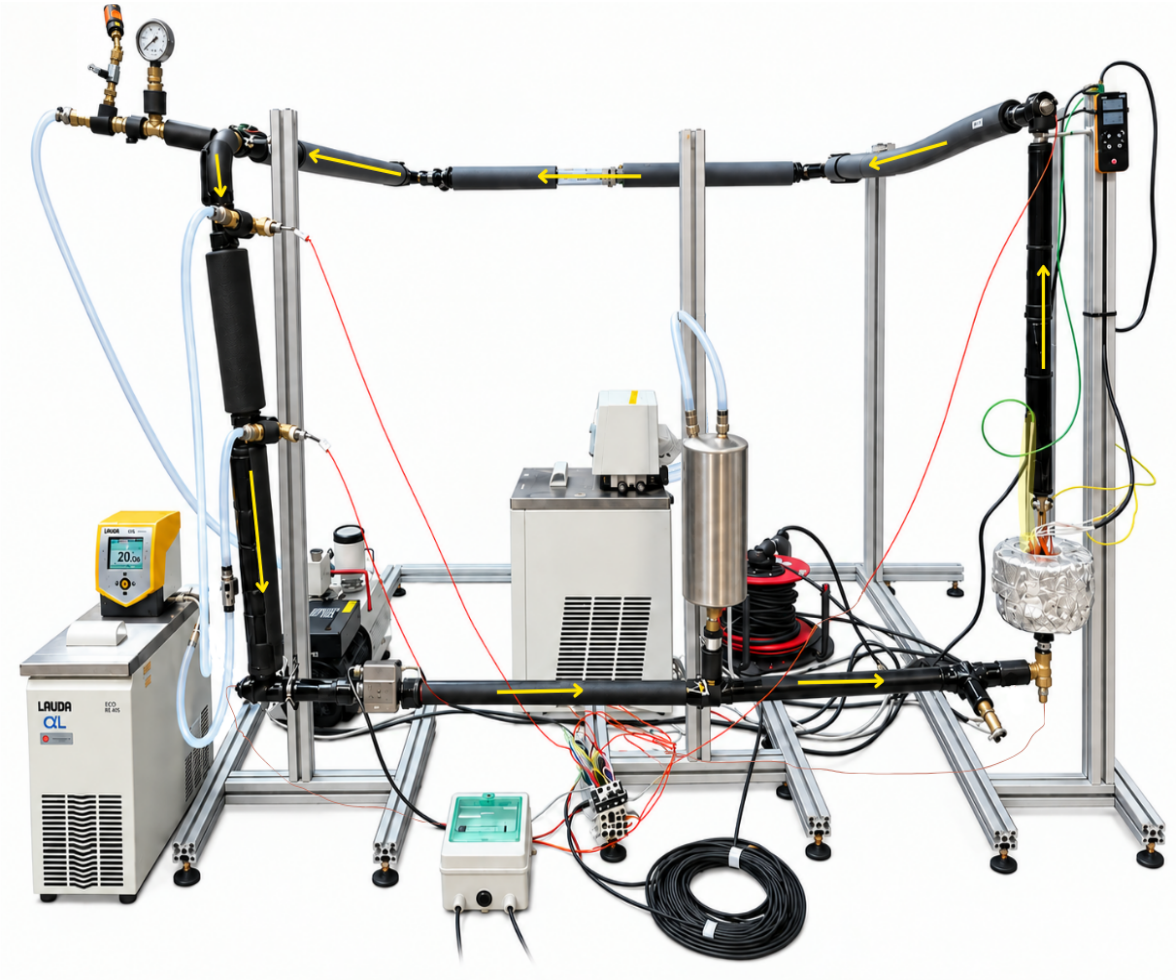
Thesis Report

Development of an MRI-Compatible Sub-Atmospheric Two-Phase Thermosyphon Loop

TU Delft



Author:

Themistoklis Kyriakos (5523443)

Supervisors:

Dr. M. Raquel Serial

Dr. Ir. Emanuele Zanetti

Ir. M.E. Miguel Cabral Dionísio

November 17, 2026

Abstract

Two-phase flow carries most of the heat in evaporators and condensers, yet its internal structure is hard to measure without disturbing it. Magnetic resonance imaging resolves liquid distribution and velocity inside an opaque pipe without contact, but every published two-phase MRI study uses air and water near atmospheric pressure. Water and steam below atmospheric pressure, where the liquid is twelve thousand times denser than its vapor, have never been imaged.

This thesis covers the design, construction and operation of a thermosyphon loop built for that purpose. The loop is a closed circuit in which water boils, rises, condenses and falls back under gravity, with no pump. A hot test section degrades the image, so the loop runs at $0.12\,\mathrm{bar}$, where water boils at $50\,^{\circ}\mathrm{C}$, with the heater $1.7\,\mathrm{m}$ away and nothing ferromagnetic near the magnet. A steady-state pressure balance model set the pipe diameter and the sizing of the heater, the condenser and the expansion vessel that fixes the pressure of the loop.

The loop was built and operated on eleven days, at charges between 48 and 70 percent of its volume. The circulation had two states. A freshly charged loop circulates quickly and unsteadily at $170\,\mathrm{mL\,min^{-1}}$ to $230\,\mathrm{mL\,min^{-1}}$, and after roughly four hours it changes, with no way back, to a slow and steady $12\,\mathrm{mL\,min^{-1}}$. Both states repeat across days, but the change between them is governed by the time since charging rather than by any setting, and its cause is not resolved. The model describes only the slow state.

The result is a working thermosyphon whose measurement section can go inside a scanner. The fast state is the one with liquid structures worth imaging, but it drifts throughout and the model does not describe it, so the condition being imaged would not be known. In the slow state, which is steady and predictable, no liquid film is visible in the section at all. The changes needed before a scan can be planned are identified.

Acknowledgements

This thesis was carried out at the Process and Energy department of TU Delft. Without the people around it, it would have been a weaker piece of work, and the year behind it a lonelier one.

I would first like to thank my supervisors, Dr. M. Raquel Serial, Dr. Ir. Emanuele Zanetti and Ir. M. E. Miguel Cabral Dionísio, for their guidance throughout this project. Our weekly meetings were never just progress reports. They usually turned into brainstorming sessions, and I almost always walked out with more ideas than I walked in with. The help also went well beyond the meeting room. All three came down to the frame, helped assemble the setup and were there while the experiments were running, and I learned a great deal from working next to them.

The setup itself would not exist without the DEMO team at the Process and Energy department. Martijn Karsten and Bart van der Lubbe helped me order the components, guided me through the assembly, and carried out the welding and cutting I was not permitted to do myself. Most of what I now know about building an experiment, as opposed to designing one on paper, I learned from them.

A large part of this year was spent at the frame alongside Daniel Casaril, a fellow EFPT student running his own experiment next to mine. We shared the space, the equipment and a long run of twelve hour days, and somewhere in among them we became friends. The work would have been finished either way, but it would have been a far duller year without him there. Thank you for the help, the second opinions, and for making the lab somewhere I actually wanted to be.

Finally, to everyone else in the lab who lent a hand, answered a question in passing, or helped me hunt down a fitting at the wrong end of the building, thank you.

Themis Kyriakos
Delft, October 2026

Contents

Abstract 1

Acknowledgements i

1 Introduction 2

2 Fundamentals of Two-Phase Flow 3
- 2.1 Describing a Two-Phase Flow . . . 3
- 2.2 Two-Phase Flow Regimes in Horizontal Pipes . . . 3
- 2.3 Flow Pattern Maps . . . 4
  - 2.3.1 Superficial velocity maps . . . 5
  - 2.3.2 Superficial Reynolds number maps . . . 5
  - 2.3.3 Mechanistic maps . . . 6
  - 2.3.4 Mass flux and quality maps . . . 6
- 2.4 Flow Visualization Techniques in Two-Phase Flow . . . 8
  - 2.4.1 Optical Imaging . . . 8
  - 2.4.2 X-ray Visualization . . . 8
  - 2.4.3 Other Measurement Techniques . . . 8

3 Magnetic Resonance Imaging 9
- 3.1 Working Principle of MRI . . . 9
- 3.2 MRI in Multiphase Flow Studies . . . 10
- 3.3 Experimental Constraints of MRI-Compatible Systems . . . 11

4 Fundamentals of Thermosyphon Loops 12
- 4.1 Operating Principle of a Thermosyphon . . . 13
- 4.2 Influence of Filling Ratio . . . 16
- 4.3 Heat Transfer Limits . . . 16
  - 4.3.1 Dryout Limit . . . 16
  - 4.3.2 Flooding Limit . . . 17
  - 4.3.3 Operating Window . . . 17

5 Research Questions 17

6 Design and modeling of the thermosyphon loop 18
- 6.1 Design Constraints and Operating Conditions . . . 18
- 6.2 Thermosyphon Loop MATLAB Model . . . 19
- 6.3 Parametric Study for Diameter Selection . . . 21
- 6.4 Flow Regime Prediction . . . 23
  - 6.4.1 Superficial velocity map . . . 23
  - 6.4.2 Superficial Reynolds number map . . . 23
  - 6.4.3 Reconstruction of a mass flux against quality map . . . 24
  - 6.4.4 Mechanistic map . . . 27
  - 6.4.5 Common conclusion and limitations . . . 28
- 6.5 Final Diameter Selection . . . 28
- 6.6 Condenser Sizing . . . 29
- 6.7 Heater Sizing . . . 32
- 6.8 Expansion Vessel Concept . . . 32

6.9 Building the experimental setup . . . 34

7 Experimental Campaign **38**
7.1 Sensor Calibration . . . 38
7.1.1 Calibration Procedure . . . 38
7.1.2 Fitting and uncertainty . . . 39
7.1.3 Limitations . . . 40
7.2 Operating Procedure . . . 40
7.3 Data Acquisition and Monitoring Software . . . 41
7.3.1 Measurement Uncertainty . . . 43

8 Results **44**
8.1 Circulation & Two-phase flow visualization . . . 44
8.2 Saturation control and the limit at low charge . . . 46
8.3 Repeatability . . . 47
8.4 The rapid circulation state . . . 49
8.5 The transition . . . 51
8.6 The slow circulation state . . . 51

9 Discussion **52**
9.1 The design and the model . . . 53
9.2 The observed behavior . . . 54
9.2.1 Mechanism that can be excluded . . . 54
9.2.2 Incomplete wetting of the heated wall with progressive liquid redistribution . . . 54
9.2.3 Slow net drift of liquid into the expansion vessel . . . 55
9.2.4 Accumulation of non-condensable gas . . . 55
9.2.5 The common difficulty . . . 55

10 Conclusions & Recommendations **56**
10.1 Instrumentation needed to resolve the transition . . . 56
10.2 Further changes to the setup . . . 57
10.3 Conclusions and answering the research questions . . . 58

A Extended Calibration Data **60**
A.1 Setpoints . . . 60
A.2 Choice of polynomial degree . . . 60
A.3 Independence of the samples . . . 61
A.4 Results . . . 61

B Description of the Experimental Setup **63**
B.1 Piping and connectors . . . 63
B.2 Heating System and Control . . . 63
B.3 Cooling and Condenser System . . . 64
B.4 Pressure Control and Expansion System . . . 66
B.5 Instrumentation and Data Acquisition . . . 68
B.6 Mechanical Structure and Mounting . . . 70
B.7 Safety Systems . . . 70

References **71**

List of Figures

1 Two-phase flow regimes in horizontal pipes . . . 4
2 Flow pattern map of Mandhane et al. for horizontal air-water flow . . . 5
3 Superficial Reynolds number map of Ghajar and Tang . . . 6
4 Generalized flow regime map of Taitel and Dukler . . . 7
5 Mass flux against vapor quality map for refrigerant R-410A . . . 7
6 Hydrogen nuclei inside an applied magnetic field . . . 10
7 Air-water slug flow visualization using MRI . . . 11
8 Schematic diagram of a thermosyphon loop . . . 12
9 3D model of experimental setup . . . 19
10 Flowchart of MATLAB model . . . 20
11 Outlet vapor quality and mean liquid velocity. . . . 22
12 Void fraction and circulation rate. . . . 22
13 Operating trajectories on the Mandhane map. . . . 24
14 Operating trajectories on the Ghajar-Tang map. . . . 25
15 Operating trajectories on the Thome and Cioncolini map. . . . 26
16 Operating trajectories on the Taitel-Dukler map. . . . 27
17 Tube-in-tube counterflow heat exchanger . . . 30
18 Total fluid mass against heat input for different tank configurations. . . . 33
19 P&ID of the complete system . . . 35
20 Experimental setup . . . 36
21 Monitoring software interface. . . . 42
22 The measurement section over $1.4\,\mathrm{s}$, in consecutive frames $0.1\,\mathrm{s}$ apart. . . . 45
23 Measured circulation rate in both states on Day 9, against the model prediction. . . . 46
24 Day 2, $FR = 0.48$. TS05 following successive changes to the expansion vessel bath setpoint. . . . 47
25 Liquid redistribution at low charge, leading to partial exposure of the heated wall. . . . 48
26 Circulation rate and heat input on the repeated runs. . . . 48
27 Day 9, rapid state. Measurement section through one pulse cycle. . . . 50
28 Day 9. Circulation rate, loop temperatures and heat input through the transition. . . . 51
29 Day 9, slow state. Condensate on the wall of the measurement section. . . . 52
30 Division of the heat input between sensible and latent transport at $300\,\mathrm{W}$. . . . 53
31 Expansion vessel fitted with a sight glass. . . . 56
32 Autocorrelation of each channel at the $50\,^{\circ}\mathrm{C}$ setpoint. . . . 61
33 Residuals after the first degree correction. Shaded band marks the $\pm 0.1\,\mathrm{K}$ acceptance criterion. 62
34 Three-part union coupling ($Cu \rightarrow$ PVC-U) . . . 63
35 Heater assembly. (a) The machined aluminum block. (b) The block with the six cartridge heaters installed and the rock wool insulation fitted around it. . . . 64
36 Heater power control and block temperature monitoring. (a) Control box containing the variable autotransformer. (b) Type K thermocouple used to monitor the aluminum block temperature, with its separate display. . . . 65
37 Tube-in-tube condenser and water bath for controlling inlet temperature and flow rate. . . . 65
38 Vaccum pump, Manometer and vaccum gauge . . . 66
39 Expansion Tank . . . 67
40 Data acquisition hardware . . . 69
41 Krohne AF-E 400 electromagnetic flow meter installed on the bottom pipe. . . . 69
42 Mechanical support components used for mounting the thermosyphon loop. . . . 70

Nomenclature

Abbreviations

- MRI : Magnetic Resonance Imaging
- NMRI : Nuclear Magnetic Resonance Imaging
- RF : Radio Frequency
- TPLT : Two-Phase Loop Thermosyphon
- PIV : Particle Image Velocimetry
- EIT : Electrical Impedance Tomography
- EPI : Echo Planar Imaging

Units

- Pa, MPa : Pressure units
- T : Tesla
- W : Power
- m : Length
- m^2, m^3 : Area, Volume
- kg/s : Mass flow rate
- kg/m^2s : Mass flux
- m/s : Velocity
- kg/m^3 : Density

Symbols and Units

- ΔP : Pressure difference (Pa)
- f : Darcy friction factor (-)
- L : Pipe length (m)
- D : Pipe diameter (m)
- ρ : Density (kg/m^3)
- ρ_{mix} : Mixture density (kg/m^3)
- Re : Reynolds number (-)
- ϵ : Pipe roughness (m)
- g : Gravitational acceleration (m/s^2)
- H : Loop height (m)

MRI Parameters

- B_0 : Static magnetic field strength (T)
- ω_0 : Larmor frequency (rad/s)
- γ : Gyromagnetic ratio (rad/s·T)
- ϕ : Phase shift (rad)
- v : Velocity (m/s)
- G_{MRI} : MRI gradient strength (T/m)
- δ : Gradient pulse duration (s)
- Δ : Gradient pulse separation time (s)

Flow Parameters

- x : Vapor quality (-)
- α : Void fraction (-)
- G : Mass flux (kg/m^2s)
- u_{drift} : Drift velocity (m/s)
- C_0 : Distribution parameter (-)
- σ : Surface tension (N/m)
- Fr : Froude number (-)
- We : Weber number (-)
- $\dot{m}$: Mass flow rate (kg/s)
- Q : Heat input (W)
- h_{fg} : Latent heat of vaporization (J/kg)
- FR : Filling ratio (% or -)

1 Introduction

Two-phase flows play a central role in many engineering applications involving heat transfer, including power generation, electronics cooling, and thermal management systems. The interaction between liquid and vapor phases enables highly efficient heat transport, but also gives rise to complex flow behavior that is difficult to predict [1].

Despite their widespread occurrence, two-phase flows remain challenging to characterize in detail. The interaction between phases leads to complex flow regimes, spatially varying phase distributions, and strong coupling between heat transfer and fluid dynamics. Because the arrangement of liquid and vapor governs how effectively heat leaves the wall and how freely liquid returns to it, the flow regime sets the limits of operation [2]. Accurate understanding and prediction of these effects are therefore essential for reliable design and optimization.

Experimental investigation of two-phase flow is inherently difficult. The coexistence of two phases with different densities, velocities, and interfacial structures results in a flow that is highly transient and spatially complex [1], making it difficult to characterize using conventional measurement approaches. Such techniques are often intrusive or limited to surface observations, making it challenging to obtain detailed information about the internal structure of the flow. In particular, capturing both the spatial phase distribution and the velocity of liquid and vapor phases within opaque systems remains a significant limitation [3].

Magnetic Resonance Imaging (MRI) offers a unique solution to this problem. As a non-invasive imaging technique based on the detection of hydrogen nuclei, MRI enables visualization of liquid-phase distribution inside opaque geometries without disturbing the flow [4]. Since gas phases produce negligible signal, MRI inherently provides strong contrast between liquid and vapor, making it particularly suitable for studying gas–liquid two-phase systems [5]. In addition, MRI can provide velocity and temperature information through phase-sensitive measurement techniques, allowing simultaneous access to both structural and dynamic flow properties [6].

To investigate two-phase flow under controlled conditions, a thermosyphon loop is employed as an experimental platform. Thermosyphons rely on natural circulation driven by density differences between liquid and vapor phases, providing a well-defined system in which two-phase flow develops without the need for mechanical pumps. This makes them particularly suitable for studying flow behavior and phase distribution in a closed-loop configuration.

The application of MRI to such systems introduces a new set of challenges. Experimental setups must be fully compatible with the strong magnetic field and electromagnetic environment of the scanner, imposing strict constraints on materials, geometry, and instrumentation. At the same time, the system must maintain stable two-phase natural circulation under controlled conditions.

This thesis addresses these challenges through the design, modeling and experimental investigation of an MRI-compatible two-phase loop thermosyphon. The loop is sized with a steady-state pressure balance model, built from materials and in a geometry that allow the measurement section to be placed inside a scanner bore, and operated across a range of charges and heat inputs to establish how it behaves and how far that behavior can be predicted. The measurements are then compared against the model to find where it holds and where it fails, and the loop as built is weighed against what an MRI measurement of two-phase flow would require, which sets out the changes still needed before it is taken to the scanner.

2 Fundamentals of Two-Phase Flow

A liquid and its vapor flowing together in the same pipe do not travel at the same velocity, and they do not spread themselves evenly over the cross-section. The way the two arrange themselves depends on how much of each is flowing, and it governs the pressure drop and the heat transfer that follow from it. Describing such a flow therefore requires a set of quantities of its own, and predicting the form it will take requires tools built for the purpose. Both are set out below.

2.1 Describing a Two-Phase Flow

The state of a two-phase flow at any point in a pipe is described by a small number of quantities, which are defined here and used throughout the remainder of this work.
The total mass flow rate circulating in the pipe is the sum of the liquid and vapor contributions,

$$\dot{m} = \dot{m}_l + \dot{m}_v, \tag{1}$$

and dividing it by the pipe cross-sectional area gives the mass flux,

$$G = \frac{\dot{m}}{A}, \tag{2}$$

which expresses the same flow rate per unit area and is the variable in which most two-phase correlations are written. The vapor quality is the fraction of that mass flow carried as vapor,

$$x = \frac{\dot{m}_v}{\dot{m}_l + \dot{m}_v}, \tag{3}$$

so that $x = 0$ corresponds to pure liquid and $x = 1$ to pure vapor [7].
The void fraction α is the share of the pipe cross-section taken up by vapor,

$$\alpha = \frac{A_v}{A}, \tag{4}$$

where A_v is the area occupied by vapor and A is the total cross-sectional area of the pipe [7]. Since the interface between the phases moves continuously, A_v is understood as an average over the measurement interval rather than an instantaneous value. A void fraction of zero means the pipe is running full of liquid and a value of one means it is running full of vapor.
Quality and void fraction describe different things and are easily confused: quality is a fraction of the mass flowing, while the void fraction is a fraction of the area occupied. The void fraction is the quantity most often sought from a flow visualization measurement, and several of the techniques described in chapter 2.4 are judged by how well they resolve it.

2.2 Two-Phase Flow Regimes in Horizontal Pipes

The spatial distribution of liquid and vapor inside a pipe is classified into a small number of recognizable flow regimes. Which regimes are encountered depends on the relative flow rates of the two phases, but also on the orientation of the pipe: gravity acts along the flow direction in vertical sections and perpendicular to it in horizontal sections, producing distinctly different regime structures. Since the pipe section imaged in the present work is horizontal, the description below focuses on horizontal two-phase flow regimes, illustrated in Figure 1.
In *bubbly flow*, small vapor bubbles are dispersed in a continuous liquid phase, tending towards the upper half of the pipe under the action of buoyancy unless the liquid velocity is high enough to redistribute them across the cross-section. At low gas and liquid velocities, buoyancy dominates completely and the two phases

separate: vapor flows along the top of the pipe and liquid along the bottom with a smooth interface (*stratified flow*). As the gas velocity increases, interfacial shear produces traveling waves on the liquid surface that do not reach the top of the pipe (*stratified-wavy flow*).
At intermediate vapor fractions, elongated bubbles develop along the top of the pipe. When these bubbles are smaller than the pipe diameter and the liquid remains continuous beneath them, the flow is in the *plug* regime, and as the bubbles grow to approximately the pipe diameter, the flow is in the *slug* regime. Plug and slug flow are often grouped together as *intermittent flow*, characterized by a cyclic alternation of small and large amplitude waves where the larger waves intermittently wet the top of the pipe.
At higher gas flow rates the liquid forms a continuous annular film around the pipe perimeter with a high-velocity vapor core in the center (*annular flow*). In horizontal flow the film is typically thicker at the bottom of the pipe than at the top due to gravity, with the asymmetry decreasing as the gas velocity increases. Finally, at very high gas velocities the interfacial shear strips the liquid film from the wall and entrains it as droplets in the gas core (*mist flow*), in which the gas is the continuous phase and the wall is essentially dry.
The flow regime present in the measurement section at any operating condition directly influences what MRI can resolve: bubbly, plug, slug and intermittent regimes contain coherent liquid structures on the millimeter scale that are well within the scanner's spatial resolution, while annular and mist regimes contain progressively thinner liquid films and dispersed droplets that approach the resolution limit. The choice of pipe diameter and operating point in the design chapter must therefore account for which regimes are expected in the measurement section under the planned operating range [6, 8].

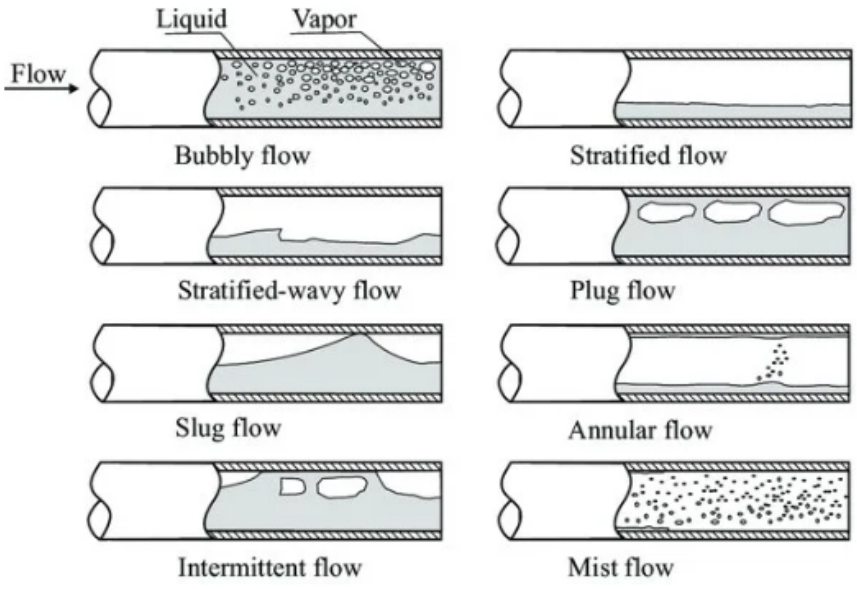


Figure 1: Two-phase flow regimes in horizontal pipes [9].

2.3 Flow Pattern Maps

For a given pipe and a given pair of flow rates, the regime that actually develops has to be predicted, and this is done with a flow pattern map. A map is a plot whose two axes are flow parameters and whose interior is divided into regions by transition boundaries, so that an operating point placed on it falls inside one region and is expected to show that regime. Maps differ in two respects, the coordinates chosen for the axes and whether the boundaries were drawn through observed data or derived from a force balance. Both choices decide how far a map can be carried away from the conditions it was built for.

2.3.1 Superficial velocity maps

The oldest and most widely used coordinates are the superficial velocities of the two phases. The superficial velocity of a phase is the velocity that phase would have if it flowed alone through the whole pipe cross-section,

$$J_l = \frac{\dot{m}\,(1-x)}{\rho_l A}, \qquad J_g = \frac{\dot{m}\,x}{\rho_v A}, \tag{5}$$

where A is the cross-sectional area of the pipe. Neither phase in fact has the pipe to itself, so these are not the real phase velocities, but they are convenient because they follow directly from the mass flow rate and the quality without requiring the void fraction.

Mandhane et al. [10] placed the boundaries between the bubbly, elongated bubble, slug, stratified, wavy, annular and dispersed regimes on these axes, using a large bank of air-water observations taken close to atmospheric pressure. Their map is shown in Figure 2. It is purely empirical, in the sense that the lines are fitted to where the regime was seen to change and carry no dependence on fluid properties beyond those of the data they came from. Applying the map to a different fluid pair, or to the same pair at a very different pressure, is therefore an extrapolation. Kong and Kim [11] used the same coordinates for horizontal air-water flow and identified stratified, slug, churn and annular regimes at conditions comparable to those expected in the present setup.

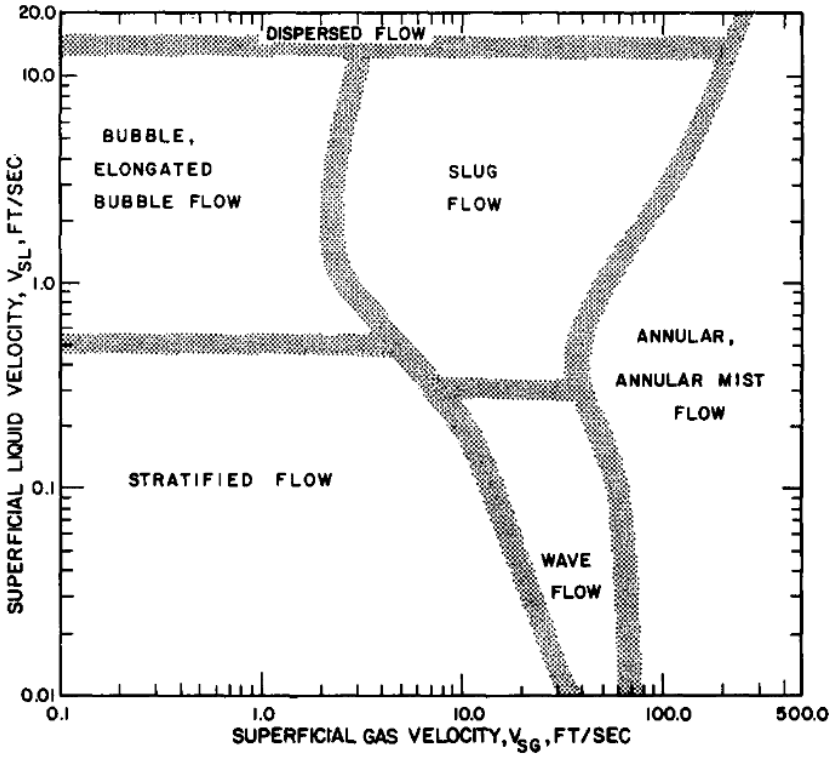


Figure 2: Flow pattern map of Mandhane et al. for horizontal air-water flow [10].

2.3.2 Superficial Reynolds number maps

Replacing the superficial velocities by the superficial Reynolds numbers of the two phases,

$$Re_{SG} = \frac{\dot{m}\,x\,D}{A\,\mu_v}, \qquad Re_{SL} = \frac{\dot{m}\,(1-x)\,D}{A\,\mu_l}, \tag{6}$$

folds the pipe diameter and the viscosity of each phase into the coordinates themselves. The axes are then dimensionless, so pipes of different size and fluids of different viscosity can be compared on a single plane, which a velocity-based map does not allow. Ghajar and Tang [12] presented maps in these coordinates for

non-boiling air-water flow in horizontal and slightly inclined pipes, shown in Figure 3.

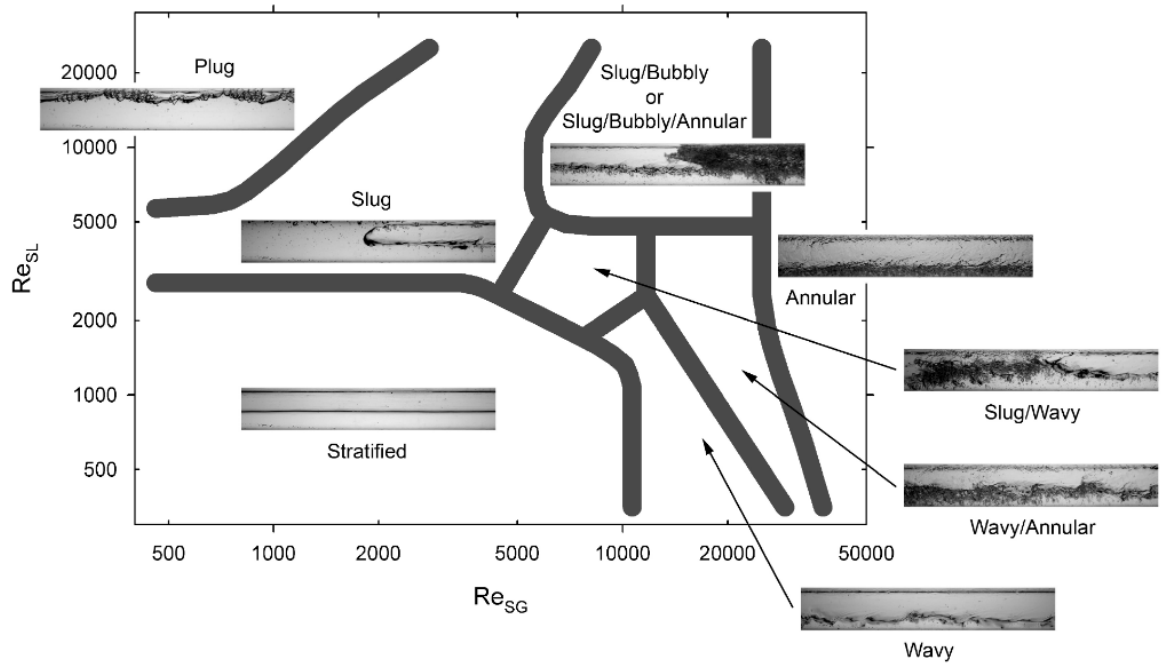


Figure 3: Superficial Reynolds number map of Ghajar and Tang [12].

2.3.3 Mechanistic maps

Taitel and Dukler [13] took a different route and derived each transition from a physical argument rather than from a fit. The starting point is the equilibrium liquid level that an imaginary stratified layer would settle at, obtained from a balance between the interfacial shear and the weight of the liquid. Each transition is then tested against that geometry. Whether stratified flow survives is decided by whether a disturbance on the interface grows, which they treat with a Kelvin-Helmholtz argument, and whether the result of its breakdown is intermittent or annular flow is decided by how deep the equilibrium layer is. The transition to dispersed bubble flow is decided by whether turbulence in the liquid is strong enough to overcome buoyancy and hold bubbles in suspension.

Because these criteria are written in terms of the actual densities, viscosities and surface tension, the map can be constructed for any fluid pair at any pressure, which is the main advantage of a mechanistic map over an empirical fit. The result is normally presented in dimensionless groups rather than in velocities, with the Martinelli parameter on one axis and the three groups that govern the individual transitions on the other, as in Figure 4.

2.3.4 Mass flux and quality maps

In a boiling system neither the superficial velocities nor the superficial Reynolds numbers are known in advance, since both depend on the quality, which is itself an outcome of the heat input. Maps drawn on the mass flux G and the vapor quality x, both defined in Section 2.1, are easier to use in that situation, because those are the two quantities by which an evaporator is specified. The price is that the boundaries are no longer universal curves. They have to be evaluated from closed-form expressions using the properties of the fluid and the diameter of the tube, so every fluid, pressure and diameter has a map of its own.

The framework of this type used in the present work is the one given by Thome and Cioncolini [14], which takes the geometric treatment of Taitel and Dukler and recasts the transitions as explicit expressions for the mass flux at which each one occurs. It was developed and validated on a database of refrigerants, with

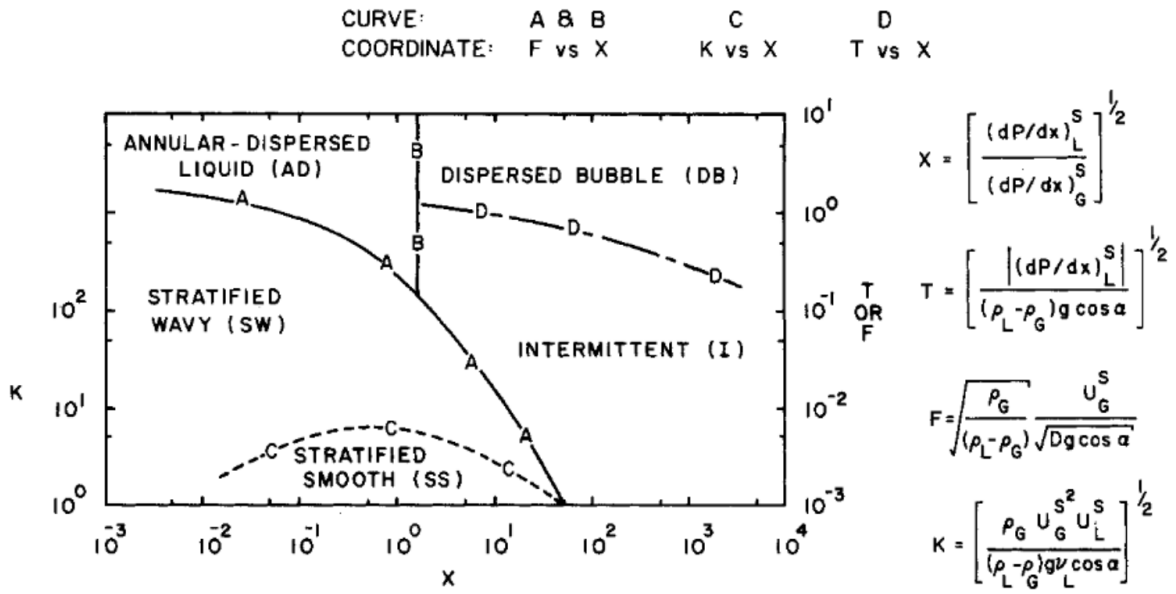


Figure 4: Generalized flow regime map of Taitel and Dukler [13].

liquid-to-vapor density ratios of roughly 30 to 250, and an example of the resulting map is shown in Figure 5. No equivalent map has been published for water and steam at sub-atmospheric pressure, where the density ratio is two orders of magnitude larger, so one is constructed from the same equations for the conditions of the present loop in Section 6.4.3.

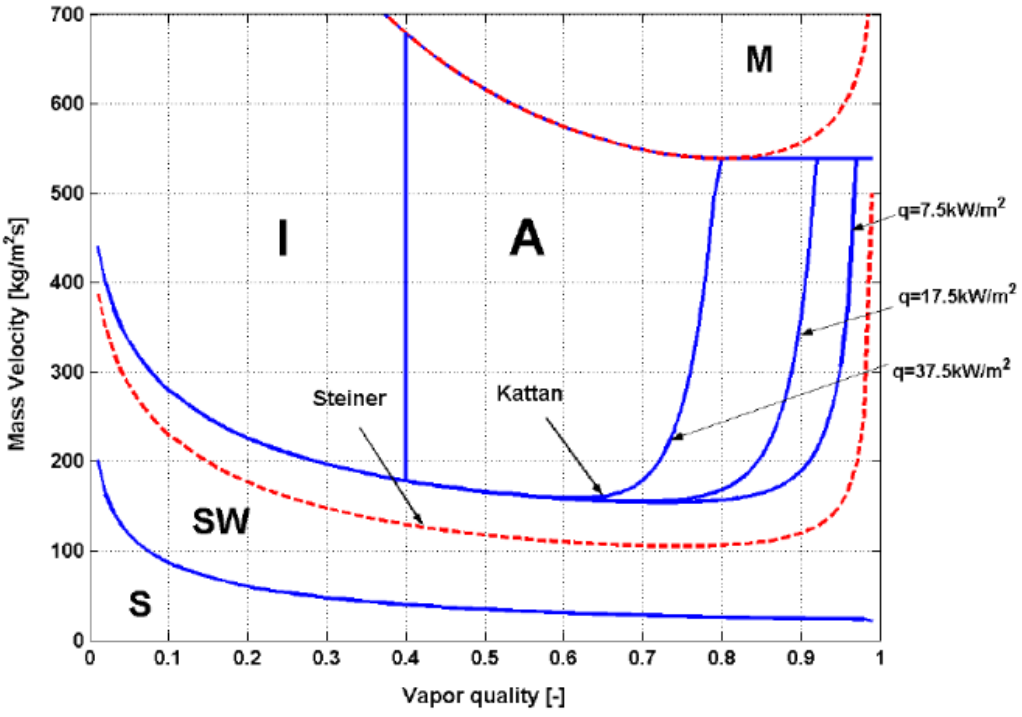


Figure 5: Mass flux against vapor quality map for refrigerant R-410A [14].

Every map described above was built from measurements taken at or above atmospheric pressure, with either air and water or a refrigerant as the working pair. None of them covers water and steam at sub-atmospheric pressure, which is the condition of interest here. How far they can still be used, and what the very large density ratio of the present system does to the position of the boundaries, is examined in Section 6.4.3.

2.4 Flow Visualization Techniques in Two-Phase Flow

Understanding the internal flow behavior of two-phase systems is essential for validating theoretical models and improving system design. Over the years, a range of experimental visualization techniques have been developed and applied to capture the complex phenomena that occur within these systems, including phase distribution and flow pattern transitions. This section provides an overview of the most prominent methods used in the literature, along with their results, ranging from classical optical approaches to advanced visualization methods.

2.4.1 Optical Imaging

Optical imaging techniques were among the first methods applied to visualize two-phase flows and remain widely used due to their relatively simple implementation and high spatial resolution. These methods rely on the transmission or reflection of visible light through a transparent test section, making them well suited to laboratory-scale systems with glass or acrylic pipe walls.
High-speed photography and video recording allow direct observation of flow patterns such as bubbly flow, slug flow, annular flow, and their transitions. When combined with backlighting, the contrast between the liquid and vapor phases is sufficient to identify phase boundaries and measure bubble sizes and velocities. Particle Image Velocimetry (PIV) extends this capability by seeding the liquid phase with tracer particles and using cross-correlation of successive image frames to reconstruct two-dimensional velocity fields [15].
The primary limitation of optical methods is that they require optical access to the flow. This means they can only be used with transparent pipe materials and relatively simple geometries. In addition, optical techniques are inherently limited to the pipe surface as they cannot resolve the internal phase distribution across the full cross-section of an opaque or thick-walled pipe, and their performance degrades significantly at high void fraction, where the dense bubble population scatters light.

2.4.2 X-ray Visualization

X-ray based techniques overcome the optical access limitation by using high-energy radiation that can penetrate opaque pipe walls. The attenuation of X-rays as they pass through the flow depends on the local density, creating contrast between the liquid and vapor phases. This makes X-ray imaging particularly effective for measuring void fraction distributions and identifying flow patterns in metallic and optically opaque systems. Aliseda et al. [16] demonstrated the use of X-ray imaging for quantitative void fraction measurement in two-phase flows, showing that the technique can resolve cross-sectional phase distributions with high temporal resolution. X-ray tomography further extends this to three-dimensional reconstructions of the phase distribution, but at a higher cost and increased system complexity.
The main practical limitations of X-ray methods are the radiation safety requirements, the cost and size of the equipment, and the fact that the technique measures line-averaged attenuation rather than a direct local measurement unless full tomographic reconstruction is performed.

2.4.3 Other Measurement Techniques

Several other techniques have been developed for cases where optical access is unavailable and X-ray equipment is impractical.

- Electrical Impedance Tomography (EIT) reconstructs the cross-sectional distribution of electrical conductivity from boundary measurements, exploiting the difference in conductivity between liquid and vapor phases. It offers high temporal resolution and does not require optical access, but its spatial resolution is limited compared to optical and X-ray methods [17].
- Ultrasonic techniques use acoustic pulses transmitted through the pipe wall to measure local void fraction and velocity. They are non-invasive and applicable to metallic pipes, but their accuracy degrades

once the bubble density is high, since each pulse is reflected off many bubble interfaces before it returns. [18].

- Neutron radiography operates similarly to X-ray imaging but uses neutron beams, which interact differently with matter as hydrogen-rich fluids attenuate neutrons strongly, providing excellent contrast for water-based flows inside metallic containment. However, the technique requires access to a neutron source, typically a nuclear reactor facility, which severely limits its practical applicability [19].
- Magnetic Resonance Imaging (MRI) represents a further alternative that combines the ability to penetrate opaque walls with the capability to measure both phase distribution and velocity fields simultaneously and non-invasively [20], making it particularly well suited to water-based two-phase flows. Its physical principles and application to multiphase flow measurement are discussed in detail in the following chapter.

3 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a non-invasive imaging technique based on the interaction of magnetic moments with an externally applied magnetic field. Unlike optical or radiographic methods, MRI does not rely on light transmission or density differences, but instead detects signals from hydrogen atoms.
MRI has become a widely used imaging technique in the medical industry mainly because it relies on the detection of hydrogen nuclei, which are abundant in water. Since the human body consists of ∼70% water, it provides a strong and easily detectable MRI signal. The presence of hydrogen in biological tissues makes MRI particularly effective for internal imaging of the human body.
The same principle can be extended to engineering systems that contain hydrogen-rich fluids. Since liquid water produces a strong MRI signal due to its high hydrogen density, water-based flows can also be investigated using MRI. This allows the technique to be applied not only in medical imaging but also in flow visualization and velocity measurements in fluid systems.

3.1 Working Principle of MRI

MRI is based on the magnetic behavior of hydrogen nuclei. Each hydrogen nucleus consists of a single proton, which behaves like a small magnet due to its intrinsic spin. When a strong static magnetic field B_0 is applied, these spins tend to align with the field. Instead of remaining fixed, they precess around the field direction at a specific frequency known as the Larmor frequency, as seen in Figure 6:

$$\omega_0 = \gamma B_0 \tag{7}$$

where γ is the gyromagnetic ratio and B_0 is the magnetic field strength.
To generate a measurable signal, a radiofrequency (RF) pulse is applied at the Larmor frequency. This pulse temporarily disturbs the alignment of the protons, tipping them away from the direction of B_0. When the RF pulse is turned off, the protons gradually return to their original state. During this process, they emit a signal that can be detected by a receiver coil.
The strength and decay of this signal depend on how the hydrogen nuclei interact with their surroundings. In liquids such as water, where hydrogen is abundant and mobile, the signal is strong. In contrast, gas phases contain far fewer hydrogen nuclei and therefore produce little to no signal. This natural contrast between liquid and gas phases is particularly useful for two-phase flow studies.
To form an image, spatial information must be encoded into the signal. This is achieved by applying magnetic field gradients, which slightly vary the magnetic field strength across space. As a result, protons at different locations rotate at slightly different frequencies. By analyzing these frequency differences, the spatial distribution of hydrogen, and therefore the fluid, can be reconstructed.

In this way, MRI converts the behavior of hydrogen nuclei in a magnetic field into a spatial map, allowing non-invasive visualization of fluid distribution inside opaque systems.

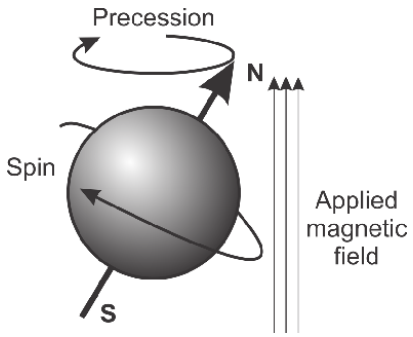


Figure 6: Hydrogen nuclei inside an applied magnetic field

MRI can also be used to obtain information about fluid motion, in addition to producing images. By applying specific magnetic field gradients, moving hydrogen nuclei exhibit a change in the phase of the detected signal. This phase shift is directly related to the velocity of the fluid, allowing spatially resolved velocity fields to be obtained.
In practice, modern MRI systems implement this through dedicated pulse sequences and reconstruct velocity maps automatically, meaning that no manual calculation is required by the user. Although the primary focus of this work is on flow visualization, the ability of MRI to provide both structural and velocity information makes it a powerful tool for studying complex flow phenomena [6, 8].

3.2 MRI in Multiphase Flow Studies

The application of MRI to flow measurement has a long history rooted in medical imaging. Early studies demonstrated that phase-contrast velocity encoding could measure blood flow velocities quantitatively in arteries and veins, including flow direction, flow regime, and volumetric flow rate [21, 22, 23]. These medical studies established the fundamental measurement methodology and produced the image reconstruction and post-processing software that later became available for engineering applications.
The technique was subsequently extended to industrial single-phase and porous media flows. Chen et al. [24] applied Nuclear Magnetic Resonance imaging (NMRI) to multiphase flow in porous geological media, exploiting the high proton density of oil relative to gas to distinguish phases and measure saturation distributions. Kose [25, 26] applied ultra-fast spin-echo MRI sequences to turbulent pipe flow, successfully imaging instantaneous velocity distributions during transitional and turbulent regimes and demonstrating that multiple spin-echo EPI sequences could capture all three velocity components simultaneously. These studies showed that MRI could be applied beyond the medical domain to flowing engineering fluids, provided the system materials and geometry were compatible with the scanner environment.
The transition to gas-liquid two-phase flow studies required further methodological development, since the highly inhomogeneous and time-varying nature of two-phase flow demands both high spatial resolution and short acquisition times. Gladden [8] demonstrated that phase-sensitive MRI techniques using pulsed field gradients could provide simultaneous quantitative measurements of velocity fields and phase distribution in gas-liquid systems.
This work highlighted that MRI is non-invasive, penetrates opaque pipe walls without disturbing the flow, and is particularly effective for water-based systems due to the strong hydrogen signal from the liquid phase. Gas phases produce negligible signal, which is a limitation for directly imaging the gas, but is simultaneously an advantage for void fraction measurement, since the absence of signal in a voxel, the small volume element of fluid that each point of the reconstructed image represents, directly indicates the presence of gas. The study also identified the key practical limitations of the technique: the requirement for MRI-compatible

materials, the cost and complexity of the scanner, and the temporal resolution constraints in rapidly evolving flow regimes [8].

One of the most detailed early applications of MRI to two-phase flow in pipes was presented by Reyes et al. [6]. In this study, air–water flow patterns such as stratified, slug, and annular flow were successfully visualized inside a horizontal pipe using a clinical MRI scanner. The results demonstrated that MRI is capable of capturing the spatial distribution of the liquid phase with sufficient temporal resolution to resolve dynamic flow structures. An example of their results can be seen in figure 7.

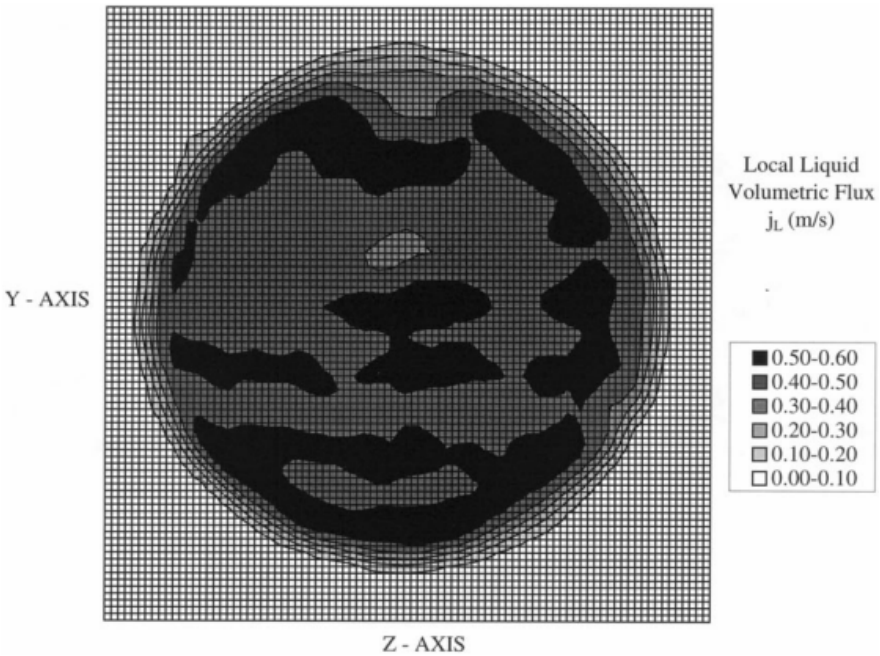


Figure 7: Air-water slug flow visualization using MRI [6].

From the acquired images, several key flow parameters were obtained, including liquid distribution, interfacial structure, and void fraction. The measurements showed good agreement with conventional intrusive techniques, validating the accuracy of MRI for quantitative two-phase flow analysis.

A limitation identified in this work is that the MRI signal intensity alone represents a combination of liquid hold-up and velocity, meaning that these two effects cannot be separated without additional techniques. This can be addressed by combining standard imaging with velocity-sensitive methods, allowing both phase distribution and flow dynamics to be characterized independently [6].

3.3 Experimental Constraints of MRI-Compatible Systems

The MRI environment imposes a set of physical constraints on any experimental system operating inside or near the scanner that must be carefully respected to ensure both measurement quality, and equipment/operator safety.

The most fundamental constraint is the requirement for non-magnetic, non-conducting materials in the vicinity of the scanner bore. Ferromagnetic materials are attracted to the magnet and represent a serious safety hazard. Electrically conducting materials, even non-magnetic ones such as aluminum or copper, distort the static field B_0 and the gradient fields through eddy currents, which are loops of current induced in the metal by the changing fields and which generate stray magnetic fields of their own, degrading image quality and introducing systematic errors in velocity measurements. For this reason, the components of an MRI-compatible flow system that sit in or near the bore must be made from non-magnetic and preferably non-conducting materials, such as plastics or ceramics [6].

The bore of a clinical MRI scanner has a fixed diameter, which imposes a hard geometric limit on the size of any experimental apparatus inserted into it. The usable imaging region is further restricted to the homogeneous field volume at the center of the bore, meaning that the section of the flow system intended for measurement must be positioned precisely at the center of the magnet.
Radiofrequency (RF) interference is an additional concern. Electronic components such as sensors, heaters with resistive elements, or data acquisition systems can emit RF noise at or near the Larmor frequency, corrupting the MRI signal. Therefore, any electrical equipment used in conjunction with MRI measurements must be carefully shielded or operated at a safe distance from the MRI scanner.
Temperature is a further constraint. The resonance signal is itself temperature dependent, so a hot or thermally non-uniform test section alters the measured signal and degrades the image [27]. Any apparatus placed in the bore should therefore operate as close to ambient temperature as the experiment allows.
Finally, the strong static magnetic field can affect the behavior of the flow itself if the working fluid has significant magnetic susceptibility. For water at the field strengths typical of clinical MRI systems (1.5–3 T), these magnetohydrodynamic effects are negligible, but they must be considered when working with fluids containing magnetic particles or dissolved paramagnetic species.

4 Fundamentals of Thermosyphon Loops

A two-phase loop thermosyphon (TPLT) is adopted in this work as the platform on which two-phase flow is to be produced under controlled conditions and later visualized by MRI. The TPLT provides a simple naturally circulating system in which vapor–liquid flow develops without mechanical pumping, making it well suited for MRI-based visualization and analysis. A two-phase loop thermosyphon is a heat transfer device that transports thermal energy from a heat source to a heat sink without any mechanical pump or external energy input. It does so by boiling and condensing its working fluid. Turning liquid into vapor absorbs a large amount of energy at constant temperature, called the latent heat of vaporization, and releasing that same energy is what condenses the vapor back to liquid. TPLTs are used across a wide range of engineering applications, including data center and electronics cooling, where they provide a reliable passive cooling solution for heat dissipation in server racks [28], and solar thermal systems, where heat is transferred from solar collectors to storage tanks without any pumping energy [29].

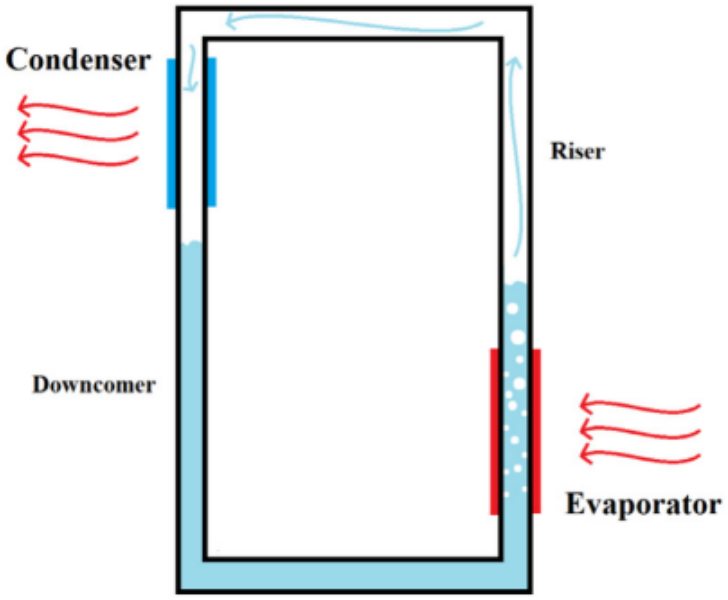


Figure 8: Schematic diagram of a thermosyphon loop [30].

A schematic of the system is shown in Figure 8, and its main components are the evaporator to supply heat, and the condenser to act as a heat sink. The rest of the system is piping, with the important sections being the riser, a vertical pipe located directly after the evaporator, and the downcomer, located directly after the condenser. As can be seen in the figure, the condenser is elevated at a certain height, allowing the condensed liquid to drop directly into the downcomer after cooling with the help of gravity, while the evaporator is placed at the bottom, allowing the produced vapor to rise through the riser, thus creating a natural circulation flow in the closed thermosyphon loop.

4.1 Operating Principle of a Thermosyphon

The operating cycle begins in the evaporator, where the working fluid absorbs heat from the source and boils. The generated vapor rises naturally through the riser due to its density difference with the surrounding liquid, and travels upward towards the condenser. In the condenser, the vapor releases its latent heat to a cooling medium and returns to the liquid phase. The condensate then flows back down through the downcomer under gravity, returning to the evaporator to complete the cycle.

The density difference between the vapor-rich riser and the liquid-rich downcomer creates a pressure difference in the loop that keeps the fluid moving around naturally, meaning that as long as there is a heat source and a heat sink, the system will circulate on its own.

The steady-state operation of the loop is governed by a closed pressure balance. As the fluid completes one full cycle and returns to its starting point, there can be no net pressure change around the loop, therefore, any pressure gained in one section must be exactly recovered elsewhere. This condition can be written as:

$$\Delta P_{total} = 0 \tag{8}$$

Each section of the loop contributes its own pressure change depending on the flow direction and the phase of the fluid present. In the **bottom horizontal pipe**, the single-phase liquid loses pressure to friction. In the **riser**, the two-phase mixture loses pressure to friction and additionally loses pressure due to the hydrostatic weight of the mixture, which opposes the flow direction. Since heat is continuously added along the riser, the fluid evaporates as it rises, therefore the mixture becomes lighter, and the flow accelerates. This acceleration requires an additional pressure drop, known as the accelerational or momentum pressure drop. In the **top horizontal pipe**, the two-phase mixture again loses pressure to friction. In the **downcomer**, the liquid loses pressure to friction but recovers pressure from gravity, as the dense liquid column falling downward generates a hydrostatic pressure gain. Equations 9 to 12 express the pressure contribution of each section individually:

$$\Delta P_{bottom} = -\Delta P_{friction,bottom} \tag{9}$$

$$\Delta P_{riser} = -\Delta P_{friction,riser} - \Delta P_{grav,riser} - \Delta P_{acc,riser} \tag{10}$$

$$\Delta P_{top} = -\Delta P_{friction,top} \tag{11}$$

$$\Delta P_{downcomer} = -\Delta P_{friction,downcomer} + \Delta P_{grav,downcomer} \tag{12}$$

The frictional pressure drop in each section is calculated using the Darcy-Weisbach [31] equation:

$$\Delta P_{friction} = f \cdot \frac{L}{D} \cdot \frac{\rho \cdot u^2}{2} \tag{13}$$

where f is the Darcy friction factor, L is the length of each pipe, D is the pipe diameter, ρ is the fluid density,

and u is the mean fluid velocity.
For the sections of the system where the flow is in a single phase, the Darcy friction factor is obtained from the Churchill correlation [32], which provides a continuous explicit expression for f valid for all single-phase flow regimes:

$$f = 8\left[\left(\frac{8}{Re}\right)^{12} + \frac{1}{(A+B)^{1.5}}\right]^{1/12} \tag{14}$$

$$A = \left[2.457\ln\left(\frac{1}{\left(\frac{7}{Re}\right)^{0.9} + 0.27\frac{\varepsilon}{D}}\right)\right]^{16}, \qquad B = \left(\frac{37530}{Re}\right)^{16} \tag{15}$$

where Re is the Reynolds number, which compares inertial to viscous forces and determines whether the flow is laminar or turbulent, ε is the pipe roughness, and D is the pipe diameter.
For the two-phase sections, the Friedel correlation is used [33], which accounts for the interaction between the vapor and liquid phases by applying a multiplier ϕ^2 to the liquid pressure drop:

$$\Delta P_{friction,two-phase} = \phi^2 \cdot \Delta P_{friction,liquid-phase} \tag{16}$$

where $\Delta P_{friction,liquid-phase}$ is the pressure drop calculated as if the flow were entirely liquid. The multiplier ϕ^2 depends on the vapor quality x, the density ratio ρ_l/ρ_v, the viscosity ratio μ_l/μ_v, the Froude number (Fr), which compares inertial to gravitational forces, and the Weber number (We), which compares inertial forces to surface tension.

The gravitational pressure drop in the riser depends on the local density of the two-phase mixture. Since the vapor quality changes along the riser as heat is added, the mixture density varies with position and must be evaluated locally. The mixture density is expressed as:

$$\rho_{mix} = \alpha\rho_v + (1-\alpha)\rho_l \tag{17}$$

where α is the void fraction defined in Equation (4). Vapor and liquid do not travel at the same velocity. The vapor is far less dense and moves faster, so it occupies a larger share of the cross-section than its share of the mass flow would suggest, which is why α is larger than x. The ratio between the two velocities is the slip ratio,

$$s = \frac{u_v}{u_l} = \frac{x}{1-x}\cdot\frac{1-\alpha}{\alpha}\cdot\frac{\rho_l}{\rho_v}, \tag{18}$$

which equals one only when the two phases move together [7]. The void fraction must therefore be obtained from a correlation that accounts for this slip rather than from the quality alone. Here the drift-flux model of Rouhani and Axelsson [34] is used, which treats the mixture as a single stream and corrects for the relative motion between the phases:

$$\alpha = \frac{x}{\rho_v}\left[C_0\left(\frac{x}{\rho_v} + \frac{1-x}{\rho_l}\right) + \frac{u_{drift}}{G}\right]^{-1} \tag{19}$$

where G is the mass flux, C_0 is the distribution parameter accounting for the non-uniform velocity and void fraction profiles across the cross-section, and u_{drift} is the drift velocity representing the buoyancy-driven relative motion between the phases:

$$u_{drift} = 1.18\left[\frac{\sigma g(\rho_l - \rho_v)}{\rho_l^2}\right]^{0.25} \tag{20}$$

with σ representing the surface tension. The distribution parameter takes the value $C_0 = 1.54$ when the mass flux $G < 200$ kg/m^{2}s, which corresponds to the natural circulation regime of the present system, and

$C_0 = 1.12$ otherwise [34]. However, the Rouhani-Axelsson correlation was developed and validated for steam-water flows at pressures significantly higher than those used in this work. Chexal and Lellouche [35], reviewing the most widely used void fraction correlations, report validated pressure ranges starting from 1 bar. The operating pressure of the present system is 0.12 bar, which lies an order of magnitude below this lower bound. At this pressure, the liquid-to-vapor density ratio is $\rho_l/\rho_v \approx 12{,}200$, far exceeding the range of conditions represented in the experimental databases used to develop existing correlations. The Rouhani-Axelsson model is therefore adopted as the best available option for low-mass-flux natural circulation, while acknowledging that its predictions carry significant uncertainty at sub-atmospheric conditions. The experimental void fraction measurements obtained in a future work will provide data in a pressure regime that is not currently covered by the validation databases of existing correlations.
The gravitational pressure drop in the riser is then obtained by integrating the local mixture weight along the pipe length:

$$\Delta P_{grav,riser} = \int_0^L \rho_{mix}(z)\, g\, dz \tag{21}$$

The accelerational pressure drop arises from the change in momentum of the two-phase flow as the fluid evaporates and accelerates along the heated section. In the separated-flow formulation, each phase carries momentum proportional to its own velocity, which differs from the other phase due to slip. The accelerational pressure drop is therefore expressed as:

$$\Delta P_{acc} = G^2 \left[\frac{(1-x)^2}{(1-\alpha)\rho_l} + \frac{x^2}{\alpha\rho_v}\right]_{out} - G^2 \left[\frac{(1-x)^2}{(1-\alpha)\rho_l} + \frac{x^2}{\alpha\rho_v}\right]_{in} \tag{22}$$

where the terms in brackets represent the momentum flux of the separated liquid and vapor phases at the outlet and inlet of the heated section respectively. The liquid term $(1-x)^2/[(1-\alpha)\rho_l]$ accounts for the momentum carried by the liquid fraction flowing at its actual velocity, and the vapor term $x^2/[\alpha\rho_v]$ accounts for the momentum carried by the vapor fraction, with α evaluated using the Rouhani-Axelsson model from equation 19] at each location. This term should only appear in the heated section of the riser, since quality, and therefore momentum, only changes where heat is applied.

The driving force of the loop is the hydrostatic pressure recovered in the downcomer. The dense liquid column falling under gravity generates a pressure gain:

$$\Delta P_{grav,downcomer} = \rho_l \cdot g \cdot H \tag{23}$$

where H is the vertical height of the loop. This term is the engine of natural circulation: as long as the downcomer contains denser liquid than the riser contains two-phase mixture, a net pressure difference drives the flow around the loop.
At steady state, the sum of all pressure contributions around the complete loop must equal zero [eq. 8]:

$$\Delta P_{bottom} + \Delta P_{riser} + \Delta P_{top} + \Delta P_{downcomer} = 0 \tag{24}$$

For a given heat input Q, the outlet quality is linked to the mass flow rate through the energy balance at the evaporator:

$$x_{out} = x_{in} + \frac{Q}{\dot{m} \cdot h_{fg}} \tag{25}$$

where h_{fg} is the latent heat of vaporization. This condition, combined with the continuity of mass flow around the loop, fully determines the steady-state operating point for a given heat input, fluid properties, and loop geometry.

4.2 Influence of Filling Ratio

The filling ratio (FR) is defined as the ratio of the liquid volume charged into the system to the total internal volume of the loop:

$$FR = \frac{V_{liquid}}{V_{total}} \cdot 100\% \tag{26}$$

It is one of the most important parameters governing the behavior of a thermosyphon loop, as it determines how the liquid inventory is distributed across the system during operation.
At rest, before the system starts, the liquid fills the loop from the bottom upward and the level equalizes on both sides of the loop following the principle of communicating vessels. When the system begins operating and boiling occurs in the evaporator, the liquid inventory redistributes, part of it evaporates and occupies the two-phase sections as vapor, while the remainder stays as liquid in the single-phase sections. The filling ratio therefore directly controls how much liquid is available to sustain circulation.
Experimental studies have shown that over a wide range of filling ratios, the heat transfer coefficient of the evaporator remains nearly constant [2]. This is because the parameter directly governing evaporator performance during operation is the liquid level in the evaporator itself, rather than the total amount of liquid charged into the system. As long as the liquid level is sufficient to keep the heated wall wetted, and the condenser provides adequate area for condensation, the system operates under stable conditions and performance is independent of the exact filling ratio. However, when the filling ratio becomes too low or too high, the system deviates from this stable behavior and its performance drops.

4.3 Heat Transfer Limits

The heat transfer performance of a thermosyphon is strongly influenced by the prevailing two-phase flow regime and the distribution of the liquid phase along the heated wall. In bubbly and slug flow, the presence of liquid in contact with the wall promotes stable nucleate boiling and efficient heat transfer. In annular flow, heat transfer is sustained by a thin liquid film along the wall, which can provide high heat transfer rates but is more susceptible to instability.
Under certain operating conditions, these mechanisms break down, leading to limitations in heat transfer performance. The most important limits are dryout and flooding, which define the operational boundaries of the thermosyphon.

4.3.1 Dryout Limit

The dryout limit is reached when the filling ratio is too low relative to the applied heat input. When the liquid inventory is insufficient, the heated wall of the evaporator can no longer be kept fully wetted by a liquid film. As the heat flux increases, the vapor generation rate rises, the void fraction in the heated section increases, and the liquid film becomes thinner until it breaks down locally. Once a part of the heated wall is exposed directly to vapor, the local heat transfer coefficient collapses, vapor conducts heat far less effectively than a boiling liquid film, and the wall temperature rises sharply and irregularly.
Ding et al. [2] observed this behavior experimentally in a separated heat pipe system. At a filling ratio of 65% and a heating capacity of 597.2 W, the wall temperature at the evaporator exit became irregular and fluctuating, in contrast to the stable temperatures observed at lower heat inputs at the same filling ratio. The corresponding flow pattern at the evaporator exit transitioned from stable annular flow to a more chaotic regime, consistent with intermittent liquid film breakdown.
The dryout limit therefore defines a lower bound on the filling ratio for a given heat input. Below this bound, the system cannot transfer heat reliably and the evaporator wall is at risk of overheating.

4.3.2 Flooding Limit

The flooding limit is reached when the filling ratio is too high. When the liquid inventory is excessive, the condenser becomes partially or fully submerged in liquid before the system starts. When heat is applied and boiling begins, vapor and liquid flow together into the downcomer pipe, but the available condensation area in the condenser is insufficient to condense all the vapor leaving the evaporator. The heat pipe therefore cannot reach steady state and instead operates in a cyclical on-off fashion: vapor builds up, a small amount condenses, the liquid level in the condenser temporarily drops exposing a slightly larger wall area, brief condensation occurs, and the cycle repeats.
Ding et al. [2] characterized this behavior experimentally and identified three observable symptoms when the flooding limit is reached: the overall temperature difference in the heat pipe increases, the liquid pipe becomes overcooled, and the temperature of the evaporator wall oscillates periodically with time. At a filling ratio of 293%, where the entire system except the top horizontal pipe was submerged, these oscillations were clearly visible in the thermocouple records.
The flooding limit therefore defines an upper bound on the filling ratio. Above this bound, the condenser cannot function continuously and steady-state circulation cannot be maintained. It should be noted that the filling ratio of 293% reported by Ding et al. [2] is defined relative to the volume of the downcomer alone, rather than the total internal volume of the loop. Using the definition adopted in the present work, where $FR = V_{liquid}/V_{total}$, this value corresponds to a filling ratio well above the flooding limit but below 100%.

4.3.3 Operating Window

Between the dryout and flooding limits there exists a range of filling ratios within which the system operates normally. Within this window, as noted above, performance is relatively insensitive to the exact filling ratio. The width of this window depends on the geometry of the loop, the working fluid, and the applied heat input. A higher heat input narrows the window from below, as a larger liquid inventory is needed to sustain the evaporation rate without drying out.

The two limits can be estimated geometrically from the loop dimensions. The dryout limit corresponds to the minimum filling ratio required to keep the bottom pipe and the heated section of the riser fully filled with liquid at rest, ensuring the heated wall remains wetted at startup:

$$FR_{dryout} = \frac{A \cdot (L_{bottom} + L_{heated,riser})}{V_{total}} \tag{27}$$

The flooding limit corresponds to the maximum filling ratio below which the top horizontal pipe sections remain free of liquid, preserving a vapor path from the riser to the condenser:

$$FR_{flood} = 1 - \frac{A \cdot L_{top}}{V_{total}} \tag{28}$$

where A is the pipe cross-sectional area, L_i are the lengths of the respective pipe sections, and V_{total} is the total internal volume of the loop. These expressions provide conservative geometric bounds; the true operational limits may differ depending on the heat input and the dynamic redistribution of liquid during operation.

5 Research Questions

The objective of this thesis is to design, build and operate a sub-atmospheric two-phase thermosyphon loop for magnetic resonance imaging of two-phase flow. Two sets of requirements apply at once. The loop must sustain natural circulation with a two-phase flow developed enough to be worth imaging, and it must be built from materials and in a geometry that can be placed inside a scanner bore. The work therefore runs from the design and modeling of the loop, through its construction and operation, to an assessment of how

close the result stands to what an MRI measurement would require. This leads to the following research questions.

1. How can a sub-atmospheric two-phase thermosyphon be designed and modeled to achieve stable natural circulation while satisfying the constraints required for MRI-based flow visualization?

This question covers the selection of the working fluid, operating pressure, pipe diameter, heat input and charge mass, together with the material and dimensional limits imposed by the scanner. It is addressed with a steady-state model that closes the pressure balance around the loop and predicts the circulation rate and outlet quality at a given heat input, used as a design tool to find the region in which circulation is both stable and observable.

2. How does the operating behavior of the developed thermosyphon vary under different operating conditions, and to what extent can this behavior be predicted?

This question is addressed experimentally. The charge mass, the starting condition of the fluid and the heat input are varied in turn, and the circulation rate, the loop temperatures and the vapor quality derived from them are measured at each condition. These measurements are then compared against the model of the first question, to establish how far the loop can be predicted and where it departs from the prediction.

3. To what extent does the developed thermosyphon satisfy the requirements for future MRI-based visualization of two-phase flow, and what modifications are required before MRI measurements can be performed?

This question weighs what the loop can deliver against what an MRI measurement would require, in the flow conditions it can reach and hold, in the materials and geometry of the sections that pass through the bore, and in the instrumentation available to interpret an image once one is taken. Recommendations are then made for the changes the loop would need before it is taken to the scanner.

6 Design and modeling of the thermosyphon loop

Following the literature review on thermosyphon operation and flow visualization, the next step is the translation of these concepts into a physically realizable experimental system. The design of the thermosyphon loop is driven by both the physical requirements of two-phase natural circulation and the constraints imposed by operation within an MRI environment. In parallel, a set of models is developed to support the design process and to predict the operating conditions of the system.

6.1 Design Constraints and Operating Conditions

The system is designed to operate at sub-atmospheric conditions. The main design choice is the operating pressure, which is set to approximately 0.12 bar. At this pressure, the saturation temperature of water is significantly reduced compared to atmospheric conditions($T_{sat} = 50°C$), allowing boiling to occur at relatively low temperatures. This reduces the required heat input in the evaporator and minimizes the thermal loads close to the MRI scanner, which is critical for safe and stable operation.

In addition to thermal considerations, the MRI environment imposes strict constraints on the design. As mentioned in section 3.3, all components located near the scanner must be non-magnetic and preferably non-conductive to avoid distortion of the magnetic field and interference with image acquisition. Furthermore, the geometry of the setup must be compatible with the limited bore size of the scanner, and the region of interest must be positioned within the homogeneous magnetic field zone.

To satisfy these constraints, the thermosyphon loop is designed with an extended geometry, as seen in figure 9. The heating section is located at a sufficient distance from the MRI measurement region, connected through long horizontal pipe sections of 1.65 meters, set by the available lab space and the distance required

from the scanner. This configuration ensures that the heat source and any associated electrical components are positioned away from the sensitive zone, while the two-phase flow can still be observed inside the scanner.

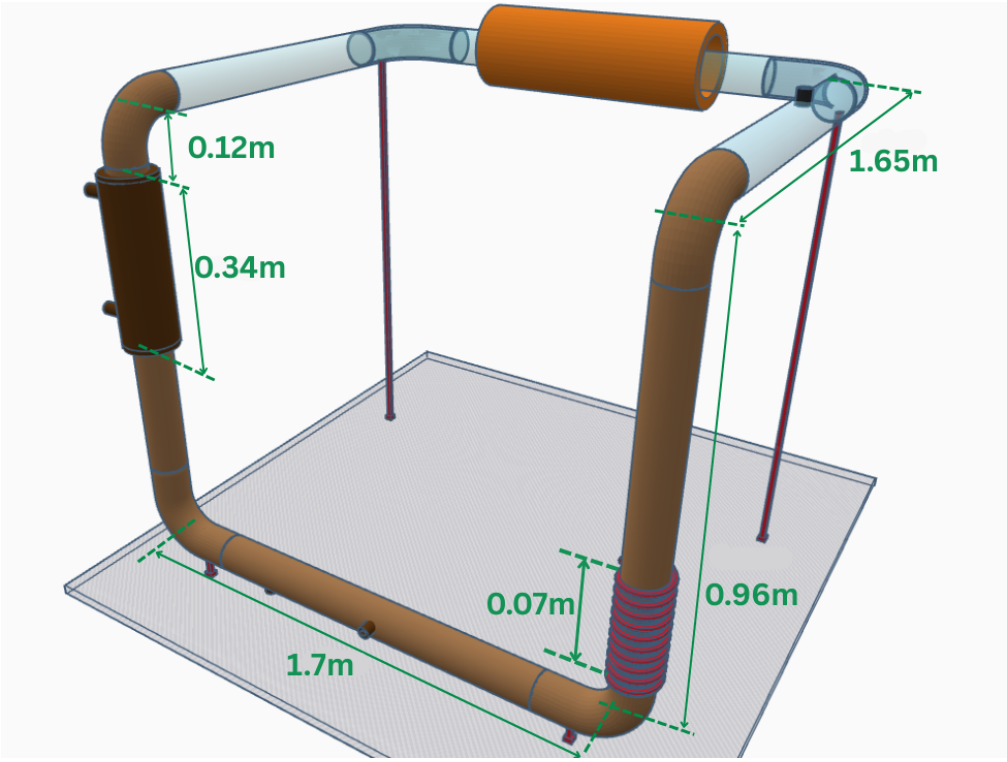


Figure 9: 3D model of experimental setup

6.2 Thermosyphon Loop MATLAB Model

The steady-state behavior of the thermosyphon loop is predicted using a lumped-parameter pressure balance model implemented in MATLAB. The model is built around the fundamental condition that governs natural circulation at steady state: the net pressure change around one complete cycle of the fluid must equal zero, as seen in the following Equation (eq. 29).

$$\sum_i \Delta P_i = 0 \tag{29}$$

where the sum runs over all pipe sections in the loop. The physical interpretation is straightforward: the hydrostatic pressure gain recovered in the liquid-filled downcomer must exactly balance all resistances, such as frictional losses in every section, the hydrostatic weight of the two-phase mixture in the riser, the accelerational pressure drop in the heated section, and localized bend losses at each corner. The loop geometry, fluid properties and heat input are all prescribed inputs; the only unknown is the circulation mass flow rate $\dot{m}$.

For a given heat input Q, the vapor quality at the exit of the heated section is linked to $\dot{m}$ through the energy balance at the evaporator:

$$x_{out} = x_{in} + \frac{Q}{\dot{m} \cdot h_{fg}} \tag{30}$$

where h_{fg} is the latent heat of vaporization and $x_{in} = 0$ since the flow enters the heated section as saturated liquid. Substituting this energy balance into the pressure balance produces a single nonlinear algebraic equation in the single unknown $\dot{m}$, which is solved numerically using the `fzero` root-finding function in MATLAB.

The overall solution procedure is illustrated in Figure 10.

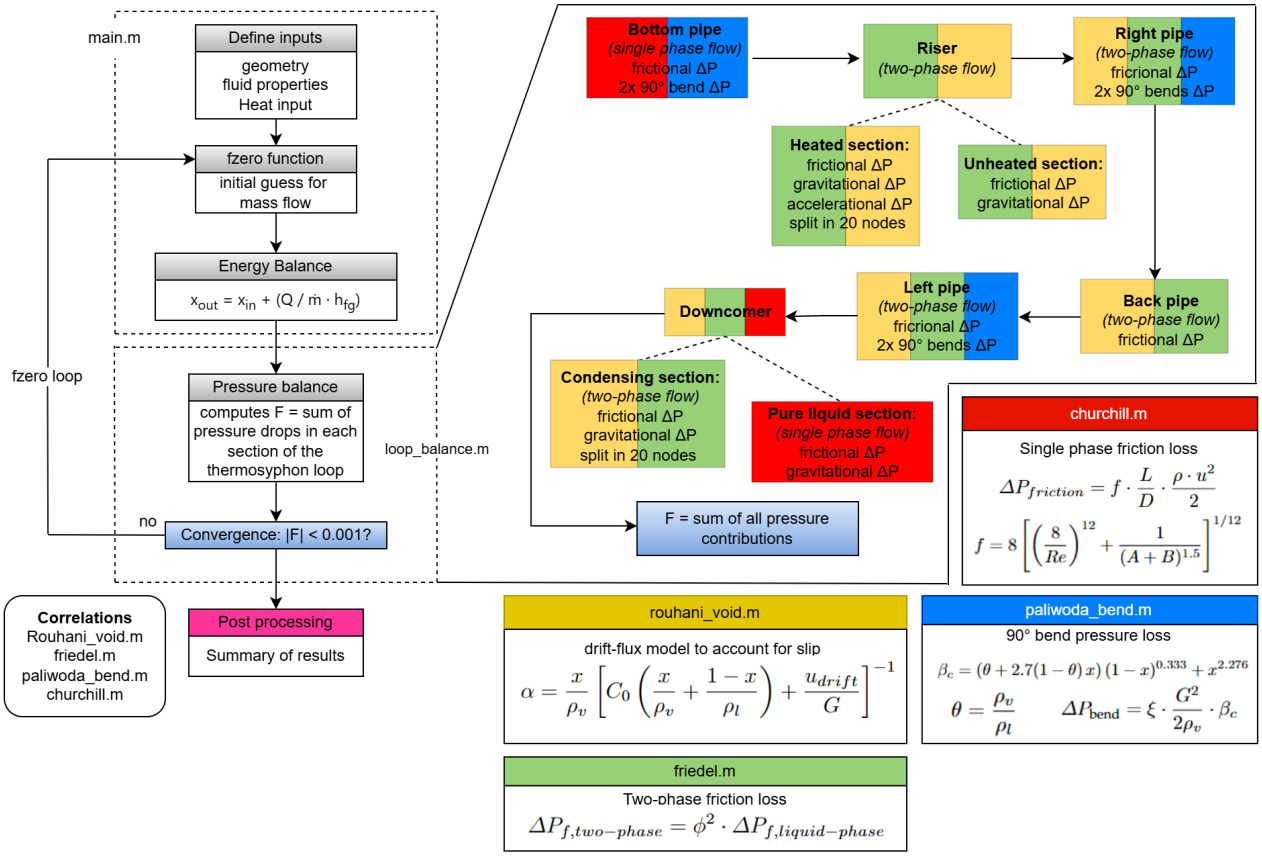


Figure 10: Flowchart of MATLAB model

The pressure balance function `loop_balance.m` evaluates the pressure contribution of each of the six pipe sections in sequence. The bottom pipe and the lower part of the downcomer carry single-phase liquid; their frictional losses are computed using the Churchill correlation [32], which provides a continuous expression for the Darcy friction factor valid for all flow regimes. The riser, horizontal top pipes, and upper condensing zone of the downcomer carry two-phase flow; their frictional losses are computed using the Friedel correlation [33], which applies a two-phase multiplier ϕ^2 to the all-liquid pressure drop. Two-phase bend losses are computed using the Paliwoda correlation [36]. In all two-phase sections, the local void fraction is evaluated using the drift-flux model of Rouhani and Axelsson [34], from which the local mixture density is obtained and used to compute hydrostatic pressure contributions.
The heated section of the riser and the condensing zone of the downcomer are both treated with a nodal approach: each is divided into $N = 20$ equally spaced nodes, and the local void fraction, mixture density, frictional and hydrostatic pressure contributions are evaluated at each node individually. This captures the variation in mixture properties along these sections without requiring a full computational fluid dynamics treatment. In the heated section, an accelerational pressure drop is additionally computed to account for the momentum change as the fluid evaporates and accelerates along the pipe, as seen in section 4.1, equation 22.

Once `fzero` has identified the mass flow rate $\dot{m}$ that satisfies the pressure balance, all other quantities of interest are computed as post-processing steps. The outlet vapor quality x_{out}, the void fraction distribution along the riser, and the required filling ratio FR are all derived from the converged solution. The dryout and flooding limits are also evaluated: the dryout limit defines the minimum FR required to keep the evaporator wall wetted, and the flooding limit defines the maximum FR above which the vapor path through the top horizontal pipes would be blocked by liquid.

6.3 Parametric Study for Diameter Selection

The internal pipe diameter D is one of the most important geometric choices in the design of the loop, since it sets both the circulation rate reached at a given heat input and the vapor quality of the flow that develops. It influences the system through two competing effects. A smaller diameter increases the mass flux $G = \dot{m}/A$ for a given circulation rate, which raises frictional pressure losses and can limit $\dot{m}$, but also produces higher void fractions and more vigorous two-phase structures. A larger diameter reduces friction losses and allows a higher mass flow rate, but at the same heat input the outlet vapor quality x_{out} is lower, potentially producing weak two-phase flow with limited MRI contrast. An intermediate diameter is therefore expected to best satisfy the requirements simultaneously.
Three criteria are used to evaluate each candidate diameter:

1. **Stable natural circulation**: the model must converge to a positive, physically meaningful $\dot{m}$ across the intended heat input range. Diameters that produce negligible or unstable circulation are excluded.

2. **Practical flow velocity and heat input**: the resulting velocities should remain high enough to sustain circulation, but low enough to avoid excessive pressure losses and overly rapid flow structures. In addition, the heat input required to develop a two-phase flow must remain within the capacity of the electrical heater.

3. **Access to multiple flow regimes**: the diameter should place the measurement section in more than one of the horizontal flow regimes described in Section 2.2, such as stratified, wavy, plug, slug and annular flow, across the available heat input range, rather than locking the loop into a single regime, so that a transition between regimes can be observed. The regimes reachable by each candidate diameter are evaluated in Section 6.4.

A parametric sweep is conducted over inner diameters of $D \in \{10, 13, 19.8, 30\}$ mm. The first three are standard commercially available copper pipe sizes from the selected supplier, and the fourth is included to show where the trend with increasing diameter leads. Each diameter is applied to the whole circuit, so the loop is taken to have that same bore in the riser, the horizontal pipes and the downcomer. All other dimensions, including the loop height, the length of each section, the position of the heater along the riser and the position of the condenser along the downcomer, are held at their design values, as seen in figure 9. The model is solved over a heat input range of $Q = 10$ to 1100 W, a reasonable upper limit for electrical heating of a loop of this size. The resulting outlet vapor quality and mean liquid velocity are shown in Figure 11, and the void fraction and circulation rate in Figure 12.
All four diameters sustain natural circulation, but over very different spans of heat input. The 10 mm and 13 mm curves end at roughly 130 W and 430 W respectively. This is not a failure of the solver but the point at which the outlet quality reaches the upper limit of validity of the two-phase correlations, above which no pressure balance is returned. A diameter whose curve ends early can only be operated over the part of the heater range below that point.
The four diameters separate clearly on how quickly quality develops with heat input. The 10 mm bore passes from single-phase liquid to almost complete evaporation within about 130 W, which is an eighth of the installed heater capacity. The whole range of two-phase conditions is compressed into a narrow band of heat input, so a small change in power produces a large change in the state of the flow and the operating point is correspondingly difficult to hold. The 13 mm bore spreads the same range over roughly 430 W, around three times wider. The 19.8 mm bore reaches only about half the quality range at the top of the sweep, so its upper half cannot be reached at all, and the 30 mm bore stays below $x_{out} \approx 0.03$ throughout, never leaving the nearly single-phase state.
The circulation rate behaves inversely, as the pressure balance requires. At low heat input the loop runs fast and wet, and as the quality rises the circulation rate falls by two to three orders of magnitude. The 30 mm

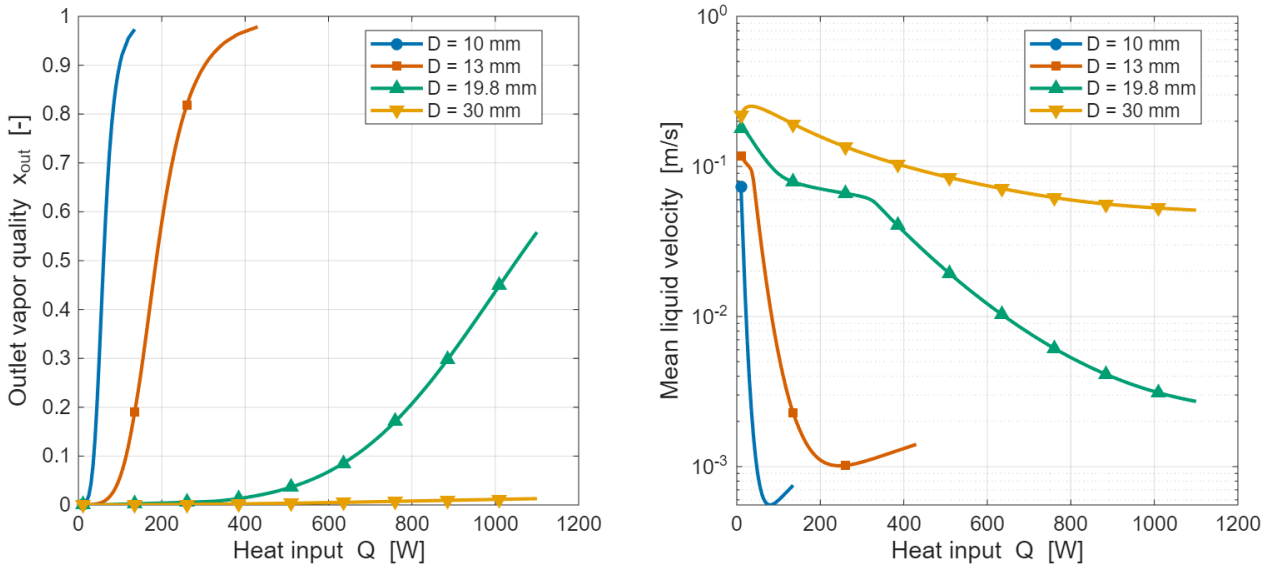


Figure 11: Outlet vapor quality and mean liquid velocity.

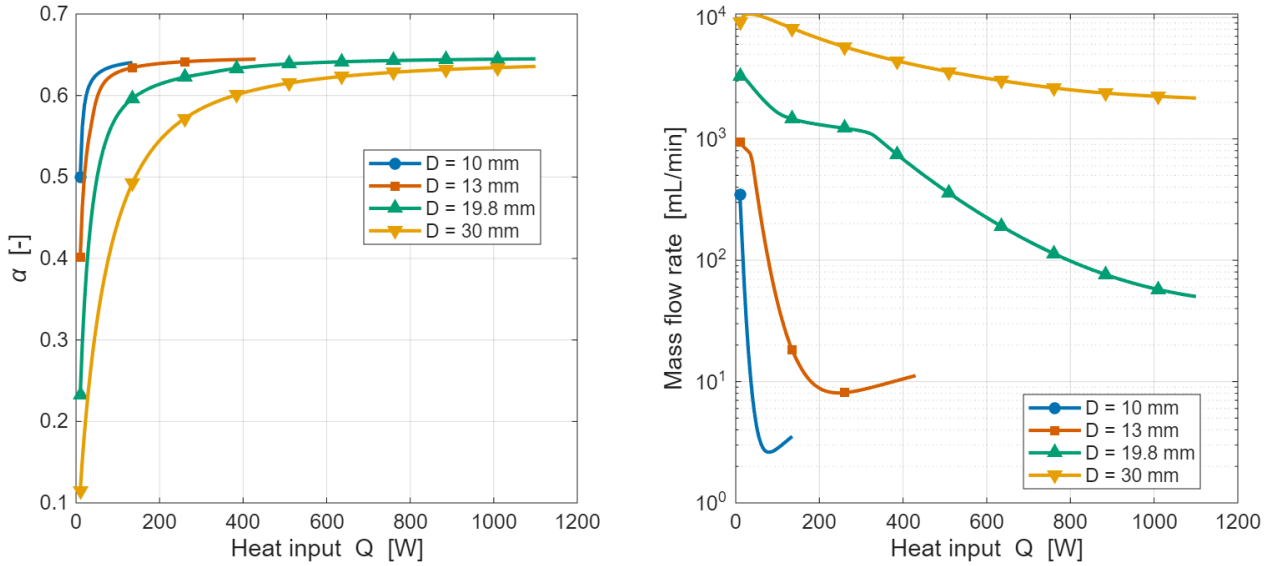


Figure 12: Void fraction and circulation rate.

bore never leaves the fast branch, holding above 2000 mL/min at full power with a mean liquid velocity of roughly 0.05 m/s. At the other extreme the 10 mm bore falls below 3 mL/min, with a liquid velocity under 10^{-3} m/s, where the flow is slow enough that the assumption of a steady circulation becomes questionable. The 13 mm bore passes through a minimum of about 8 mL/min near 260 W and recovers slightly beyond it. The void fraction panel of Figure 12 shows every diameter rising steeply at low heat input and then flattening near $\alpha = 0.65$. This ceiling is a property of the drift-flux correlation rather than of the flow. With the distribution parameter $C_0 = 1.54$ adopted for the low mass flux branch in Section 4.1, Equation (19) approaches $1/C_0 = 0.649$ as the quality approaches one, so the model cannot report a void fraction above that value at any quality. The true void fraction at the densities of this system continues to rise towards unity. The predicted void fraction is therefore conservative at high quality, and the riser mixture it implies is heavier than the real one, which in turn makes the predicted circulation rate conservative as well. The panel is retained because the approach to the ceiling still distinguishes the diameters at low heat input, but the flat portion carries no information about the flow.

The 30 mm bore is excluded on the second criterion. Reaching a developed two-phase flow would require a heat input several times the installed capacity, and within the available power the loop simply circulates

subcooled liquid with a trace of vapor. It is carried through the figures only to show the direction of the trend, and is not considered further.

6.4 Flow Regime Prediction

The primary measurement section of the thermosyphon loop is the back pipe, which passes horizontally through the MRI bore. The two-phase flow pattern present in this section determines what can be visualized and measured, and therefore directly influences the choice of pipe diameter. Unlike the riser, where quality increases continuously along the heated length, the back pipe carries two-phase flow at approximately a fixed quality x_{out} if heat loss along the top pipes is neglected.

The third criterion is checked by evaluating, for every diameter and across the full heat input range, the coordinates of each of the four maps introduced in Section 2.3, and placing the resulting trajectory on that map.

Four diameters are examined, each treated as a uniform bore, so the loop is taken to have the same internal diameter in the riser, the back pipe, the downcomer and every other section rather than changing from one to the next. Three of them, 10, 13 and 19.8 mm, are the commercially available copper sizes considered in the parametric study above. The fourth, 17 mm, is the bore of the polymer pipe available for the sections that pass through the scanner.

Three of the four maps are used as published, by digitizing their boundaries from the original figures. The fourth, the mass flux against quality map, has no published form for water and steam at sub-atmospheric pressure and is constructed here from its governing equations.

Across the whole operating range the superficial gas velocity exceeds the superficial liquid velocity by three to five orders of magnitude. This follows from the density ratio of roughly 12,200 at 0.12 bar, since a small vapor mass fraction fills almost the entire cross-section and must travel quickly to carry its share of the mass flow.

6.4.1 Superficial velocity map

The superficial velocities of Equation (5) were evaluated for each diameter and plotted on the map of Mandhane et al. [10] in Figure 13. Across the four diameters the gas velocity spans roughly 0.55 to 130 ft/s, which sits comfortably inside the range the map covers, while the liquid velocity spans 1.6×10^{-4} to 0.80 ft/s and therefore reaches two orders of magnitude below the lowest liquid velocity the map was drawn for. Slightly more than half the computed operating points lie below the bottom edge of the published map and cannot be placed on it at all.

The points that do fall on the map describe a consistent picture. Each trajectory begins at its lowest heat input near the boundary between the stratified and slug regions, then descends steeply to the right as the heat input rises, since increasing the heat input raises the vapor velocity while reducing the liquid velocity. The 10, 13 and 17 mm trajectories pass into the stratified region almost immediately and remain there, crossing into the wavy region only at the highest heat inputs. The 19.8 mm trajectory sits roughly an order of magnitude higher in liquid velocity and tracks along the slug and wavy boundaries rather than descending below them, which reflects the much larger circulation rate that diameter sustains at a given heat input.

6.4.2 Superficial Reynolds number map

Recasting the same operating points in the superficial Reynolds numbers of Equation (6) gives Figure 14, placed on the map of Ghajar and Tang [12]. The gas Reynolds number spans 26 to 3000 and the liquid Reynolds number 1 to 8700, against a map whose boundaries are drawn from $Re_{SG} = 457$ upward. The operating points therefore occupy the lower left of the plane, below and to the left of most of the boundary structure.

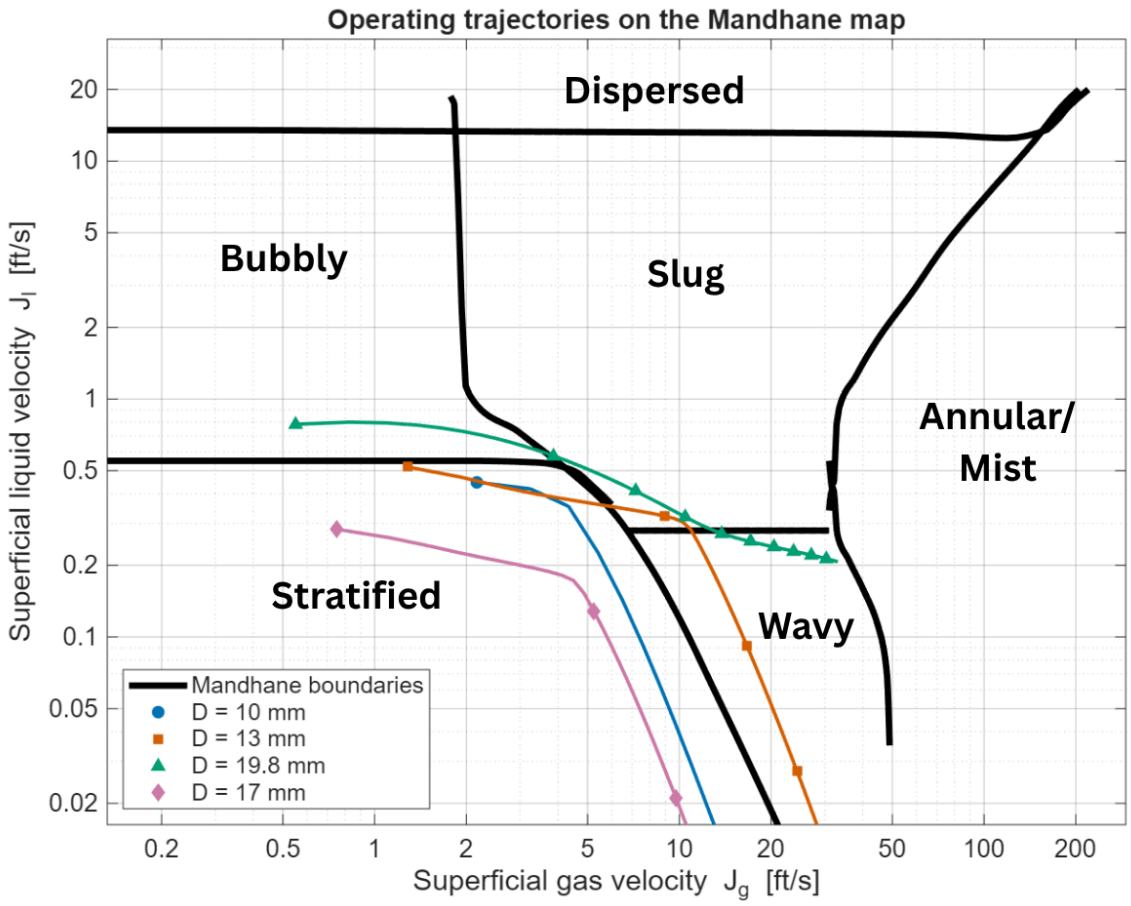


Figure 13: Operating trajectories on the Mandhane map.

This is itself informative. The boundary that separates the stratified region from the slug region above it runs horizontally at $Re_{SL} \approx 2800$ over the range of gas Reynolds numbers the loop reaches. The 10, 13 and 17 mm trajectories lie below that line throughout, and all three fall steeply as the heat input rises, ending one to two orders of magnitude below it. Only the 19.8 mm trajectory begins above the line, at $Re_{SL} \approx 8700$, and it crosses downward through it within the first part of the heat input range. The conclusion matches the superficial velocity map, which is expected since the two differ only by the diameter and viscosity grouping that converts one coordinate system into the other.

6.4.3 Reconstruction of a mass flux against quality map

The two maps above are empirical and were drawn for air and water. The mass flux against quality family is the only one surveyed whose boundaries can be evaluated for an arbitrary fluid, so a map of that type is constructed here for saturated water and steam at the operating pressure of the loop. The construction follows the explicit form given by Thome and Cioncolini [14].

The transitions are evaluated using the dimensionless areas and liquid level of an equivalent stratified flow geometry. The cross-sectional void fraction α is obtained from the Rouhani-Axelsson drift-flux model introduced in Section 4.1 (Eqs. 19 and 20), reproduced here in compact form:

$$\alpha = \frac{x}{\rho_v}\left[C_0\left(\frac{x}{\rho_v} + \frac{1-x}{\rho_l}\right) + \frac{u_{drift}}{G}\right]^{-1} \tag{31}$$

The stratified angle θ_{strat}, defined as the angle subtended at the center of the cross-section by the upper dry perimeter, is then obtained from the explicit approximation of Biberg[14]:

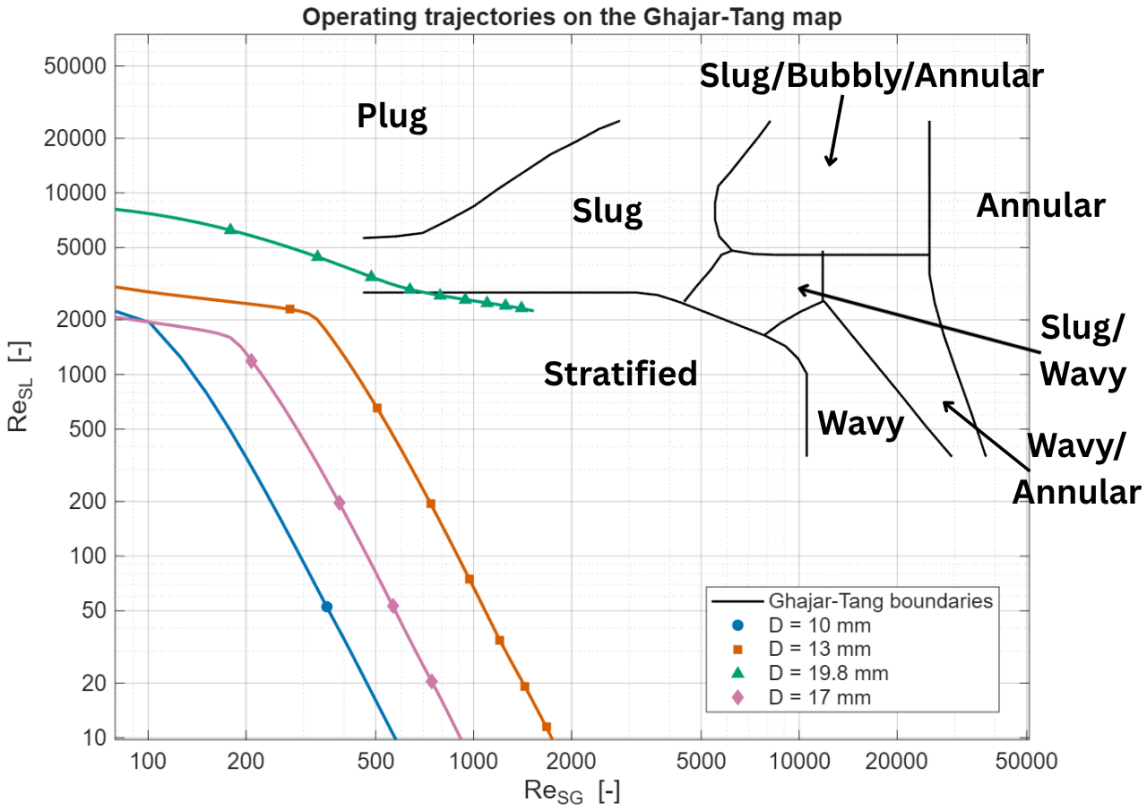


Figure 14: Operating trajectories on the Ghajar-Tang map.

$$\theta_{strat} = 2\pi - 2\Big\{ \pi(1-\alpha) + (3\pi/2)^{1/3}\Big[1 - 2(1-\alpha) + (1-\alpha)^{1/3} - \alpha^{1/3}\Big] \\ - \tfrac{1}{200}(1-\alpha)\alpha\big[1-2(1-\alpha)\big]\big[1+4((1-\alpha)^2+\alpha^2)\big]\Big\} \tag{32}$$

From the void fraction and stratified angle, the dimensionless liquid and vapor cross-sectional areas A_{ld}, A_{gd} and the dimensionless liquid height h_{ld} follow directly:

$$A_{ld} = (1-\alpha)\frac{\pi}{4}, \qquad A_{gd} = \alpha\,\frac{\pi}{4}, \qquad h_{ld} = \tfrac{1}{2}\Big[1 - \cos\big((2\pi - \theta_{strat})/2\big)\Big] \tag{33}$$

Three transition curves are constructed on the (G, x) plane. The stratified to stratified-wavy transition is given by:

$$G_{strat} = \left[\frac{(226.3)^2\, A_{ld}\, A_{gd}^2\, \rho_v(\rho_l - \rho_v)\,\mu_l\, g}{x^2(1-x)\,\pi^3}\right]^{1/3} \tag{34}$$

The stratified-wavy to intermittent or annular transition, in its adiabatic form, is:

$$G_{wavy} = \left\{\frac{16\, A_{gd}^3\, g\, d\, \rho_l\, \rho_v}{x^2\pi^2\sqrt{1-(2h_{ld}-1)^2}}\left[\frac{\pi^2}{25\, h_{ld}^2\,(We/Fr)_l} + 1\right] + 50\right\}^{1/2} \tag{35}$$

where $(We/Fr)_l = \rho_l\, g\, d^2/\sigma$ is the liquid Weber-to-Froude number ratio. The vertical intermittent to annular boundary is set by a fixed-value criterion on the turbulent-turbulent Martinelli parameter, $X_{tt} = 0.34$, where

$$X_{tt} = \left(\frac{1-x}{x}\right)^{0.875}\left(\frac{\rho_v}{\rho_l}\right)^{0.5}\left(\frac{\mu_l}{\mu_v}\right)^{0.125} \tag{36}$$

Inverting this criterion for x gives the closed-form expression for the transition quality:

$$x_{IA} = \left[0.2914 \left(\frac{\rho_v}{\rho_l} \right)^{-1/1.75} \left(\frac{\mu_l}{\mu_v} \right)^{-1/7} + 1 \right]^{-1} \tag{37}$$

The curves were evaluated for saturated water and steam at $p = 0.12$ bar, with $T_{sat} \approx 49.4\,^{\circ}\text{C}$ and IAPWS-IF97 properties, for a tube diameter of 13 mm, and are shown in Figure 15 with the operating trajectories overlaid.

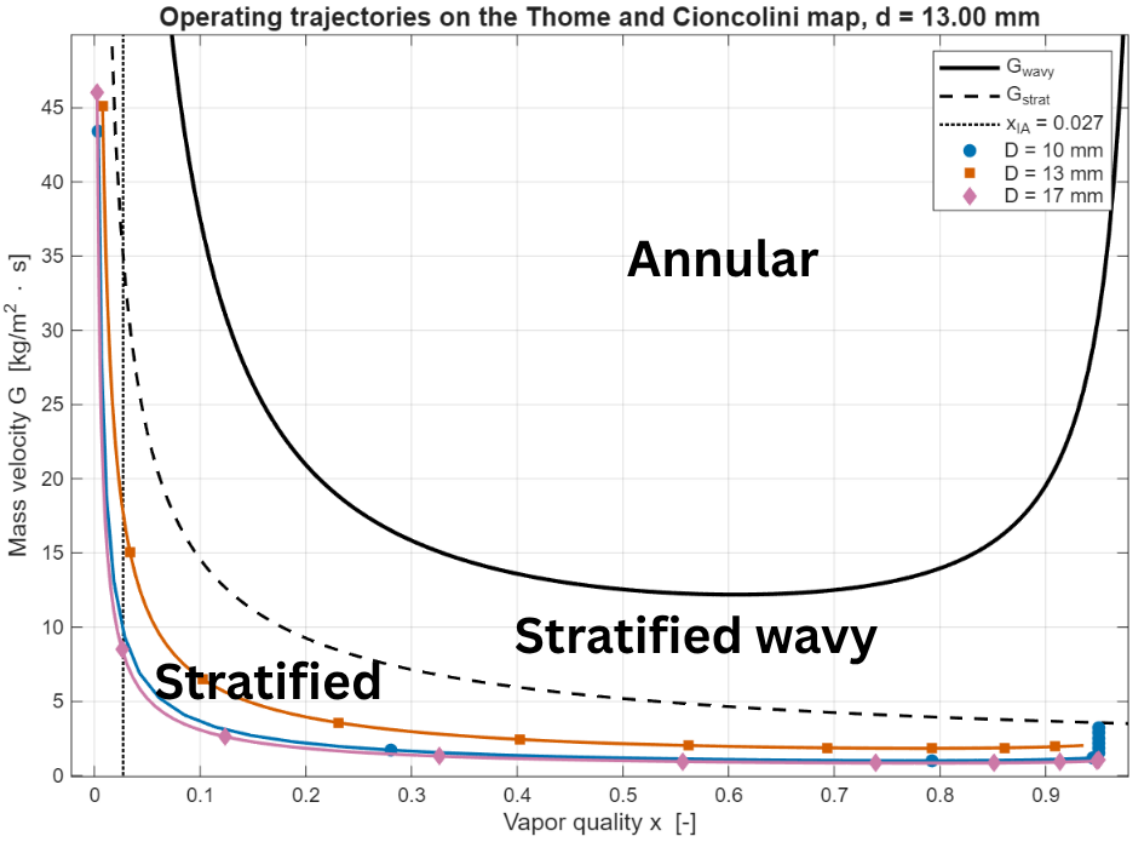


Figure 15: Operating trajectories on the Thome and Cioncolini map.

Only G_{wavy} depends on the tube diameter at all. Equations (34) and (37) contain no diameter term, so those two boundaries are identical for every diameter. The diameter enters G_{wavy} both explicitly and through the $(We/Fr)_l$ group, and evaluating Equation (35) across the range examined shifts the curve by a few percent near the middle of the quality range and by about a quarter of a decade at its extremes. That shift is far smaller than the distance between the operating points and the boundary, and testing each diameter separately placed the flow in the same regime in every case, so a single map drawn for 13 mm is used for all of them.

Each trajectory enters the map at its lowest heat input with a high mass flux and a quality near zero, then falls steeply as quality rises, and runs along the bottom of the plane at a mass flux of 2 to 5 kilograms per square meter per second for the remainder of the range. Every trajectory lies below G_{strat} across the whole range of quality, which is the fully stratified region. The boundary is never crossed, not even at the lowest qualities where the mass flux is at its highest and the trajectory comes closest to it.

Two features of the map itself are notable, independent of where the loop sits on it. First, the intermittent to annular transition collapses to a very low quality, $x_{IA} \approx 0.027$, against roughly 0.36 for R-22 at 5 °C. This follows directly from the density ratio entering Equation (37) through $(\rho_v/\rho_l)^{-1/1.75}$. Because the density ratio at 0.12 bar is two orders of magnitude smaller than that of a typical refrigerant, the quality required to reach the same volumetric vapor occupancy is correspondingly smaller. Evaluating Equation (31) at x_{IA} for both fluids confirms that the transition occurs at essentially the same void fraction in both cases, so the

Martinelli criterion is in effect a void fraction criterion written in mass-quality coordinates, and the apparent collapse of x_{IA} reflects only the change of coordinates.

Second, the transition mass velocities are roughly an order of magnitude lower than for refrigerants. The minimum of G_{wavy} is approximately 12 kg/m^{2}s for water and steam at 0.12 bar, compared with 121 kg/m^{2}s for R-22 at the same diameter [14]. This again traces back to the density ratio, since G_{wavy} scales with $\sqrt{\rho_l \rho_v}$ and the very small vapor density pulls the whole curve down. The combination of a low boundary and a low operating mass flux is what places the loop in the stratified region rather than above it.

6.4.4 Mechanistic map

The last comparison is against the mechanistic map of Taitel and Dukler [13], shown in Figure 16. Reading this map requires more care than the other three. Its four boundaries share a single horizontal axis, the Martinelli parameter X, but each is read against a different vertical group: curves A and B against the Froude-type group F, curve C against the Jeffreys-type group K, and curve D against the turbulence group T. A single operating point therefore appears at three different heights depending on which transition is being tested, which makes a combined plot difficult to read. The comparison is instead split into three panels sharing the horizontal axis, each carrying the boundary it belongs with.

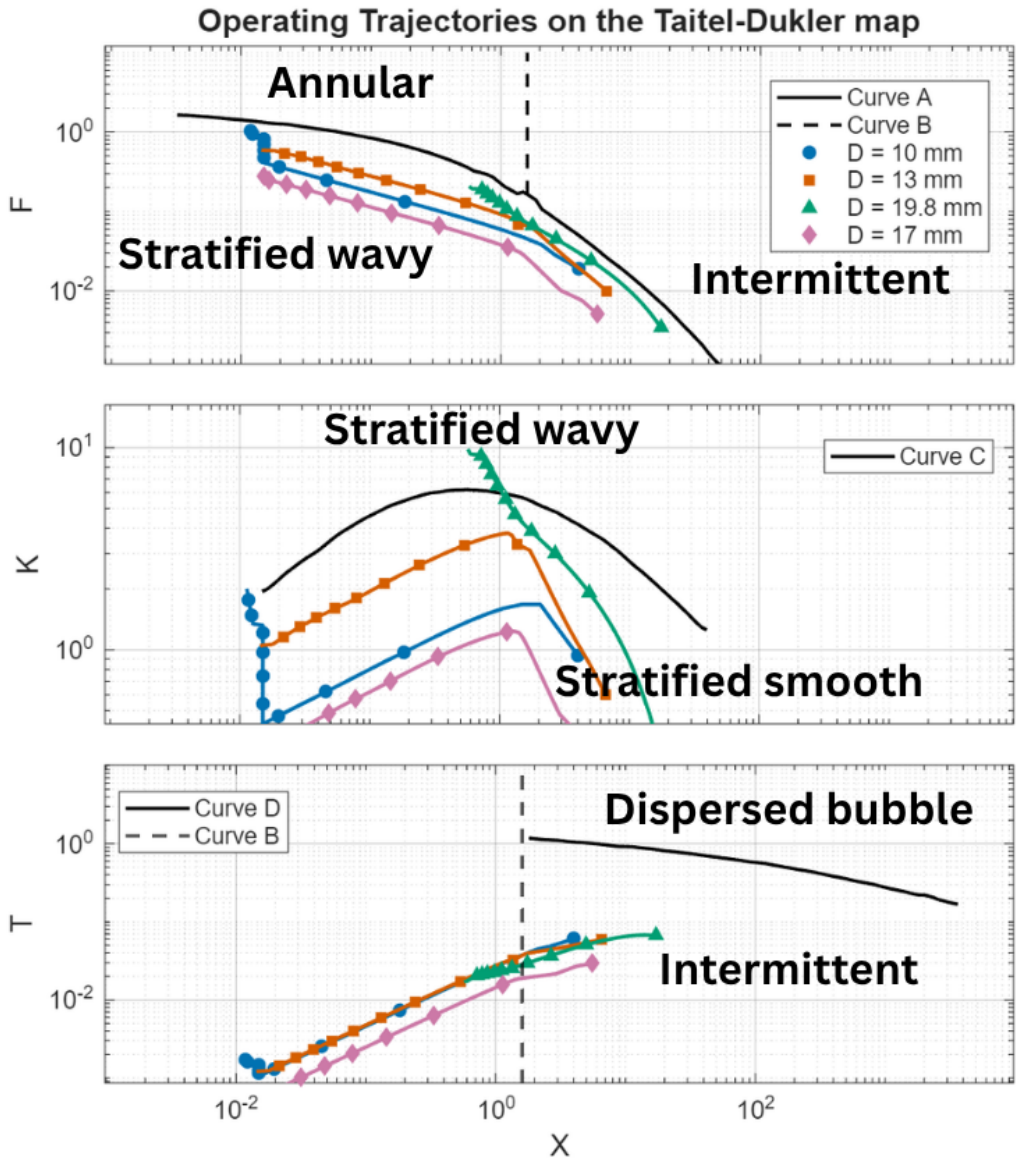


Figure 16: Operating trajectories on the Taitel-Dukler map.

In the first panel every trajectory lies below curve A across the whole heat input range, which places the flow on the stratified side of the stratified to non-stratified transition for all four diameters. Curve B is therefore never reached, since it only separates regimes above curve A. The second panel resolves the stratified flow into its smooth and wavy forms. The 10, 13 and 17 mm trajectories remain below curve C throughout, which is stratified smooth flow, while the 19.8 mm trajectory rises above curve C beyond a heat input of about 335 W and is predicted to be stratified wavy from there on. In the third panel every trajectory lies far below curve D, so dispersed bubble flow is excluded everywhere, which is consistent with the stratified prediction of the first panel.

6.4.5 Common conclusion and limitations

All four maps place the measurement section in the stratified region across the entire heat input range, for every diameter examined. The two empirical maps, the mechanistic map of Taitel and Dukler and the constructed one agree on this despite being built on different coordinate systems, different fluids and, in the case of Taitel and Dukler, a mechanistic rather than an empirical derivation. Within that agreement the only distinction between diameters is how far into the stratified region they sit: the 19.8 mm case lies closest to the boundaries and is the only one predicted to reach stratified-wavy flow, while the smaller diameters sit deeper in the smooth stratified region as the heat input rises.

This outcome does not satisfy the third criterion as it was stated. None of the diameters examined gives access to several distinct flow regimes across the available heat input range, because the very large density ratio at sub-atmospheric pressure produces a liquid velocity too small for any of the transitions out of stratified flow to be crossed. The criterion is therefore not a discriminator between the diameters, and the diameter selection rests on the first two criteria.

The confidence that can be placed in this prediction is limited, and the limitation is the same for all four maps. None of them was developed for water and steam at sub-atmospheric pressure. The operating conditions of the present system, $P = 0.12$ bar and $\rho_l/\rho_v \approx 12{,}200$, fall outside the validated range of every available correlation. The Thome and Cioncolini framework was developed and validated on a refrigerant database, covering R-22, R-134a, R-123, R-410A, R-402A, R-404A, R-502, R-407C, R-507A and ammonia, with density ratios between roughly 30 and 250, and the authors state explicitly that the map is not validated for air-water or water-steam systems. The empirical maps carry the same restriction from the opposite direction, having been built on air-water data at atmospheric pressure. A further practical limitation is visible in Figures 13 and 14: a substantial part of each trajectory falls outside the area the published maps cover, so for those points the regime is being inferred by extrapolating boundaries beyond where they were ever drawn.

These maps are therefore used as qualitative guidance rather than as quantitative prediction. What supports the conclusion is not the accuracy of any one of them but the fact that four independent frameworks, applied to the same operating points, place the flow in the same region. An experimental measurement of the flow pattern in this section would provide data in a pressure regime that none of the existing correlations covers, which is one of the motivations for the work.

6.5 Final Diameter Selection

A bore of $D = 13$ mm is selected for the copper sections of the loop, supplied as 15 mm outer diameter tube with a 1 mm wall. Among the diameters examined it is the only one that develops the full range of two-phase conditions, from single-phase liquid to near-complete evaporation, across a span of heat input wide enough to be set and held with the available heater. The 10 mm bore covers the same range but compresses it into roughly 130 W, which leaves the operating point too sensitive to the power setting and drives the circulation rate below a few milliliters per minute. The 19.8 mm bore is controllable but cannot be driven past about half the quality range within the installed capacity, and the 30 mm bore produces almost no vapor at all.

The horizontal sections that pass through the scanner cannot be made of copper, since the magnetic resonance environment requires non-magnetic and non-conducting materials as described in Section 3.3. These are made from PVC-U instead, in the largest bore available in the required pressure rating, 20 mm outer diameter with a 1.5 mm wall giving a 17 mm bore. Selecting the polymer sections wider than the copper has a deliberate effect on the measurement: for a given circulation rate the mass flux in the back pipe is lower than in the riser, which places the measurement section further into the stratified region and produces a thicker and more stable liquid layer than the riser bore would give. Section 6.4 shows that the 17 mm bore sits in the same regime as the 13 mm bore across the whole heat input range, so the change of diameter between the heated section and the measurement section does not change the flow pattern being observed.
The flow regime analysis also shows that the third criterion cannot be satisfied by any diameter. The measurement section is predicted to remain stratified throughout, and no choice of bore within the range considered gives access to a second regime. That outcome is a property of the operating pressure rather than of the geometry: at a liquid-to-vapor density ratio of roughly 12,200 the liquid velocity is too small for any of the transitions out of stratified flow to be crossed. The consequence for the measurement is not that there is nothing to observe. Stratified flow has a well defined liquid layer with a free surface, and film thickness, interface shape and the velocity profile within the liquid are all quantities MRI can resolve, none of which have been reported for water and steam at sub-atmospheric pressure.
No single heat input is adopted as a design point. The loop is intended to be operated by sweeping the heat input, so that the circulation rate and the flow pattern can be followed as they change, and it must therefore reach a steady state at every power setting within the range rather than at one chosen condition. The condenser and heater sizing calculations that follow are carried out against the limits of that range: the condenser must reject the largest heat input the loop will be driven at while holding the saturation temperature, and the heated section must still sustain boiling at the smallest.

6.6 Condenser Sizing

The condenser must reject the heat supplied in the evaporator while fitting within the available downcomer length. Since the loop is intended to be operated by sweeping the heat input rather than at one nominated condition, the condenser is sized at the top of the range over which the selected 13 mm bore develops two-phase flow. Section 6.4 shows that the outlet vapor quality reaches $x_{out} = 1$ at around 430W, so a design heat load of $Q = 500$ W is adopted, leaving some margin above it.
A tube-in-tube counterflow heat exchanger is used, in which the condensing working fluid flows downward inside the inner copper tube while cooling water flows upward through the surrounding annulus, as shown in Figure 17.

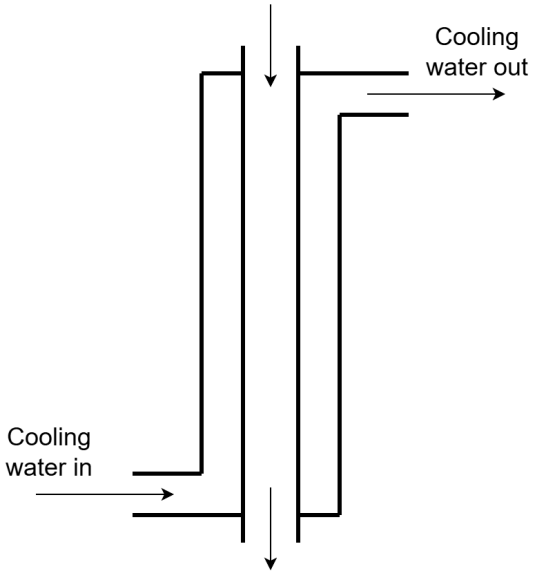


Figure 17: Tube-in-tube counterflow heat exchanger

To reach the cooling water, the heat must cross three barriers one after the other: the liquid film that the condensing vapor leaves on the inside wall, the copper wall itself, and the slow-moving layer of water against the outside of the tube. Each barrier has a thermal resistance, and because they act in series the resistances simply add:

$$\frac{1}{U} = \frac{D_o/D_i}{\alpha_{cond}} + R_{wall} + \frac{1}{\alpha_{ann}} \tag{38}$$

where α_{cond} and α_{ann} are the heat transfer coefficients on the condensing and cooling water sides. All three terms are referred to the outer surface of the inner tube, which is why the condensing term carries the area ratio D_o/D_i: condensation happens on the inner surface, while the exchanger area is counted on the outer one.

The wall term follows from conduction through a cylinder:

$$R_{wall} = \frac{D_o}{2k_{wall}} \ln\left(\frac{D_o}{D_i}\right) \tag{39}$$

with k_{wall} the thermal conductivity of copper. For a 1 mm copper wall it is under one percent of the total and is kept only for completeness.

As the vapor condenses it forms a thin liquid film that runs down the inside wall, and all of the heat released must pass through that film. The thinner the film, the more easily the heat escapes. Nusselt theory gives the coefficient for such a film on a vertical surface, on the assumption that the film flows laminarly [7]:

$$\alpha_{cond} = 0.943 \left(\frac{\lambda_l^3 \cdot \rho_l \cdot g \cdot h_{fg} \cdot (\rho_l - \rho_v)}{L_{cond} \cdot \mu_l \cdot (T_{sat} - T_w)}\right)^{1/4} \tag{40}$$

where T_w is the wall temperature on the condensing side, λ_l and μ_l are the conductivity and viscosity of the liquid, and h_{fg} is the latent heat. The film collects more condensate as it runs down, so the coefficient depends on the condenser length L_{cond}, which is the quantity being solved for. Neither L_{cond} nor T_w is known in advance, so both are guessed and the calculation is repeated iteratively until they stop changing.

Whether the film stays laminar is checked through its own Reynolds number, formed from the condensate flow spread around the tube perimeter:

$$Re_{film} = \frac{4\,\dot{m}_{cond}}{\pi D_i \mu_l}, \qquad \dot{m}_{cond} = \frac{Q}{h_{fg}} \tag{41}$$

At $Q = 500$ W only 2.1×10^{-4} kg/s of vapor condenses, which gives $Re_{film} = 37$, far below the value of about 1800 at which such a film becomes turbulent [37].
On the cooling water side the annulus Reynolds number, formed on the hydraulic diameter D_h, is 2227, so the water flow is laminar. The annulus is short enough that the water has not settled into a steady temperature profile by the time it leaves, so the Hausen relation for a thermally developing laminar flow is used [37]:

$$Nu_{ann} = 3.66 + \frac{0.0668\, Gz}{1 + 0.04\, Gz^{2/3}}, \qquad Gz = \frac{D_h}{L_{cond}}\, Re\, Pr \tag{42}$$

$$\alpha_{ann} = \frac{Nu_{ann}\, k}{D_h} \tag{43}$$

The Graetz number Gz also contains L_{cond}, so this coefficient is updated in the same iteration. At the design condition it gives $Nu_{ann} = 11.1$.
Once U is known the length follows from the standard heat exchanger relation:

$$Q = U\, A\, \Delta T_{lm}, \qquad A = \pi D_o L_{cond} \qquad \Rightarrow \qquad L_{cond} = \frac{Q}{U \pi D_o \Delta T_{lm}} \tag{44}$$

The driving force ΔT_{lm} is the temperature difference between the two streams averaged along the exchanger. The average is logarithmic rather than arithmetic because that difference does not vary linearly with position:

$$\Delta T_{lm} = \frac{\Delta T_1 - \Delta T_2}{\ln(\Delta T_1 / \Delta T_2)} \tag{45}$$

where ΔT_1 and ΔT_2 are the differences at the two ends. The condensing fluid stays at its saturation temperature throughout, so only the cooling water changes temperature along the length.

The condenser geometry consists of two concentric copper tubes, summarized in Table 1.

Table 1: Condenser tube geometry

Tube	Role	Outer dia.	Inner dia.	Wall thickness
Inner tube	Condensing fluid	15 mm	13 mm	1 mm
Outer jacket	Cooling water	22 mm	19.8 mm	1.1 mm

The annulus between the outer surface of the inner tube and the inner surface of the outer jacket provides the flow channel for the cooling water, with a hydraulic diameter of:

$$D_h = D_{i,outer} - D_{o,inner} = 19.8 - 15 = 4.8 \text{ mm} \tag{46}$$

A low cooling water flow rate of 250 L/h, equal to $\dot{m}_{cool} = 0.069$ kg/s, is selected together with an inlet temperature of 15°C. The water then moves through the annulus at 0.53 m/s and warms by only 1.7 K, so ΔT_{lm} is 34.1 K and is set almost entirely by the gap between the saturation temperature and the water inlet. For this geometry the model predicts an overall heat transfer coefficient of $U = 1158$ W/m^{2}K and a required length of:

$$L_{cond} = 0.27 \text{ m} \tag{47}$$

The three resistances are very unevenly matched. The annulus accounts for 85% of the total, the condensing film for 15% and the copper wall for less than 1%, with $\alpha_{cond} = 8941$ W/m^{2}K against $\alpha_{ann} = 1366$ W/m^{2}K. The condenser is therefore limited almost entirely by the cooling water side.
To provide margin against uncertainty in the heat transfer correlations and against operation above the design load, a safety factor of 1.3 is applied to the required length. The design length is therefore:

$$L_{cond,design} = 0.35 \text{ m} \tag{48}$$

which occupies 36% of the 0.96 m downcomer, leaving the remainder for the liquid column below the condenser and for the two-phase section above it.

It should be noted that the required condenser length is not determined by pipe diameter alone. It is influenced by several coupled parameters, including the imposed heat duty Q, the cooling water inlet temperature and mass flow rate, the annulus gap, the tube wall thickness, the material conductivity, and the resulting flow regime on the cooling side.

6.7 Heater Sizing

The evaporator must supply a controlled heat input to the working fluid in the riser across the whole range over which the loop is to be operated. Since the loop is swept rather than run at a fixed condition, the heater capacity sets the upper end of the accessible range, and the parametric study of Section 6.3 was carried out against an upper bound of 1100 W. For the selected 13 mm bore the outlet vapor quality reaches the top of the range the steady-state model covers at around 430 W, so the heat input needed to span the useful operating range is well inside that bound. The heater is specified to 1100 W, which gives a factor of roughly two over the requirement. The surplus covers heat losses from the heated section and the unheated riser, allows the upper end of the range to be reached without the heater running at its limit, and leaves room to operate the loop at a larger bore should the design be revisited.

Two heater configurations were considered: external heating tape wrapped around the riser, and embedded cartridge heaters within an aluminum heating block. After consultation with the laboratory technical staff, the cartridge heater configuration was selected. Cartridge heaters embedded in a thermally conductive block provide significantly lower heat losses than external heating tape, because the heat is delivered through a tightly contained metallic conduction path rather than through a thin layer adhered to the outer pipe surface and exposed to ambient air. The aluminum block also acts as a thermal homogenizer, distributing the heat from the discrete cartridges into a uniform radial heat flux on the copper riser.

At the conceptual design stage the heater is therefore represented as a set of cylindrical cartridge heaters embedded symmetrically around the copper riser within a machined cylindrical aluminum block, redistributing the discrete heat sources into a uniform radial heat flux on the tube wall. The heat input is varied continuously by regulating the AC supply voltage delivered to the cartridges rather than by switching individual elements on and off, which preserves a symmetric heat distribution at all power levels and gives smooth control instead of discrete steps. The specific number and rating of the cartridges, the final block geometry, and the associated power-control and safety hardware are fixed during construction and are described together with the as-built heater assembly in Section B.2.

6.8 Expansion Vessel Concept

As the loop heats from ambient to the operating condition of $T_{sat} = 50\,^{\circ}\text{C}$, the fluid inside it needs more room. Thermal expansion of the liquid accounts for about 1% and is negligible here. What matters is evaporation: at $P = 0.12$ bar the density ratio is $\rho_l/\rho_v \approx 12{,}200$, so every gram of liquid that boils occupies a vapor volume more than four orders of magnitude larger. In a sealed loop with nowhere for that volume to go, the pressure would climb with heat input and carry the operating point away from the design condition.

An expansion vessel provides the room, but it must also do something harder, which is to fix the saturation pressure. For a pure substance T_{sat} and P_{sat} lie on a single curve, so fixing one fixes the other. The design pressure of 0.12 bar corresponds to $T_{sat} = 50\,^{\circ}\text{C}$, and that condition has to hold across the whole range of heat inputs if measurements taken at different operating points are to be compared.

Which vessel configuration achieves this follows from how the fluid mass held in the loop depends on heat

input and saturation temperature. At any steady operating point the pressure balance model gives the mass flow rate $\dot{m}$ that satisfies the loop hydraulics at a given Q and T_{sat}, from which the outlet quality $x_{out} = Q/(\dot{m}\, h_{fg})$ follows. The total fluid mass M_{loop} is then obtained by integrating the local density over every pipe section: liquid sections contribute $\rho_l \cdot A \cdot L$, while two-phase sections contribute the mixture density $[\alpha\, \rho_v + (1-\alpha)\, \rho_l] \cdot A \cdot L$, with the void fraction α evaluated from the Rouhani-Axelsson model. Repeating this across a grid of $(Q,\, T_{sat})$ values produces a family of constant-T_{sat} curves of M_{loop} against Q. These are the isotherms of the system, and because T_{sat} and P_{sat} are locked together, they are equivalently isobars.
At $Q = 0$ every section is liquid, so all isotherms begin at nearly the same mass. As Q rises vapor forms, the mixture lightens and M_{loop} falls, faster on the hotter isotherms because ρ_v grows with temperature. The mass charged at filling, $M_{initial}$, is fixed once the loop is sealed and can be drawn as a horizontal line across this map. The system operates where that line meets an isotherm, and the isotherm it meets sets T_{sat} and therefore P_{sat}. The upper panel of Figure 18 shows the loop without a vessel: the isotherms are steep, the horizontal line crosses a different one at every heat input, and the pressure floats freely with Q.

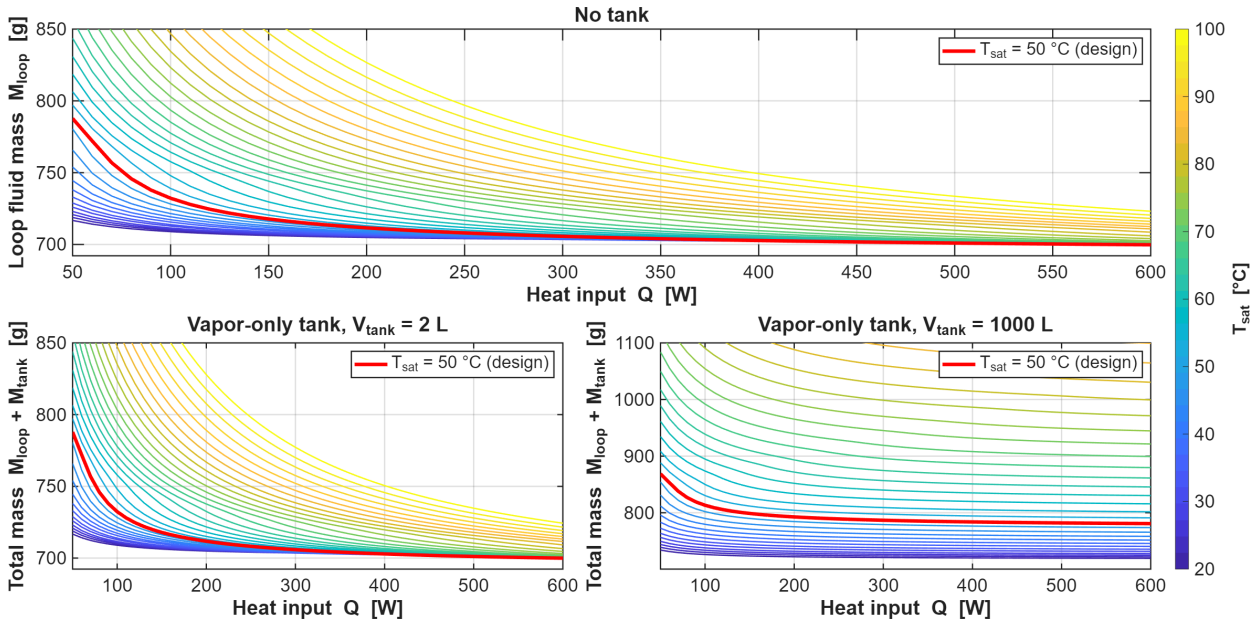


Figure 18: Total fluid mass against heat input for different tank configurations.

Adding a vapor-filled vessel of volume V_{tank} at the top of the loop raises the system mass by $V_{tank} \cdot \rho_v(T_{sat})$, which shifts the isotherms upward and, for a large enough volume, flattens them. The lower panels of Figure 18 show two sizes. At 2 L almost nothing changes, because $\rho_v \approx 0.08$ kg/m^3 at these conditions means the vessel holds only 0.16 g of vapor against a loop mass of well over a kilogram. Only at 1000 L do the isotherms approach flat, and a vessel of that size is not practical.
Size is not the only objection. Flat isotherms mean that one particular value of $M_{initial}$ is needed to sit on the $50\,^{\circ}$C isotherm, and being a few grams off places the system on a different isotherm at a different pressure. Filling to that precision is not achievable, given the uncertainty in the modeled geometry, in the fluid properties and in the charging procedure itself. A liquid-filled vessel fails in the same way and for the same reason. A single-phase vessel of either kind therefore offers no robust pressure control, it only converts an uncontrolled pressure drift into an uncontrolled pressure offset.
A vessel holding saturated liquid and saturated vapor at the same time works on a different principle, which is why it succeeds at a size where the single-phase vessels fail. It does not flatten the isotherms at all. The vapor-liquid interface inside it imposes $P = P_{sat}(T_{tank})$ on the entire connected system, so the pressure is set by the temperature of the vessel rather than by how much mass was charged. As heat input increases and more vapor forms in the loop, liquid is displaced into the vessel and the level there rises. For as long as both phases remain present the interface survives and the pressure stays pinned. Any fill mass that keeps

the vessel two-phase across the operating range therefore gives the correct pressure, which is a far more forgiving requirement than hitting an exact target.

The vessel temperature is held at $50\,^{\circ}\mathrm{C}$ by circulating thermostatted water through a stainless-steel coil inside it, fed from a dedicated bath. With T_{tank} fixed, the interface enforces $P = P_{sat}(50\,^{\circ}\mathrm{C}) = 0.12$ bar throughout the experiment, independently of the heat applied to the loop.

The volume the vessel needs follows from how much liquid the loop expels when it starts boiling. Three sections stay liquid at any heat input: the bottom pipe, the riser below the heater, and the downcomer below the condenser. Together these hold $0.311\,\mathrm{L}$ of the $1.427\,\mathrm{L}$ internal volume of the loop, leaving $1.116\,\mathrm{L}$ that can become a two-phase mixture. If those sections reach a void fraction α, the liquid they no longer hold has been displaced into the vessel,

$$\Delta V_{liq} = \alpha \cdot V_{2\phi}. \tag{49}$$

The Rouhani-Axelsson model with $C_0 = 1.54$ cannot return a void fraction above $1/C_0 = 0.649$, which gives $\Delta V_{liq} = 0.725\,\mathrm{L}$.

The vessel only fixes the pressure while both phases are present inside it, so the liquid level must stay within a working band rather than reaching either end. Allowing the level to move between 20% and 80% of the vessel volume,

$$V_{tank} \geq \frac{\Delta V_{liq}}{f_{max} - f_{min}} = \frac{0.725}{0.6} = 1.21\,\mathrm{L}. \tag{50}$$

This value inherits whatever uncertainty the void fraction correlation carries at sub-atmospheric pressure, so it is worth checking the limiting case. If the two-phase sections vaporized completely, at $\alpha = 1$, the displaced liquid would be the full $1.116\,\mathrm{L}$ and the requirement would rise to $1.86\,\mathrm{L}$. The vessel should therefore be sized above that bound, at roughly $2\,\mathrm{L}$, so that the two-phase condition holds across the whole operating range regardless of how the void fraction behaves.

6.9 Building the experimental setup

Figure 19 shows the complete piping and instrumentation diagram (P&ID) of the assembled loop, and Figure 20 shows the physical implementation.

The loop resembles the design developed in Section 6, adapted where necessary to the components that were actually available and installed. The loop is a single closed path driven by natural convection. The working fluid boils in the heater block (Appendix B.2), the resulting vapor rises through the copper riser into the MRI-compatible polymer sections, the tube-in-tube condenser (Appendix B.3) cools the vapor back to liquid, and the copper downcomer carries the liquid back to the heater. System pressure and fluid inventory are set and held constant by the expansion vessel (Appendix B.4), while five Pt100 RTDs (TS01–TS05) and one electromagnetic flow meter record the loop's thermal and hydraulic state throughout operation (Appendix B.5).

Each subsystem is described in full in Appendix B, covering how it was built, which materials were chosen and why, and photographs of the assembled hardware. The operating procedure, from how the loop is evacuated, charged, brought to the target saturation state, and operated, is described in Chapter 7.

Tables 2 to 7 list every component used to build the loop, with its specification, quantity, supplier and part number.

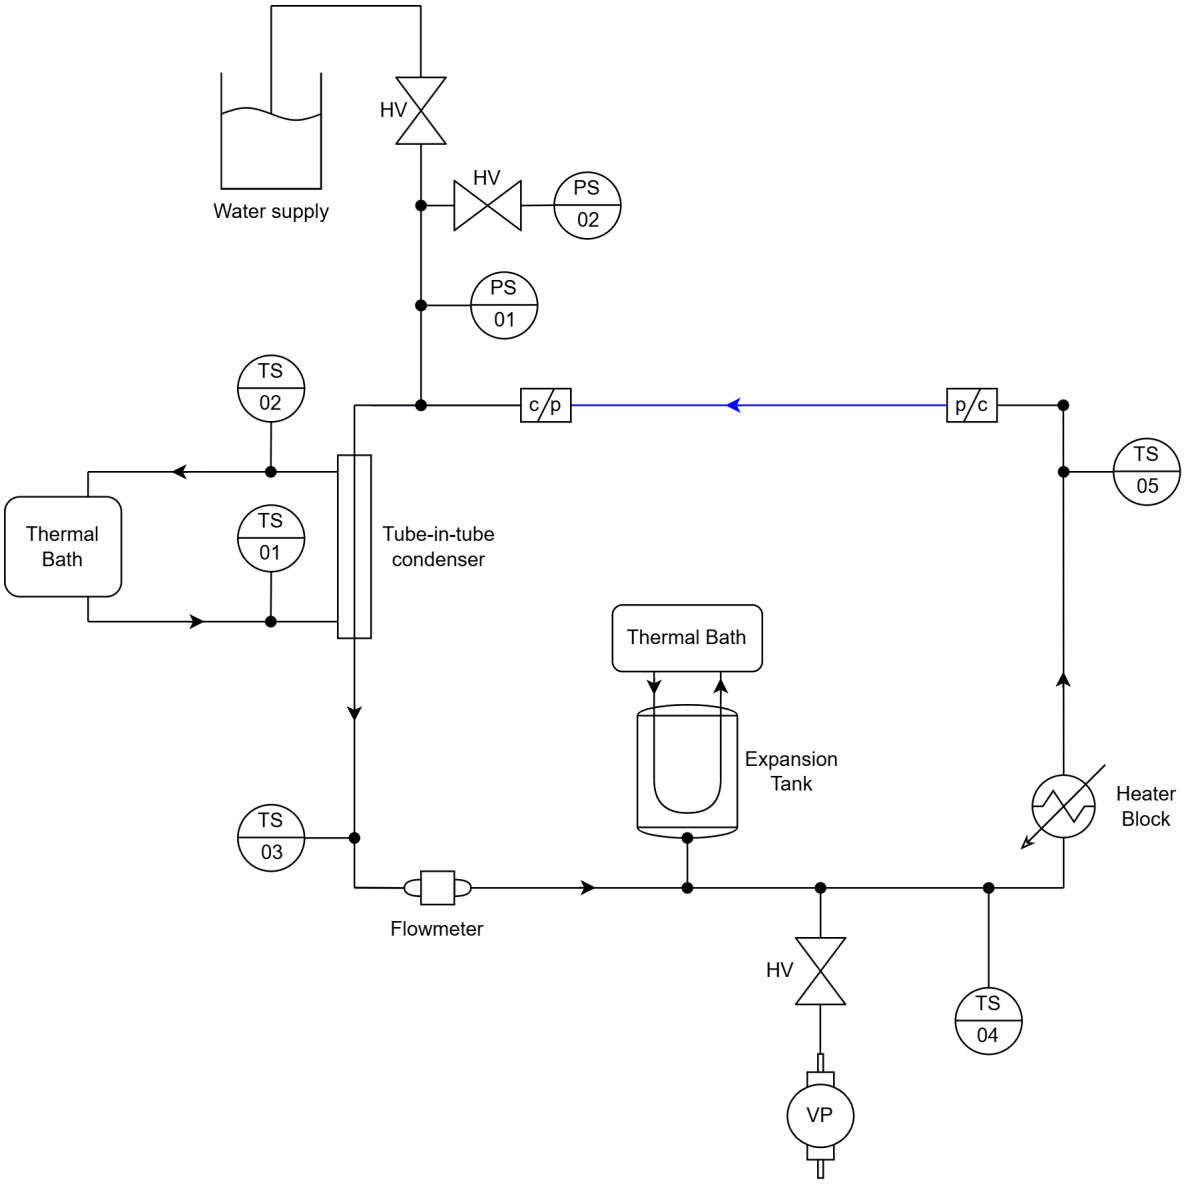


Figure 19: P&ID of the complete system

Table 2: Loop piping and fitting components

Component	**Specification**	**Qty**	**Supplier / Part No.**
Copper pipe, soft	15 mm OD / 13 mm ID, ∼6 m	1	Wildkamp, art. 890458
PVC-U pipe, opaque	16 mm OD / 12.4 mm ID, ∼5 m	1	Eriks, EC010882-0003
PVC-U pipe, transparent	20 mm OD / 17 mm ID, ∼1 m	1	Eriks, EC010882-0005
PVC-U 3-part union coupling	16 mm (PVC–PVC and Cu–PVC transitions)	6	Eriks, Serie 3.65

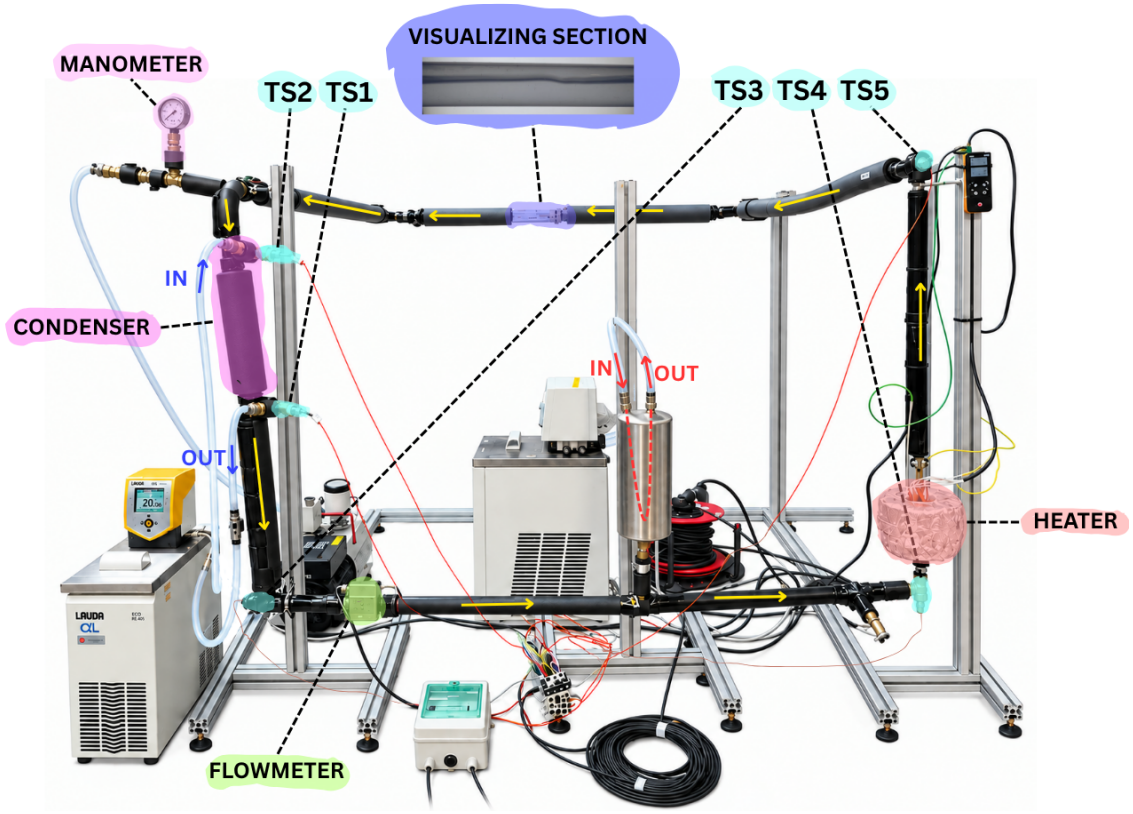


Figure 20: Experimental setup

Table 3: Heating system components

Component	Specification	Qty	Supplier / Part No.
AC cartridge heater	12.5 mm dia., AC, 180W	6	TC Direct, art. 921-112
Aluminum heater block	Machined, central 15 mm bore for riser	1	In-house
Rock wool	Heater insulation	-	Lab stock
Variac	Percentage controlled	-	Lab stock
Armaflex pipe insulation	For 15 mm pipe, ∼5 m	1	Wildkamp, art. 15583768

Table 4: Cooling and condenser system components

Component	Specification	Qty	Supplier / Part No.
Copper pipe, soft (condenser jacket)	22 mm OD / 19.8 mm ID, ∼0.5 m	1	Wildkamp, art. 890460
Brass reducing T-piece	22×15×15 mm	2	Wildkamp, art. 848708
Brass ball valve	1/4-inch, vacuum port + fill/drain	1	Wildkamp, art. 653766
Water bath	Cooling-water supply, fixed at 250 L/h	1	Lab stock

Table 5: Pressure control and expansion vessel components

Component	Specification	Qty	Supplier / Part No.
Stainless steel tube (vessel shell)	101.6 mm OD, 2 mm wall, RVS 304	1	RVS-Products
End cap (vessel)	Dia. 101 mm × 2 mm, RVS 304	2	RVS-Products
Copper pipe, soft	15 mm OD / 13 mm ID, ∼0.5 m	1	Wildkamp, art. 890458
Brass reducing T-piece (vessel branch)	15×12×15 mm	1	Wildkamp, art. 848702
Brass T-piece (pressure branch + fill/drain branch)	15 mm × 1/4-inch × 15 mm	2	Wildkamp, art. 849301
Brass ball valve	1/4-inch, vacuum port + fill/drain	1	Wildkamp, art. 653766
Brass coupling (compression × 1/4-inch M thread)	15 mm × 1/4-inch M	2	Wildkamp, art. 15381840
Straight adapter, SS 316Ti	1/4-inch BSPT M × 1/4-inch BSPT M	1	Eriks, RC0222-0071
Manometer (PS01)	Bourdon tube, 0–1 bar	1	Eriks, type 1401
Vacuum gauge (PS02)	0-2666 Pa range	1	Testo 552i
Brass Hose pillars	1/4-inch M	5	QUARKZMAN
Vacuum pump (VP)	0.001 bar	1	Lab stock
Water bath	Vessel temperature control	1	Lab stock

Table 6: Instrumentation components

Component	Specification	Qty	Supplier / Part No.
Pt100 RTD sensor (TS01–TS05)	Standard budget RTD	5	TC Direct, art. 514-103
Thermowell (RTD immersion tube)	For 3 mm sensor, 1/2-inch male thread	5	TC Direct, art. 784-750
Brass T-piece (RTD branch)	15 mm×15 mm×1/2" mm	5	Wildkamp, art. 849310
Electromagnetic flow meter	Loop circulation flow rate	1	Krohne AF-E 400
Flow meter connector	1/4" compression × 15 mm	2	Wildkamp, art. 15381840
Flow meter adapter	BSPP F × BSPP F	2	Eriks
MAX31865 RTD-to-digital converter	One per sensor, shared SPI bus	5	Lab stock
Arduino Uno	Data acquisition and logging	1	Lab stock

Table 7: Mechanical structure and mounting components

Component	Specification	Qty	Supplier / Part No.
Pipe clamp	–	6	Lab stock
Aluminum profiles	–	4	Lab stock

7 Experimental Campaign

This chapter describes how the experimental campaign is carried out. It begins with the calibration of the temperature measurement chain, which establishes the accuracy with which the readings used throughout the rest of this work can be trusted. The operating procedure is then described: how the loop is evacuated, charged and brought up to the target saturation state, and how it is held there during a run. The acquisition software developed to read the instrumentation in real time, derive the quantities of interest, and record the complete history of each experiment is described next. The results of the campaign are then presented in Chapter 8.

7.1 Sensor Calibration

Temperature is the primary measured quantity in the thermosyphon loop. At the operating pressure of $0.12\,\mathrm{bar}$ the saturation temperature of water is approximately $50\,^{\circ}\mathrm{C}$, and distinguishing saturated, sub-cooled and superheated conditions requires the temperature to be resolved to better than $0.1\,\mathrm{K}$. The uncalibrated sensor chain does not achieve this. As set out in Section B.5, the Class B Pt100 tolerance and the readout electronics together contribute an error of order $0.5\,\mathrm{K}$ to $1\,\mathrm{K}$ at $50\,^{\circ}\mathrm{C}$. Both contributions are systematic and repeatable, so both can be removed by comparison against a reference standard.

That reference is a four-wire Pt100 read by a Keysight DAQ970A on a multiplexer module, configured for four-wire resistance measurement with offset compensation enabled and an integration time of ten power line cycles. Its specified accuracy is

$$u_{\mathrm{ref}} = \pm\,(0.03 + 0.0005\,|T|)\,\mathrm{K}, \tag{51}$$

or $0.055\,\mathrm{K}$ at $50\,^{\circ}\mathrm{C}$, roughly ten times better than the sensors being calibrated. Both instruments report resistance rather than temperature, and both are converted on the host computer using the same implementation of Equation (67), so that no offset arises from differing conversion routines. A single acquisition program reads both, so every record contains a reference reading and the corresponding sensor readings taken within a few seconds of one another.

7.1.1 Calibration Procedure

The calibration follows a comparison method. The sensors and the reference are immersed together in a thermostatic bath at equal depth in a common holder, so that any temperature difference between their positions is as small as possible. The bath is brought to a series of setpoints spanning the operating range of each sensor group, and at each setpoint all channels are recorded at the same time once the bath has settled. Settling is judged by comparing each channel against its own value one look-back window earlier, using the parameters of Table 8.

The permitted drift had to be relaxed from $0.02\,\mathrm{K}$ to $0.05\,\mathrm{K}$ for the loop sensors. Below $30\,^{\circ}\mathrm{C}$ the tighter value was met routinely, but at $50\,^{\circ}\mathrm{C}$ and above the channels kept passing and failing in turn over intervals of several minutes. The bath is not the cause. At the $55\,^{\circ}\mathrm{C}$ setpoint the reference channel varied by $0.030\,\mathrm{K}$ over $40\,\mathrm{min}$ while the Arduino channels varied by $0.066\,\mathrm{K}$ to $0.109\,\mathrm{K}$. If the bath were moving, every channel

Table 8: Steady state monitoring parameters.

Parameter	Cooling water sensors	Loop sensors
Setpoint range	10 °C to 30 °C	30 °C to 70 °C
Sampling interval	10 s	10 s
Look-back window	300 s	600 s
Permitted drift	0.02 K	0.05 K

would move with it and the difference between a sensor and the reference would stay constant. That difference kept changing, so the movement comes from the Arduino measurement chain. Waiting longer does not help, because the readings are not settling towards a value but drifting slowly around one. Settling was therefore judged on the reference channel alone, and the movement of the Arduino channels was averaged over instead.

At each setpoint $N_s = 40$ samples are recorded at $10\,\mathrm{s}$ intervals and reduced to a single value by taking the mean of each channel. Averaging normally improves precision, but only when each reading carries new information. Because the Arduino readings drift over several minutes, two readings taken ten seconds apart are close to the same reading twice, and forty of them are worth far fewer than forty independent measurements. This independence was measured with the autocorrelation function, which compares each reading with the one k samples later and so shows how long the readings stay similar to one another,

$$\rho_k = \frac{\sum_{i=1}^{N_s-k}(x_i - \bar{x})(x_{i+k} - \bar{x})}{\sum_{i=1}^{N_s}(x_i - \bar{x})^2}. \tag{52}$$

Summing it gives τ, the number of samples that must pass before two readings are effectively unrelated, and dividing the record length by τ gives the number of independent measurements the record actually contains,

$$\tau = 1 + 2\sum_{k\geq 1}\rho_k, \qquad N_{\mathrm{eff}} = \frac{N_s}{\tau}. \tag{53}$$

The results are given in Appendix A. A record of forty samples on the Arduino channels is worth only two to five independent measurements, and in practice its mean was found to shift by up to $0.23\,\mathrm{K}$ depending on where in the slow drift the recording began. That shift is as likely to be positive as negative, so it cannot be subtracted out, but it does add scatter to the calibration points and therefore to the uncertainty of the fit.

7.1.2 Fitting and uncertainty

Setpoints are split into two groups. Calibration setpoints determine the coefficients, while verification setpoints are left out of the fit and used afterwards to check the correction at temperatures it was not built from. This guards against a correction that matches its own points closely but performs poorly between them.

For each sensor i a polynomial of degree n is fitted such that

$$T_{\mathrm{corr},i} = \sum_{k=0}^{n} c_k\, T_{\mathrm{raw},i}^{k} \approx T_{\mathrm{ref}}. \tag{54}$$

At each setpoint j the correction leaves a residual, the amount by which the corrected sensor still disagrees with the reference,

$$r_j = T_{\mathrm{corr},i,j} - T_{\mathrm{ref},j}, \qquad r_{\max} = \max_j |r_j|, \tag{55}$$

so $r_{\max}$ is the worst disagreement anywhere in the range. The typical disagreement is described by the

standard error of the regression,

$$S_r = \sqrt{\frac{\sum_j r_j^2}{N - p}}, \tag{56}$$

where N is the number of calibration setpoints and $p = n + 1$ the number of fitted coefficients. A correction is accepted when $r_{\max} < 0.1\,\mathrm{K}$ at both the calibration and the verification setpoints.

S_r lumps together every error that changes from one setpoint to the next, including converter drift, variation in the connection resistance, quantization, the short records discussed above and any non-uniformity of the bath, and it is not broken down further. Doubling it widens the interval so that the true value should fall inside it in roughly nineteen cases out of twenty, and combining that with the accuracy of the reference gives the uncertainty quoted on every corrected reading in this work,

$$U_{95} = \sqrt{(2S_r)^2 + u_{\mathrm{ref}}^2}. \tag{57}$$

S_r measures how well a corrected sensor agrees with the reference, not how close either is to the true temperature. The reference carries its own error, so no sensor can end up more accurate than u_{ref} however well it is calibrated.

A first degree correction, $T_{\mathrm{corr}} = c_1 T_{\mathrm{raw}} + c_0$, was fitted for each sensor over its expected operating range: $10\,^{\circ}\mathrm{C}$ to $30\,^{\circ}\mathrm{C}$ for the cooling water sensors TS01 and TS02, and $30\,^{\circ}\mathrm{C}$ to $70\,^{\circ}\mathrm{C}$ for the loop sensors TS03 to TS05. A second degree term was tried for every sensor and rejected. Where it reduced the residual at the calibration points it increased it at the verification points, which is the signature of a fit following the scatter rather than a real curvature. The calibration reduces the measurement uncertainty from the $0.55\,\mathrm{K}$ Class B tolerance at $50\,^{\circ}\mathrm{C}$ to between $0.16\,\mathrm{K}$ and $0.24\,\mathrm{K}$ for the loop sensors, and from $0.40\,\mathrm{K}$ to $0.16\,\mathrm{K}$ or better for the cooling water sensors. The full set of setpoints, the degree comparison, and the per-sensor coefficients and residuals are given in Appendix A.

7.1.3 Limitations

Each correction is valid only within the range it was fitted over, so applying the cooling water corrections above $30\,^{\circ}\mathrm{C}$, or the loop corrections below $30\,^{\circ}\mathrm{C}$, would be extrapolation. The acquisition software flags any reading that falls outside its range.

Forty samples per setpoint is worth as few as two independent measurements on the worst channels. This is the largest contribution to S_r for the loop sensors and the reason the acceptance criterion was not met for all of them.

Only one reference sensor was used, so the temperature variation across the bath was never measured. It is assumed small, on the grounds that the differences between sensors are well described by a per-board offset and gain. Two references at opposite ends of the cluster would settle it.

The verification set was small, two setpoints for the loop sensors and one for the cooling water sensors, so a single unrepresentative record is enough to fail the criterion, as happened for TS05.

Finally, the calibration removes only errors that stay the same afterwards. Any change to the cable routing or to the screw terminals at the converter boards alters the contact and lead resistance and invalidates it. The wiring was fixed in place afterwards, and each sensor should be re-checked against the reference whenever the loop is reassembled.

7.2 Operating Procedure

Each experiment begins with the loop fully drained. It is evacuated through the vacuum port using the rotary vane pump, with the pressure monitored on the Testo vacuum gauge (PS02). Pumping continues until the gauge reading stops falling, which marks the balance between pumping speed and the residual outgassing and leakage of the system rather than any fixed target. In a representative run the pressure fell below $100\,\mathrm{Pa}$

within about $35\,\mathrm{min}$ and settled near $34\,\mathrm{Pa}$ after around $90\,\mathrm{min}$, with little further improvement. The base pressure reached is recorded for each run. The valves to the pump and to the gauge are then closed, isolating the evacuated loop.

The loop is charged with water brought to the boil in an external vessel. Boiling removes the air dissolved in the water, since its solubility decreases as the water is heated [38], and at the boiling point the dissolved air content approaches zero. Without this step the dissolved air would come out of solution once the water entered the evacuated loop and would remain there as a non-condensable gas. Before the charging valve is opened, the hose between the vessel and the loop is filled completely with boiling water; any air left in the hose would otherwise be drawn in the moment the valve opens. The mass admitted is determined by weighing the external vessel before and after charging.

The charging valve is opened only slightly and admits water slowly. Water at $100\,^{\circ}\mathrm{C}$ entering a vessel at a few tens of pascal is far above its saturation temperature and flashes on contact, part of it evaporating and cooling the remainder. With the valve barely open this proceeds steadily; opened further, the flashing becomes violent enough to shake the loop. Once the intended mass has been admitted the valve is closed and the loop is left to cool from its post-charging state.

As the loop temperature falls toward the target, the heater is switched on at around $55\,^{\circ}\mathrm{C}$, so that boiling is established close to the intended saturation temperature of $50\,^{\circ}\mathrm{C}$ rather than after the loop has cooled well below it. The expansion vessel bath is set at the same time to a temperature in the region of $50\,^{\circ}\mathrm{C}$, chosen between $40\,^{\circ}\mathrm{C}$ and $65\,^{\circ}\mathrm{C}$ according to where the saturation temperature lies at that moment, so that the vessel begins holding the system near the target rather than allowing it to drift. The condenser is brought into operation at the same time.

Three quantities are then used to hold the loop at the target saturation condition: the temperature of the expansion vessel bath, the inlet temperature of the condenser cooling water, and the cooling water flow rate. None is fixed in advance. They are adjusted during each experiment in response to the observed behavior of the loop, and the combinations that proved effective differed between runs.

7.3 Data Acquisition and Monitoring Software

A dedicated monitoring application was developed to acquire, display and record the state of the loop throughout an experiment. It is started before any operation begins, while the loop is still drained and at rest, and runs continuously until the experiment ends, so that evacuation, charging, heat-up and every subsequent operating point are captured in a single time-stamped record. A day of experiments therefore produces one file containing the complete history of the loop, with no gaps between phases.

The application reads the five RTDs and the flow meter through the Arduino described in Section B.5. At each cycle the raw resistances are converted to temperature and corrected using the calibration of Section 7.1, and each reading is displayed together with its expanded uncertainty and plotted live against time. The acquisition interval is set by the sensor chain, and with eight conversions averaged on each of five converters, a complete cycle takes approximately $3\,\mathrm{s}$. Figure 21 shows a screenshot of the interface.

The saturation condition is obtained from TS05, located immediately downstream of the heated section where the fluid is a saturated two-phase mixture. Its reading is taken as the saturation temperature T_{sat}, and the corresponding saturation pressure is evaluated live from it using CoolProp [39]. Because T_{sat} and p_{sat} are uniquely coupled on the saturation curve, this gives a second, independent indication of the system pressure alongside the manometer PS01, and the difference between the two is the indicator of non-condensable gasses described in Section B.4.

The loop flow rate is recorded in two forms. The instantaneous value is displayed as reported by the meter, but the first experiments showed that the flow in this loop does not settle to a constant instantaneous value: it pulsates continuously, and under some conditions alternates between near-zero flow and short surges several times larger than the mean. An instantaneous reading is therefore not representative of the circulation. The

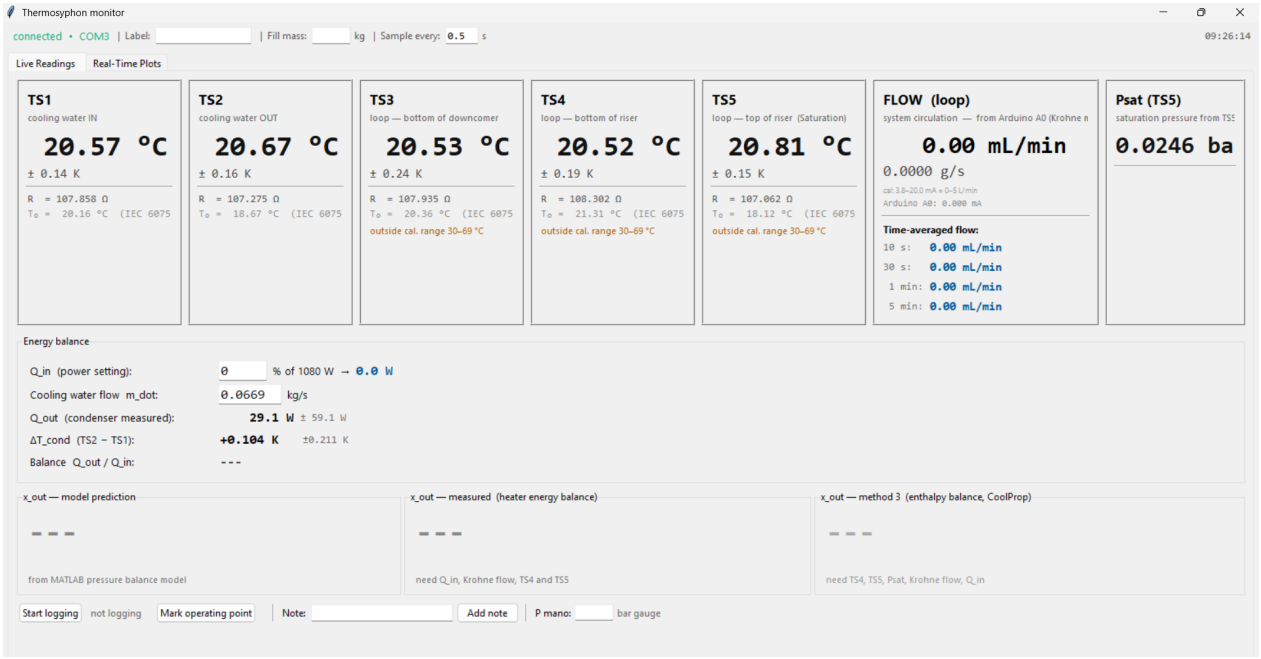


Figure 21: Monitoring software interface.

software additionally computes running means over $10\,\mathrm{s}$, $30\,\mathrm{s}$, $1\,\mathrm{min}$ and $5\,\mathrm{min}$, each evaluated from the recorded timestamps rather than from a fixed number of samples so that it remains correct if the acquisition interval changes. The one- and five-minute means are plotted live alongside the instantaneous value, and the five-minute mean is used in every derived quantity.

The resolution of the logged flow signal is set by the acquisition chain rather than by the meter. The Arduino digitizes the voltage across the $220\,\Omega$ resistance with a 10-bit converter against a $5\,\mathrm{V}$ reference, so the smallest voltage step it can resolve is $4.89\,\mathrm{mV}$, corresponding to a current step of $0.022\,\mathrm{mA}$. With the loop at rest the recorded signal is centered on $3.843\,\mathrm{mA}$ and occupies adjacent levels spaced by exactly this step, principally within one step either side, which confirms that the variation is quantization in the converter and not instability of the flowmeter. The central value was adopted as the zero-flow reference. Over the configured span one step corresponds to approximately $7\,\mathrm{mL\,min^{-1}}$. Because readings below the reference are clamped to zero while those above it are not, the indicated flow at rest averages slightly above zero, at around $2\,\mathrm{mL\,min^{-1}}$. The resolution is fixed and independent of the flow rate, so it represents a larger relative uncertainty at the low circulation rates of steady operation than at high ones; this is one of the reasons the five-minute mean, which averages over many quantization steps, is used in all derived quantities.

The software also records free-text notes. At any point the operator can type a note, which is written into the record with the time it was entered. Every action taken on the loop, from starting the vacuum pump to changing a bath setpoint, is logged this way, so that a change in the measured signals can later be matched to its cause. A second button marks operating points. Pressing it flags every record that follows until it is pressed again, which highlights a chosen interval in the log. The software makes no judgement about whether the loop is at steady state; the flag only makes the interval easy to find again during post-processing.

The heat rejected at the condenser is evaluated from the cooling water inlet and outlet temperatures TS01 and TS02,

$$\dot{Q}_{out} = \dot{m}_{cw}\, c_p\, (T_{02} - T_{01})\,, \tag{58}$$

with c_p evaluated from CoolProp at the mean cooling water temperature. No flow meter is installed on the cooling circuit, so $\dot{m}_{cw}$ is not measured live. It was instead determined beforehand by timing how long the circuit takes to fill a one-liter vessel at the bath pump setting used (pump stage 2), and is entered into the software by hand. The measured value was approximately 250 L/h. The quantity serves to check whether

the heat rejected at the condenser is consistent with the heat supplied by the heater, and is not used in any other derived result.

Its uncertainty is dominated by the temperature difference. The uncertainty in $\dot{m}_{cw}$ from repeated collection is neglected, giving

$$U_{95}\left(\dot{Q}_{out}\right) = \dot{m}_{cw}\, c_p \sqrt{U_{95,01}^2 + U_{95,02}^2}. \tag{59}$$

With the calibrated uncertainties of TS01 and TS02 this amounts to roughly $0.2\,\mathrm{K}$ in ΔT. The absolute uncertainty in $\dot{Q}_{out}$ is therefore fixed by the sensors and is independent of the heat input, so its relative significance grows as the heat input falls. At the cooling water flow rates used the temperature rise across the condenser is only of order $1\,\mathrm{K}$ at moderate heat input, so the energy balance serves as a qualitative consistency check rather than a quantitative measurement of the heat rejected.

The outlet quality of the heated section is computed live in two ways. The first is an energy balance across the heater, in which the heat input is divided between raising the entering liquid from TS04 to the saturation temperature and evaporating a fraction of it,

$$x_{out} = \frac{\dot{Q}_{in} - \dot{m}\, c_{p,l}\,(T_{sat} - T_{04})}{\dot{m}\, h_{fg}}. \tag{60}$$

The second writes the same balance in terms of specific enthalpy, $\dot{Q}_{in} = \dot{m}\,(h_{05} - h_{04})$, with the outlet state two-phase at T_{sat} so that $h_{05} = h_l + x_{out}\, h_{fg}$, and solves for the quality:

$$x_{out} = \frac{\dot{Q}_{in}/\dot{m} + h_{04} - h_l}{h_{fg}}. \tag{61}$$

Here h_{04} is the enthalpy of the subcooled liquid at the measured TS04 temperature and system pressure, while h_l and h_{fg} are evaluated at T_{sat}, all from CoolProp. In both, $\dot{m}$ is the five-minute mean loop flow rate and $\dot{Q}_{in}$ the heater power.

The two expressions are the same energy balance written in different variables, since $h_l - h_{04} \approx c_{p,l}(T_{sat} - T_{04})$ for a subcooled liquid, and they agree to within the small difference between a constant-c_p approximation and the exact enthalpy change. They are therefore not independent measurements of x_{out} but a check that each formulation has been implemented correctly. The outlet quality predicted by the loop model of section 6.2 at the applied heat input is displayed alongside them for direct comparison.

The complete record is written continuously to a spreadsheet file, and each row is simultaneously appended to a plain-text copy flushed to disk immediately, so that a software or power failure during a long run loses no data.

7.3.1 Measurement Uncertainty

Three quantities are measured directly: temperature, loop circulation rate and heater power. Everything else reported in chapter 8 is derived from them, so the uncertainty of the derived quantities follows from how the measured ones propagate.

The temperature chain is the best characterized of the three. After the correction of Section 7.1 the expanded uncertainty is between 0.16 K and 0.24 K for the loop sensors TS03 to TS05, and 0.134 K and 0.159 K for the cooling water sensors TS01 and TS02. These are the values carried into every temperature difference quoted below.

The circulation rate is the weakest measurement and limits most of what follows. Beyond the quantization described above, the meter itself displayed negative readings during the rapid state, consistent with the flow reversal seen in the transparent section, and those intervals enter the record as zero rather than as negative values. The mean circulation rate reported for the rapid state is therefore an upper bound on the net flow around the loop, not a symmetric estimate of it. In the slow state the signal sits less than two counts above

the zero reference, so the resolution rather than the bias dominates, and differences between individual operating points in that state are not resolved by the instrument.

The heat rejected at the condenser is obtained from Equation (58). Its uncertainty is dominated by the temperature difference across the cooling water, since the flow rate contributes about one percent and is neglected, and amounts to an expanded uncertainty of approximately 60 W. Because this is fixed by the sensors and the cooling water flow rate, it is independent of the heat input, and its relative weight therefore grows as the heat input falls.

The outlet vapor quality inherits the circulation rate uncertainty and amplifies it, because the mass flow rate appears in both terms of the heater energy balance of Equation (60). Propagating the standard deviations of Table 10 gives roughly 20 to 30 percent on the quality in the slow state, and more in the rapid state, where the sensible term carries most of the heat input and scales with the flow rate alongside the latent term. The quality values are therefore useful for the contrast between the two states, which spans two orders of magnitude, and not for the differences between individual operating points within either one.

8 Results

The loop was operated on eleven separate days. Two variables were changed deliberately between runs. The first is the amount of water charged, expressed as a filling ratio FR, defined in equation 26, the volume of liquid admitted divided by the total internal volume of the system, $(3.758\,\mathrm{L})$. Taking the liquid density at the operating saturation temperature of $50\,^{\circ}\mathrm{C}$, the charge was varied between $FR = 0.48$ $(1.8\,\mathrm{kg})$ and $FR = 0.70$ $(2.6\,\mathrm{kg})$. The second is the starting condition: a run either began from a freshly charged loop, with boiling water admitted into the evacuated system the same morning, or from a loop that had been charged on a previous day and left to cool to ambient overnight. Within each day the heat input was stepped through a sequence of levels, typically in increments of $50\,\mathrm{W}$ held for thirty to sixty minutes, while the expansion vessel and condenser settings were adjusted to hold the saturation condition near $50\,^{\circ}\mathrm{C}$. Several days repeat an earlier one under nominally identical conditions, to ensure that we get repeatable results. Table 9 summarizes the campaign. The darker rows, Days 2 and 9, are the runs analyzed in detail in this report; the lighter rows are shown for comparison only.

Table 9: Overview of the experimental campaign.

Day	Filling ratio	Starting condition	Purpose
1	0.48 $(1.8\,\mathrm{kg})$	Fresh charge	First operation at this charge
2	0.48 $(1.8\,\mathrm{kg})$	Overnight rest	Effect of starting condition
3	0.59 $(2.2\,\mathrm{kg})$	Fresh charge	Increased charge
4	0.70 $(2.6\,\mathrm{kg})$	Fresh charge	First operation at this charge
5	0.70 $(2.6\,\mathrm{kg})$	Overnight rest	Effect of starting condition
6	0.70 $(2.6\,\mathrm{kg})$	Fresh charge	Repeat of Day 4
7	0.70 $(2.6\,\mathrm{kg})$	Fresh charge	Repeat of Day 4
8	0.65 $(2.4\,\mathrm{kg})$	Fresh charge	First operation at this charge
9	0.65 $(2.4\,\mathrm{kg})$	Fresh charge	Repeat of Day 8
10	0.65 $(2.4\,\mathrm{kg})$	Fresh charge	Repeat of Day 8
11	0.65 $(2.4\,\mathrm{kg})$	Overnight rest	Effect of starting condition

8.1 Circulation & Two-phase flow visualization

Before the behavior is examined in detail, the first result is that the loop works. On every run made from a fresh charge the flow meter registered circulation within about fifteen minutes of the heater being switched on, with no pump and no external input beyond the heat supplied at the evaporator and removed at the

condenser. Across the campaign the circulation rate ranged from roughly $10\,\mathrm{mL\,min^{-1}}$ to $280\,\mathrm{mL\,min^{-1}}$ depending on the run and on how long it had been running, while the saturation temperature was held near $50\,^{\circ}\mathrm{C}$ throughout.

The flow in the measurement section can be followed directly through the transparent pipe. Figure 22 shows it over $1.4\,\mathrm{s}$, in consecutive frames $0.1\,\mathrm{s}$ apart. Liquid lies along the bottom of the pipe with a flat interface above it, and across that layer travel discrete fronts, each with a raised nose that crosses the field of view in a few tenths of a second and leaves the stratified appearance behind it. The direction is not fixed. Over the first half of the sequence the liquid travels one way along the pipe and over the second half it travels the other, so the flow in this section reverses within about a second.

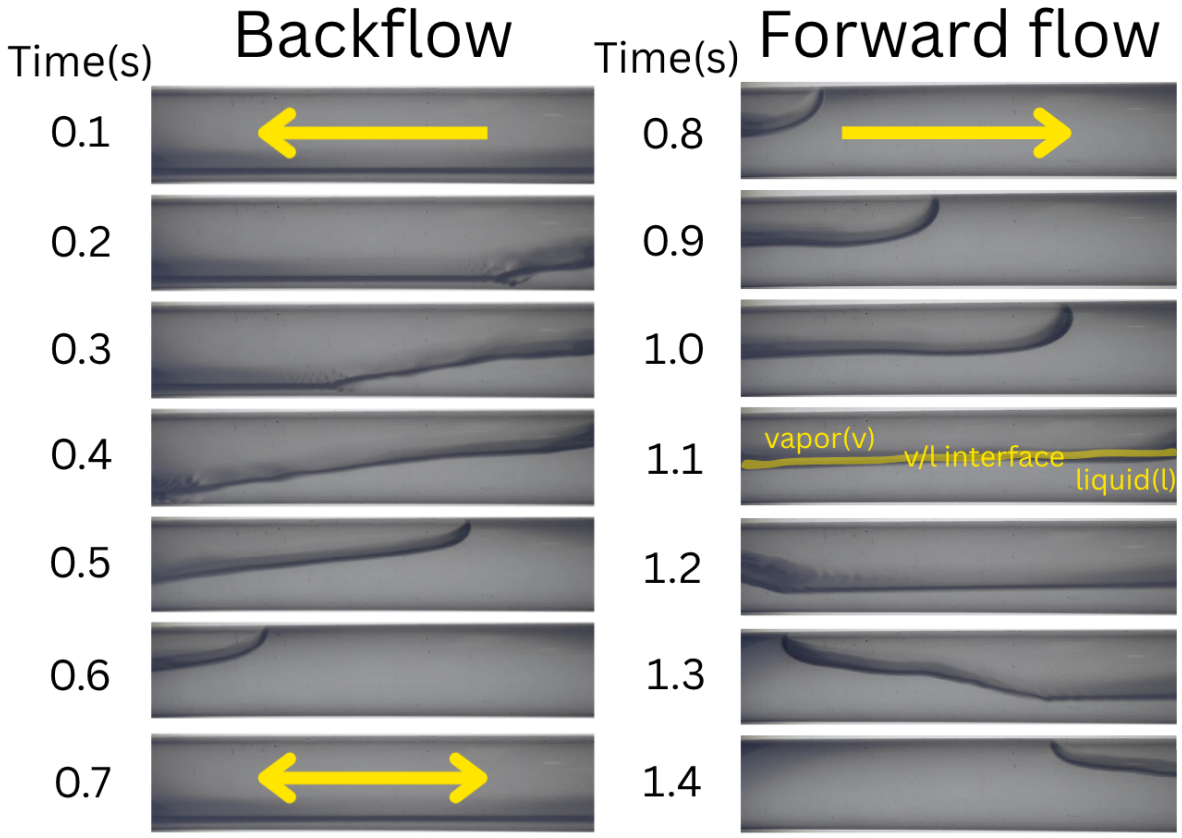


Figure 22: The measurement section over $1.4\,\mathrm{s}$, in consecutive frames $0.1\,\mathrm{s}$ apart.

Figure 23 places every operating point reached on Day 9 against the circulation rate the steady-state model of Section 6.2 predicts at the same heat input. The measured points do not fall on a single curve. They separate into two groups roughly an order of magnitude apart, one above $170\,\mathrm{mL\,min^{-1}}$ and one near $12\,\mathrm{mL\,min^{-1}}$, with no points between them, and the model passes through the lower group while lying far below the upper one. The loop therefore ran in two distinct ways on the same day and at the same heat inputs, and only one of them is described by the model that was used to design it.

The results are not presented day by day. The behavior of the loop is dominated by a single feature that recurs on every run made with a fresh charge, namely a transition from a state of rapid, strongly intermittent circulation to one of slow and almost steady circulation, and the sections below are organized around establishing that this feature is real and characterizing the two states on either side of it. Section 8.2 deals first with the saturation control and with the lowest filling ratio tested, which on Day 2 failed for a different and more easily identified reason and which establishes what dryout looks like in this loop. Section 8.3 then establishes the repeatability of the two charges that were repeated, Sections 8.4 to 8.6 characterize the two circulation states and the transition between them. What could cause that change is taken up in Chapter 9. Day 9 is used as the representative run wherever a single day is shown, since it was the longest uninterrupted

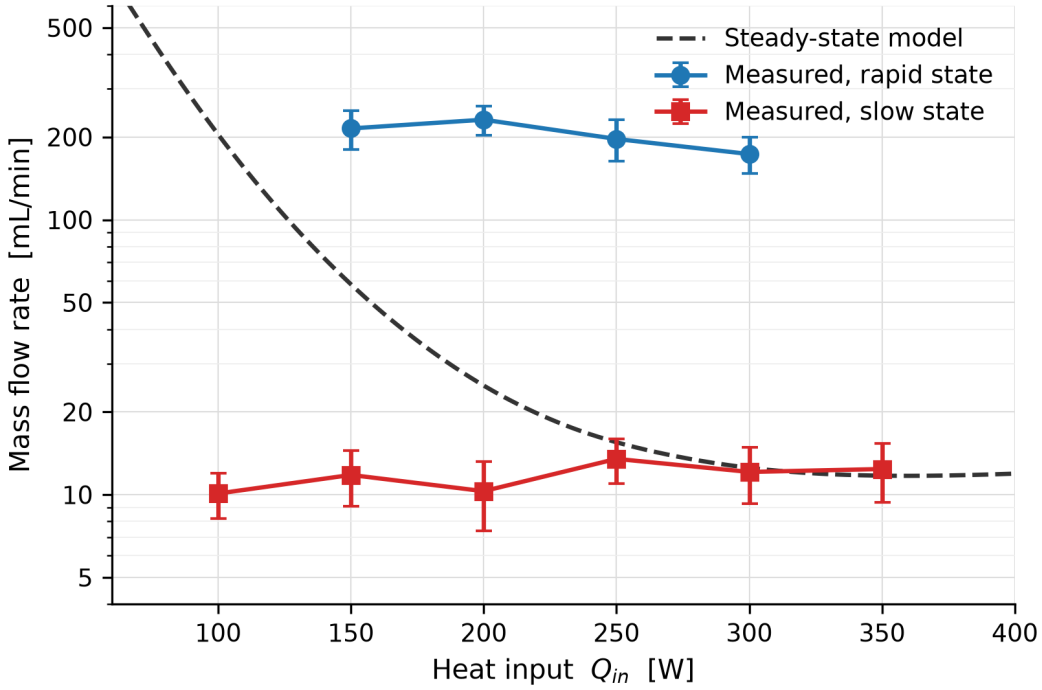


Figure 23: Measured circulation rate in both states on Day 9, against the model prediction.

sequence of heat input steps obtained at a fresh charge.

Days 1 and 3 were commissioning runs made before the acquisition software reached the form described in Section 7.3, and their records are incomplete. They are retained in Table 9 for completeness but are not analyzed in detail. Day 1 is nonetheless mentioned in Section 8.2, since it was the first run on which the transition was observed.

8.2 Saturation control and the limit at low charge

The expansion vessel is intended to fix the saturation condition of the loop independently of the heat balance, and Day 2 shows that it does. Over the course of that run the vessel bath was moved repeatedly between $45\,^{\circ}\mathrm{C}$ and $60\,^{\circ}\mathrm{C}$, and on each occasion the temperature measured at TS05, immediately downstream of the heater, followed the new setpoint and settled close to it. Figure 24 shows part of that sequence. The coupling holds only while the loop is boiling. When the heater was switched off at the end of the run and the bath was then raised from $50\,^{\circ}\mathrm{C}$ to $60\,^{\circ}\mathrm{C}$, TS05 no longer followed it, since with no vapor being generated there is no continuous vapor path between the heated section and the vessel and the interface in the vessel no longer sets the pressure that the loop sees. While the loop runs, the saturation temperature is a controlled quantity rather than an outcome.

The same day established the lower limit on charge mass, and the indication that marked it was a rise in TS04 at constant heat input. At $100\,\mathrm{W}$ the loop first ran steadily, with TS04 holding $10\,\mathrm{K}$ to $12\,\mathrm{K}$ below TS05 and the Type K thermocouple between the aluminum block and its insulation, described in Section B.2, reading about $95\,^{\circ}\mathrm{C}$. Over the last half hour of that step, with nothing altered, TS04 began to climb and the gap narrowed to under $7\,\mathrm{K}$. Raising the heat input to $150\,\mathrm{W}$ closed the gap entirely within fifteen minutes while the block temperature rose to $234\,^{\circ}\mathrm{C}$ without settling, and a later attempt at $200\,\mathrm{W}$ repeated the sequence more quickly, passing $291\,^{\circ}\mathrm{C}$ before smoke was seen at the insulation and the heater was switched off at $330\,^{\circ}\mathrm{C}$.

Both readings describe the same event, and Figure 25 sets out the three stages it passes through. At rest the liquid stands at one level throughout, in the riser, in the downcomer and in the expansion vessel, since all three are connected and nothing is driving a flow. Once boiling begins, vapor rises through the riser and

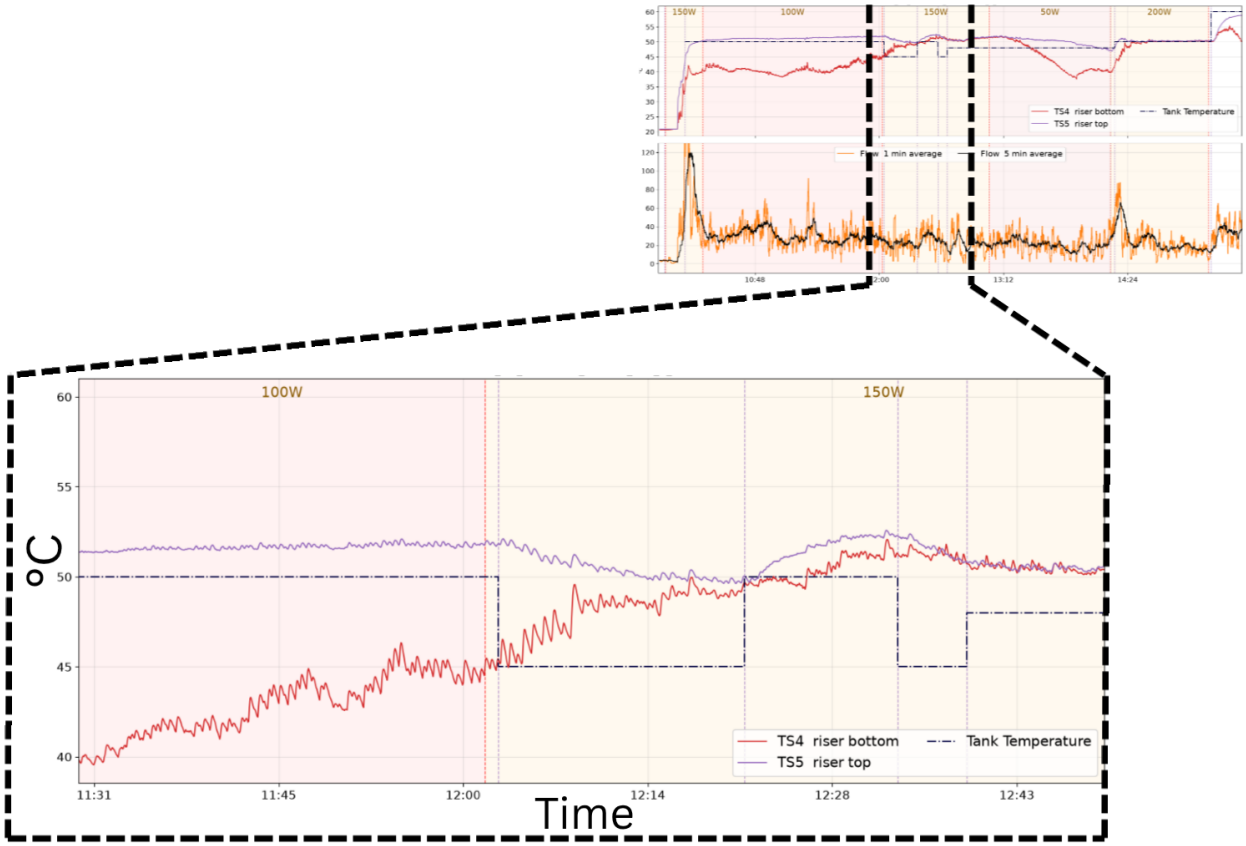


Figure 24: Day 2, $FR = 0.48$. TS05 following successive changes to the expansion vessel bath setpoint.

condensate returns down the downcomer, and the vapor now occupying part of the circuit displaces liquid into the vessel, whose level rises above the one in the riser. Through this stage the heated wall stays covered and the loop runs steadily.

The third stage is the one that ends the run. If the displaced liquid is enough to bring the level in the heated section below the top of the heated length, the exposed part of the heated wall meets vapor rather than liquid, and the heat supplied there raises the vapor temperature instead of evaporating liquid, thus turning the saturated vapor into superheated vapor. Superheated vapor cannot begin to condense until that superheat has been removed, so it travels further into the loop before it does condense, which in turn pushes even more liquid into the vessel. The level in the heater falls further, more wall is exposed, and the process feeds itself. The bottom of the riser, including TS04, is therefore no longer immersed in liquid, and reaches TS05 by conduction along the pipe wall and thermowell, while the heater block, with no boiling film to give its heat up to, overheats. At $FR = 0.48$, the liquid inventory could not sustain a wet wall around the heater at above $100\,\mathrm{W}$.

Day 1, with the same charge admitted that morning, behaved differently. The loop ran at up to $300\,\mathrm{W}$, and after about thirty minutes at that setting the circulation rate fell from roughly 300 mL/min to between 20 and 30 mL/min over about five minutes with no change to any setting, a transition of the same kind as the one described in Section 8.5. Dryout appeared only after prolonged operation at $300\,\mathrm{W}$. The incomplete record from that day does not allow it to be analyzed further, but the contrast with Day 2 is the first indication that the starting condition, and not only the charge mass, decides how the loop behaves.

8.3 Repeatability

Each day of operation was run manually from start to finish, with the vessel bath and the cooling water adjusted according to the response of the loop. The first question to settle is therefore whether the observed

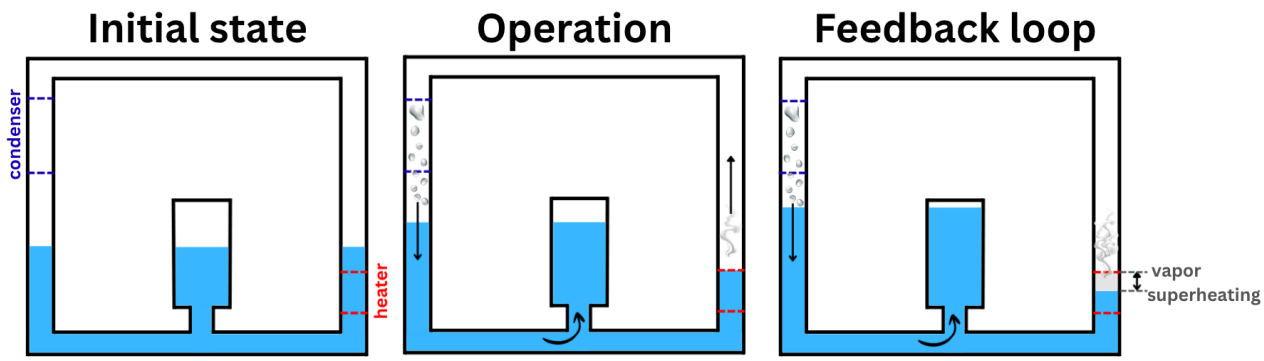


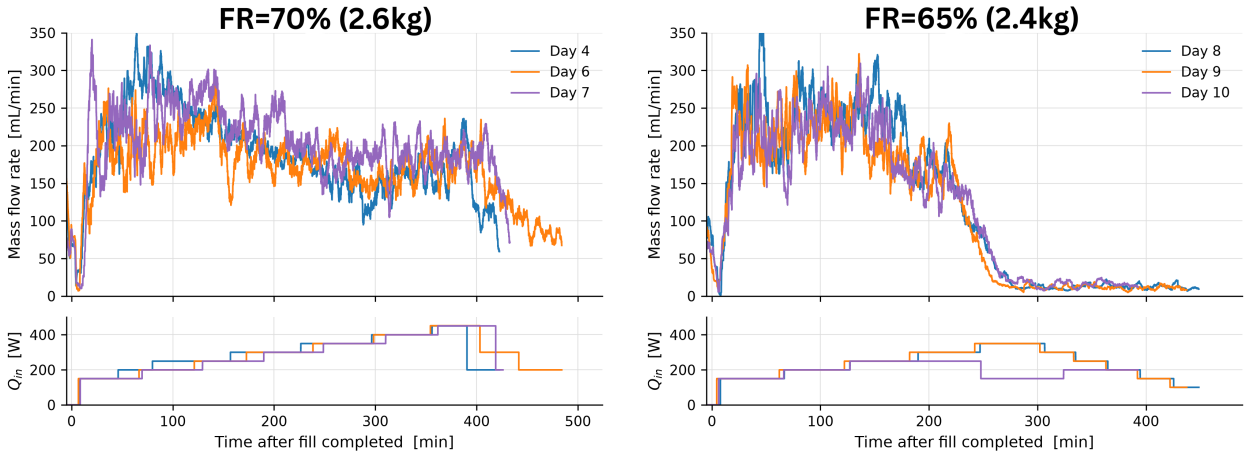
Figure 25: Liquid redistribution at low charge, leading to partial exposure of the heated wall.

behavior is a property of the loop or of the way it happened to be operated on a given day. Three days were run at $FR = 0.70$ and three at $FR = 0.65$, each from a fresh charge, and in each group the days are compared on a common time axis measured from the moment charging was completed.

Figure 26a shows the three $FR = 0.70$ days. The circulation rate rises within about fifteen minutes of the heater being switched on, reaches roughly $280\,\mathrm{mL\,min^{-1}}$, and then declines gradually over the following four hours to between 150 and $200\,\mathrm{mL\,min^{-1}}$. The three traces follow one another closely despite differences of up to an hour in when each heat input step was applied, and the decline is present on all three. The loop does not reach a constant circulation rate at any point during these runs.

Figure 26b shows the three $FR = 0.65$ days. For the first three and a half hours they are indistinguishable from the $FR = 0.70$ case, following the same rise and the same gradual decline. They then depart from it abruptly. On each of the three days the circulation rate falls from roughly $175\,\mathrm{mL\,min^{-1}}$ to approximately $12\,\mathrm{mL\,min^{-1}}$ over about thirty-five minutes, and the onset of that fall occurs at around $240\,\mathrm{min}$ after the charge was completed for all three days. Three independent runs, charged and operated on different days, begin the same transition within about twenty-one minutes of one another.

The lower plot of Figure 26b carries the result that makes this difficult to attribute to the heat input. At the moment of transition Days 8 and 9 were at $300\,\mathrm{W}$, while Day 10 had been held at $250\,\mathrm{W}$ for two hours. The transition therefore does not wait for a particular heat input. On Day 10 the heat input was in fact reduced to $150\,\mathrm{W}$ part way through the decline, and the circulation rate continued to fall exactly as before.



(a) $FR = 0.70$ (b) $FR = 0.65$

Figure 26: Circulation rate and heat input on the repeated runs.

The two runs begun from a loop charged on the previous day and left to cool overnight, Days 5 and 11, never reached the rapid circulation state at all. Their circulation rates peaked at 94 and $61\,\mathrm{mL\,min^{-1}}$ respectively

and remained below $60\,\mathrm{mL\,min^{-1}}$ for the remainder of each run, at filling ratios of 0.70 and 0.65 that produced rates three to five times higher when charged the same morning. The rapid state therefore requires a freshly admitted charge and not merely the correct charge mass, which places the condition that distinguishes the two states in the state of the fluid rather than in the geometry or the inventory alone.

8.4 The rapid circulation state

Day 9 is taken as representative. Table 10 lists, for each heat input step, the mean circulation rate over a window beginning ten minutes after the step was applied, together with the mean loop temperatures and the outlet quality derived from the heater energy balance of Equation 60. The void fraction seen in the table is not measured, since that is the quantity an MRI measurement would supply, and is instead computed from the quality and the mass flux with the Rouhani and Axelsson correlation of Equation 19. The table covers the whole day and is divided into the two states, the lower half of which is taken up in Section 8.6.

Table 10: Day 9 operating points.

State	$\dot{Q}_{in}$ [W]	$\dot{m}$ [mL min^{-1}]	T_{03} [°C]	T_{04} [°C]	T_{05} [°C]	x_{out} [-]	α [-]
Rapid	150	215 ± 35	30.3	42.7	48.5	0.008	0.615
	200	231 ± 28	29.1	41.9	49.5	0.009	0.620
	250	197 ± 34	28.7	35.3	49.2	0.009	0.616
	300	173 ± 26	28.1	32.9	50.7	0.014	0.625
Slow	350	12.4 ± 3.0	18.6	30.5	52.4	0.727	0.644
	300	12.1 ± 2.8	17.3	29.9	52.7	0.631	0.643
	250	13.5 ± 2.5	17.0	35.8	54.1	0.458	0.641
	200	10.3 ± 2.9	19.0	45.3	53.5	0.531	0.640
	150	11.8 ± 2.7	21.0	45.1	53.6	0.330	0.636
	100	10.1 ± 1.9	21.1	43.7	53.3	0.246	0.629

Two features of the rapid state are immediately visible. The first is that the circulation rate falls as the heat input rises, from $215\,\mathrm{mL\,min^{-1}}$ at $150\,\mathrm{W}$ to $173\,\mathrm{mL\,min^{-1}}$ at $300\,\mathrm{W}$. The model of Section 6.2 predicts the same sign, since a higher outlet quality raises the two-phase pressure drop faster than it lightens the riser, but it predicts a far steeper fall, from 58 to $13\,\mathrm{mL\,min^{-1}}$ over the same range. The measured branch is an order of magnitude higher than the model and much flatter. The second feature is that the standard deviation of the five-minute mean is thirteen to eighteen percent of the mean itself. The underlying instantaneous signal is far more variable still: the flow meter alternates continuously between near-zero readings and surges several times the mean, and the transparent section of the loop shows liquid and vapor slugs passing in rapid succession, with the pipe emptying and refilling and with intermittent reversal of the flow direction. Figure 27 shows the measurement section through one of these cycles. Before a pulse arrives the pipe carries a flat interface with liquid along the bottom, during its passage the section fills and the interface is lost, and once it has gone the stratified appearance returns. No single frame describes the state.

The outlet quality computed for this state is between 0.008 and 0.014, which appears at first to contradict the direct observation of a pipe largely occupied by vapor. The contradiction is apparent rather than real. At $0.12\,\mathrm{bar}$ the density of saturated vapor is of order $0.08\,\mathrm{kg\,m^{-3}}$ against roughly $988\,\mathrm{kg\,m^{-3}}$ for the liquid, a ratio of around 12200, so a quality of 0.01 corresponds to a vapor volume fraction well above one half. A mass fraction small enough to be almost invisible in an energy balance is therefore entirely consistent with a pipe that looks full of vapor, and the two observations do not need to be reconciled further.

A third feature is less easily dismissed. At the lower heat inputs of the rapid state the condenser rejects more heat than the heater supplies. At the $150\,\mathrm{W}$ operating point the cooling water carries away $272\,\mathrm{W}$, and at $200\,\mathrm{W}$ it carries away $336\,\mathrm{W}$, against an expanded uncertainty of about $59\,\mathrm{W}$ on that measurement. The loop is therefore receiving heat from somewhere other than the heater. The likely source is the expansion vessel,



(a) Before the pulse



(b) During the pulse



(c) After the pulse

Figure 27: Day 9, rapid state. Measurement section through one pulse cycle.

whose bath was held between 50 and $65\,^{\circ}\mathrm{C}$ during this period while the returning liquid was near $30\,^{\circ}\mathrm{C}$. Consistent with this, the liquid gains rather than loses temperature between TS03 and TS04, the stretch of the return line that carries the branch to the vessel, by

$$\dot{Q}_{ves} = \dot{m}\, c_{p,l}\, (T_{04} - T_{03})\,, \tag{62}$$

which amounts to $183\,\mathrm{W}$ and $205\,\mathrm{W}$ at the two points mentioned above. This attribution should be treated with some caution. The intermittent flow reversal described earlier can also carry hot fluid back down past TS04, and TS04 sits in a thermowell a short distance below the heater block, where conduction through the copper wall raises its reading above the bulk liquid temperature by up to $5\,\mathrm{K}$ to $6\,\mathrm{K}$. Both effects inflate $\dot{Q}_{ves}$, but the imbalance at the condenser does not depend on either. By $300\,\mathrm{W}$ the imbalance has closed, with $291\,\mathrm{W}$ rejected against $300\,\mathrm{W}$ supplied. The vessel, intended only to fix the pressure, appears to be delivering a heat input comparable with the heater's own.

The consequence is that the rapid state should not be regarded as a steady state. A loop rejecting more heat than it is supplied is drawing on something, and over four hours the quantity involved is far too large to be stored energy in the charge. Together with the five-minute mean, which wanders continuously and never settles, this points to an extended transient rather than an operating point. The slow state described below is, by contrast, unambiguously steady, even though it is also the condition in which almost nothing is circulating.

8.5 The transition

Figure 28 shows the transition on Day 9 in detail. The circulation rate begins to fall at 14:30, which is around 238 min after charging, during a step at 300 W that had already been held for forty-five minutes, and reaches its final value approximately thirty-five minutes later. TS05 rises by roughly 2 K across the event, consistent with the loop settling at a slightly higher saturation temperature, while TS04 falls by 3 K. The heat rejected at the condenser is essentially unchanged, 291 W immediately before and 288 W after, so the loop continues to transport very nearly the same heat.

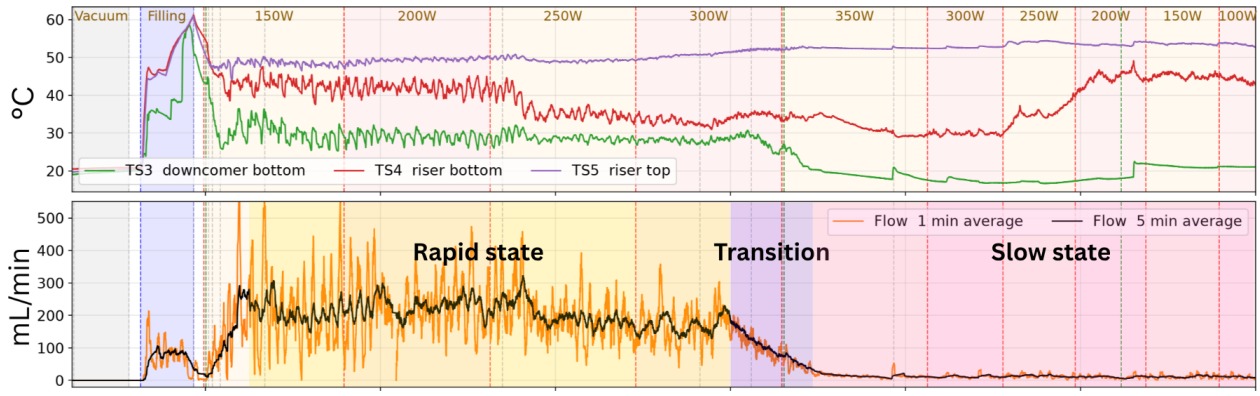


Figure 28: Day 9. Circulation rate, loop temperatures and heat input through the transition.

Three observations constrain what can be responsible. The first has already been given: the transition occurs at the same elapsed time on three days at different heat inputs, 300 W on Days 8 and 9 and 250 W on Day 10, so it is not triggered by reaching a particular heat flux. The second is that holding the heat input constant does not prevent it. On Day 10 the heater had not been touched for more than an hour when the circulation began to fall, as seen in figure 26b. The third is that the transition is not reversible by any adjustment available during operation. On Day 9 the heat input was subsequently stepped back down through 300 W, 250 W, 200 W, 150 W and 100 W, retracing the morning's sequence in reverse, and the circulation rate remained between 10 and 14 mL min^{-1} throughout. Lowering the vessel bath temperature, raising it, and changing the cooling water inlet temperature were all tried on one day or another and none restored the rapid state. Only draining the loop and charging it again with freshly boiled water did so.

Two further explanations can be excluded directly from the records. The transition is not driven by the cooling water, since on Day 10 the condenser inlet temperature measured at TS01 held at 19.9 °C for the whole run and the transition occurred regardless. Neither is it driven by the progressive reduction of the vessel bath setpoint through the day, which is common to every run: Days 4 and 7, both at $FR = 0.70$, were taken down to a vessel setpoint of 30 °C without any transition occurring, while Day 10 transitioned with the vessel still at 45 °C.

Most importantly, the transition is not dryout. Every indication that accompanied the dryout of Section 8.2 is absent here. The heater block temperature does not rise; TS04 falls rather than rising to meet TS05; the condenser continues to reject the same heat; and the loop continues to run in the resulting state for the remainder of the day, which a dried-out heated wall could not do. Whatever changes at 230 min to 250 min changes how the loop transports heat, not whether it is able to.

8.6 The slow circulation state

The lower half of Table 10 gives the operating points reached after the transition. The circulation rate is between 10 and 14 mL min^{-1} at every heat input from 100 W to 350 W, with no trend. The outlet quality, by contrast, rises overall with heat input from 0.25 to 0.73, which is what would be expected of a loop passing a fixed small liquid flow and converting an increasing fraction of it to vapor. In the transparent section no

liquid film can be seen at all. The wall is covered in condensate droplets around its whole circumference, shown in Figure 29, formed as the vapor passing through the section loses heat to the surroundings, and no moving liquid body can be followed anywhere along the pipe. There is no slugging and no reversal.



Figure 29: Day 9, slow state. Condensate on the wall of the measurement section.

A caution attaches to the flow rates in this state. As set out in Section 7.3, one count of the Arduino's converter corresponds to approximately $7\,\mathrm{mL\,min^{-1}}$, so a reading of $12\,\mathrm{mL\,min^{-1}}$ lies less than two counts above the zero reference. The five-minute mean averages over many samples and recovers some resolution below one count, which is why the standard deviations in Table 10 are smaller than the step size, but the differences between individual operating points in this state are not resolved by the instrument and should not be interpreted. The defensible statement is that the circulation rate in the slow state is approximately $12\,\mathrm{mL\,min^{-1}}$ and shows no dependence on heat input that the flow meter can detect.

Figure 23, introduced in Section 8.1, compares both states against the steady-state prediction of the model developed in Section 6.2. Above $250\,\mathrm{W}$ the model is consistent with the slow state within the resolution of the flow meter, predicting $16\,\mathrm{mL\,min^{-1}}$, $13\,\mathrm{mL\,min^{-1}}$ and $12\,\mathrm{mL\,min^{-1}}$ at $250\,\mathrm{W}$, $300\,\mathrm{W}$ and $350\,\mathrm{W}$ against measured values of $13.5\,\mathrm{mL\,min^{-1}}$, $12.1\,\mathrm{mL\,min^{-1}}$ and $12.4\,\mathrm{mL\,min^{-1}}$. Below $250\,\mathrm{W}$ the agreement fails, the model predicting $58\,\mathrm{mL\,min^{-1}}$ at $150\,\mathrm{W}$ where $11.8\,\mathrm{mL\,min^{-1}}$ was measured. The rapid state is not described by the model at any heat input, lying more than an order of magnitude above it throughout. This is consistent with the model having been formulated for a steady, fully developed two-phase circulation, which is a reasonable description of the slow state and not of the rapid one.

The difference between the two states is most usefully expressed as a difference in how the heat is carried. Applying an energy balance across the heater, the heat input divides between raising the entering liquid to the saturation temperature and evaporating a part of it,

$$\dot{Q}_{in} = \underbrace{\dot{m}\, c_{p,l}\,(T_{05} - T_{04})}_{\text{sensible}} + \underbrace{\dot{m}\, x_{out}\, h_{fg}}_{\text{latent}}. \tag{63}$$

Evaluated at the same heat input of $300\,\mathrm{W}$ on Day 9, before the transition seventy percent of the heat leaves the heater as sensible heat in the circulating liquid and thirty percent as latent heat, while afterwards six percent is sensible and ninety-four percent is latent. Figure 30 shows the comparison. The heat supplied is identical by construction and the heat rejected at the condenser agrees to within its uncertainty; what changes is the carrier. Before the transition it is moved by a large flow of barely boiling liquid, and afterwards by a small flow of which most is evaporated.

9 Discussion

The sections above report what the loop did. What follows asks what those observations say about the decisions that produced it, and what they leave unresolved.

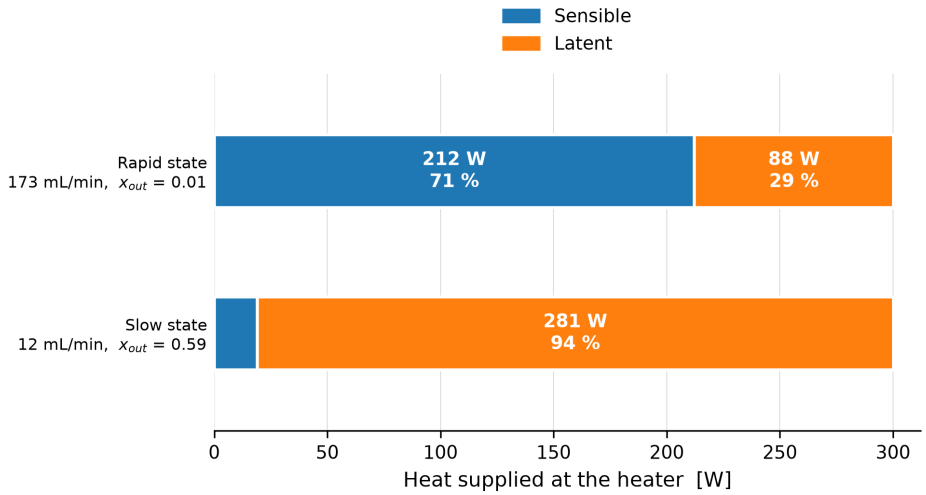


Figure 30: Division of the heat input between sensible and latent transport at $300\,\mathrm{W}$.

9.1 The design and the model

The model of Section 6.2 was built to choose between geometries rather than to predict operation, and the two uses deserve to be separated when its performance is judged. As a means of ranking candidate diameters it was decisive. It rejected the $10\,\mathrm{mm}$ bore because it compresses the whole useful quality range into roughly $130\,\mathrm{W}$, rejected the larger bores because they cannot be driven past part of that range within the installed capacity, and left $13\,\mathrm{mm}$ as the only choice that spreads the range across a span of power wide enough to set and hold. Nothing in the measurements contradicts that ranking. As a predictor of the circulation rate it is far more limited, describing the slow state above $250\,\mathrm{W}$ to within the resolution of the flow meter and nothing else.

That the rapid state is not described is a consequence of the formulation rather than of the closures. The model solves a steady pressure balance for a fully developed two-phase circulation, which the slow state is and the rapid state is not, so no choice of correlation within that framework would have reached it. The failure below $250\,\mathrm{W}$ is a different matter. There the model overpredicts the slow state by a factor of five at $150\,\mathrm{W}$, and the circulation rates it returns in that region are small enough that the steady assumption becomes questionable on its own terms.

The closures themselves are used some way outside the conditions they were built for, and the direction of the resulting error is at least known. The drift-flux model of Rouhani and Axelsson carries validated pressure ranges starting at $1\,\mathrm{bar}$, against an operating pressure of $0.12\,\mathrm{bar}$ and a liquid-to-vapor density ratio of roughly 12200. More specifically, holding the distribution parameter at $C_0 = 1.54$ on the low mass flux branch caps the void fraction the model can return at $1/C_0 = 0.649$, however high the quality rises. The true void fraction at these densities continues towards unity, so the riser mixture the model carries is heavier than the real one, the gravitational term opposing the flow is overestimated, and the circulation rate that comes out is conservative. The two-phase multiplier of the Friedel correlation is likewise dominated by the density ratio, which is the term that makes the predicted circulation rate fall as the heat input rises. The measurements reproduce that sign in the rapid state while lying an order of magnitude above the predicted values, so the mechanism the model identifies appears to be present even where its magnitudes are wrong.

Two design decisions are confirmed directly by the measurements. The first is the expansion vessel. It was conceived to fix the saturation condition independently of the heat balance, and Day 2 shows TS05 following the bath setpoint on every move between $45\,^{\circ}\mathrm{C}$ and $60\,^{\circ}\mathrm{C}$. That is what makes the rest of the campaign interpretable, since the two circulation states were observed at the same saturation condition and can therefore be compared at all. The second is the flow regime prediction. All four maps placed the measurement section in the stratified region across every diameter and every heat input examined, and the sequence in

Figure 22 shows a stratified liquid layer with fronts crossing it. The agreement is in character rather than in detail, since the maps are steady-state constructions and predict neither the intermittency nor the reversal of direction that the sequence shows. The condenser sizing was never tested, the loop having been driven no higher than $350\,\mathrm{W}$ against a design load of $500\,\mathrm{W}$.

One behavior was not anticipated at all, and it bears on everything measured in the rapid state. The expansion vessel was sized and positioned as a pressure reference, and its thermal coupling to the loop formed no part of that reasoning. At the lower heat inputs it delivers a heat input comparable with the heater's own, so the quantity the operator set at the variac is not the heat the loop received. Any comparison between measurement and model at those conditions is therefore a comparison made at the wrong heat input, and the discrepancy in Figure 23 below $250\,\mathrm{W}$ cannot be attributed to the model alone.

Three features of the loop as it was built shaped what could be measured. The horizontal sections at the top are flexible enough to sag between supports, which creates low points where liquid collects and high points where vapor is trapped, so the circuit departs from the uniform bore the model assumes and does so in a way that was neither designed nor quantified. The expansion vessel is connected through a single branch, so liquid entering and leaving it shares one pipe, and nothing in that arrangement obliges the two rates to be equal. And the flow meter resolves about $7\,\mathrm{mL\,min^{-1}}$ per count, which was specified before it was known that the loop would settle at $12\,\mathrm{mL\,min^{-1}}$, so the one state it reaches repeatedly is also the one the instrument resolves worst.

9.2 The observed behavior

The results establish three things. The loop sustains natural circulation at sub-atmospheric pressure, the saturation condition can be held near $50\,^{\circ}\mathrm{C}$ while it does so, and the behavior it settles into repeats across runs charged and operated on different days. What they do not establish is why the loop occupies two distinct states rather than one, or what decides the moment at which it passes between them. That change is the central observation of this work and it remains unexplained.

No mechanism tested during this work accounts for the transition completely. What follows sets out the candidates that remain open, the behavior each would produce, and the difficulty each one faces.

9.2.1 Mechanism that can be excluded

The rapid state occurs only when the loop has been charged the same morning, which makes the heat stored in the freshly admitted water an obvious candidate: the loop might simply be running on that heat until it is used up. The quantity is too small. The water enters at $100\,^{\circ}\mathrm{C}$ and the loop operates near $50\,^{\circ}\mathrm{C}$, so the most the charge can possibly give up is $2.4\,\mathrm{kg}$ cooled by $50\,\mathrm{K}$, which is about $500\,\mathrm{kJ}$. Spread over the four hours the rapid state lasts, that is an average of $35\,\mathrm{W}$, against the surplus of about $120\,\mathrm{W}$ identified at the lower heat inputs in Section 8.4. The real figure is smaller still, since the charge flashes on entry and reaches the operating temperature within minutes, and TS05 then holds near $50\,^{\circ}\mathrm{C}$ for the rest of the run rather than falling as a cooling store would. The surplus is more likely supplied continuously by the expansion vessel bath, as Equation (62) suggests.

This excludes stored heat as the energy that sustains the rapid state. It does not exclude the fresh charge as the condition that permits it, and the contrast with Days 5 and 11 shows that the condition is real. Admitting boiling water into an evacuated loop leaves the system in a particular state, in how the inventory is distributed and in how much dissolved air remains, and that state is lost once the loop has stood overnight. Which part of it matters is unresolved, and is taken up in the mechanisms that follow.

9.2.2 Incomplete wetting of the heated wall with progressive liquid redistribution

The first mechanism begins from the possibility that at $FR = 0.65$ the liquid level within the heated section does not cover the whole heated length. Where the wall is in contact with vapor rather than liquid, the

heat supplied there does not evaporate liquid but superheats vapor that is already present. As explained in section 8.2 and shown in figure 25, vapor leaving the heater with some superheat cannot begin to condense until that superheat has been removed, so it travels further into the loop before condensing, and a larger fraction of the loop volume is occupied by vapor at any instant. The liquid displaced by that vapor has only one place to go, since the expansion vessel is the only compliant volume in an otherwise rigid and liquid-filled circuit, and it accumulates there. The level in the heated section then falls further, more of the heated wall is exposed, and the process reinforces itself.

The connection to the circulation rate is through the driving head. Natural circulation in this loop is driven by the difference in density between the riser and the downcomer, and the downcomer must be kept full of liquid for that difference to be realized. If liquid is progressively withdrawn into the vessel, the column available to the downcomer is reduced and the driving head falls with it, so the circulation rate declines even though the heat input is unchanged. The mechanism is intrinsically progressive and therefore intrinsically time-dependent, which matches the central observation.

9.2.3 Slow net drift of liquid into the expansion vessel

The second mechanism is a weaker form of the same idea and does not require the wetting argument. If the rate at which liquid enters the expansion vessel exceeds the rate at which it returns, even slightly, the loop loses inventory continuously. Nothing in the geometry guarantees that the two are equal: the vessel is connected to the loop through a single branch, so flow into and out of it must share the same pipe, and a persistent imbalance of a few milliliters per minute would be enough to remove a significant fraction of the working inventory over three or four hours. The consequence is the same as above, a falling liquid column and a falling driving head, reached without any assumption about the state of the heated wall.

9.2.4 Accumulation of non-condensable gas

The third mechanism is time-dependent by construction, which is a property the two above have to be argued into. Air dissolved in the charge comes out of solution as the water boils, is swept to the condenser with the vapor, and cannot condense there. It therefore accumulates at the coldest part of the circuit and progressively blankets the condensing surface, which is the behavior reported by He et al. [40] and by Song et al. [41] and which motivated the evacuation standard adopted in Section B.4. The charge is boiled before admission precisely to prevent this, but boiling at atmospheric pressure removes dissolved air only to the extent that the solubility falls, and the residual is released progressively once the water enters a vessel at a few tens of pascal.

The diagnostic for this is available in principle and was described in Section B.4: the difference between the pressure read at PS01 and the saturation pressure inferred from TS05 is the partial pressure of any non-condensable gas present. It cannot be applied here. PS01 is a Bourdon gauge read by eye, and its value was not recorded during the runs analyzed in this chapter, so no time history of p_{NCG} exists. The hypothesis therefore remains untested rather than unsupported.

9.2.5 The common difficulty

Both inventory mechanisms predict that the transition should occur sooner at the larger charge, and the measurements show the opposite. A filling ratio of 0.70 leaves less vapor volume in the loop and less headroom in the expansion vessel than 0.65 does, so a given quantity of liquid redistributed into the vessel should exhaust the available inventory earlier, not later. The $FR = 0.70$ runs nonetheless continued for the whole working day without transitioning, while every $FR = 0.65$ run transitioned within four hours and ten minutes.

This does not dispose of either mechanism, since the quantity that matters is the liquid level in the heated section rather than the total inventory, and the distribution of a given charge between the loop and the vessel at the start of a run is not known. It does mean that neither mechanism explains the observations as it

stands, and that the charge dependence is the first thing any proposed explanation has to account for. The measurements that would separate the three mechanisms are set out in Chapter 10.

10 Conclusions & Recommendations

This chapter sets out the changes the loop needs before it can be taken to the MRI scanner, and then draws the conclusions of the work.

Three mechanisms were left open at the end of the previous chapter, and the current setup cannot distinguish between them. The first group of changes below would close that gap, and each one is included because it measures a quantity the present loop leaves unmeasured. The second group collects the changes that would improve the measurements more generally, or that are needed before the loop is operated inside the scanner.

10.1 Instrumentation needed to resolve the transition

A second flow meter should be installed on the branch connecting the loop to the expansion vessel. Both inventory mechanisms turn on whether liquid is accumulating in the vessel during a run, and at present nothing in the system reports that: the single flow meter on the bottom pipe measures what circulates, not what leaves the circuit. A meter on the branch would give the net transfer into and out of the vessel directly, and would separate a loop that is slowly losing inventory from one that is not. It is the single most informative addition that could be made.

A sight glass should be fitted to the expansion vessel so that the liquid level inside it can be observed during operation. Several of the questions raised by the present measurements turn on how much liquid is held in the vessel at a given moment and whether a vapor space exists above it, and neither can be answered with the instrumentation currently installed. A vessel of the type shown in Figure 31 would make the level directly visible and would allow the migration of liquid between the loop and the vessel to be followed as the experiment proceeds.



Figure 31: Expansion vessel fitted with a sight glass.

A further temperature sensor should be installed at the top of the downcomer, upstream of the condenser inlet. At present the fluid is measured as it leaves the heater at TS05 and again at TS03 at the base of the downcomer, so everything in between is a single unmeasured stretch: the three horizontal sections at the top, the condenser, and the return below it. A sensor placed before the condenser would close the horizontal run between TS05 and that point, allowing the heat exchanged with the surroundings along the top pipework to be measured rather than inferred, and separating it from the heat rejected in the condenser itself. That separation matters because the condition of the fluid arriving at the condenser is not currently known, and the present measurements suggest that a significant part of the condensation may already have taken place before the condenser is reached.

The Bourdon manometer should be replaced with a digital pressure transducer that can be read and logged continuously. The present gauge can only be read by eye, so the system pressure enters the record as occasional manual entries rather than as a measured signal. A logged pressure would allow a direct and continuous comparison with the saturation pressure inferred from the temperature at TS05, and the difference between the two gives the partial pressure of any non-condensable gas present,

$$p_{NCG} = p_{meas} - p_{sat}(T_{05})\,. \tag{64}$$

With the manometer this comparison could only be made at isolated moments and to a resolution comparable with the quantity being sought, so the presence of non-condensable gas could never be excluded with confidence. A transducer would turn it into a continuously monitored quantity, and is what the third mechanism needs in order to be tested at all.

10.2 Further changes to the setup

A flow meter should also be installed on the condenser cooling circuit. The cooling water flow rate is presently determined beforehand by timed collection at a fixed pump setting and entered into the acquisition software by hand, which means it is neither measured during operation nor recorded in the data file. Any change to the pump setting or to the valve position during a run is therefore absent from the record, and the condenser duty computed from it becomes unreliable without this being evident in the data.

The horizontal sections at the top of the loop should be replaced with a rigid material, such as borosilicate glass. The polymer pipe used here is flexible enough to sag between supports, which creates local low points where liquid collects and local high points where vapor is trapped. These irregularities add pressure drop that is neither designed nor quantified, and during operation they were associated with liquid being held in place and with momentary reversals of the flow. A rigid section would remove them and would also give a geometry that matches the one assumed in the model. The change should be accompanied by a new supporting frame that holds the entire circuit level, since the present arrangement allows the pipework to take up its own line between supports.

The AC cartridge heaters should be replaced with direct current (DC) elements before the loop is operated inside the scanner. Alternating current produces a varying magnetic field that may interfere with the imaging, and the elements themselves contain ferromagnetic material. The heater block sits roughly two meters from the bore, so the interference may well prove negligible; this should be checked by measurement before the elements are replaced, since the change is straightforward but not free.

An automatic cutout should also be added to the heater. At present the heater is switched off by hand when the block thermocouple approaches 300 °C, and during the dryout on Day 2 the block reached 330 °C before this was done. A relay driven by the same thermocouple would remove the heater supply without relying on the operator.

Finally, consideration should be given to installing a pump in the circuit so that the mass flow rate can be set directly rather than being left to establish itself. Under natural circulation the flow rate is an outcome of the heat input and the geometry, and the present work has shown that the range of flow rates reached

in practice is both narrow and not freely selectable. A pump would allow the full range of interest to be covered deliberately, and each flow rate to be held steady for as long as the imaging requires. This would require an operating strategy of its own, since the loop would no longer be self-regulating, and it would change the nature of the experiment: natural circulation and pumped circulation are different problems. It is suggested as a route to be considered if the object becomes the visualization of two-phase flow across a range of conditions rather than the behavior of the thermosyphon itself.

10.3 Conclusions and answering the research questions

A sub-atmospheric two-phase thermosyphon with an MRI-compatible measurement section was designed, built and operated over eleven days. It circulates under natural convection alone, holds its saturation condition near $50\,^{\circ}\mathrm{C}$, and repeats its behavior across runs charged and operated on different days, but it does so in two distinct states rather than one, and the model built to design it describes only the slower of them. The conclusions of the work are set out below against each of the three research questions in turn.

1. How can a sub-atmospheric two-phase thermosyphon be designed and modeled to achieve stable natural circulation while satisfying the constraints required for MRI-based flow visualization?

The design was arrived at through a steady-state model that closes the pressure balance around the loop and returns the circulation rate and outlet vapor quality at a given heat input, described in Section 6.2. Used as a design tool rather than as a predictor, it fixed the internal diameter by sweeping candidate bores across the full range of heat inputs and selecting the one that spreads the useful vapor quality range over the widest span of power, giving 13 mm copper for the riser and downcomer. The sections that pass through the scanner are PVC-U of 17 mm bore, since the magnetic resonance environment requires non-magnetic and non-conducting materials as set out in Section 3.3. The condenser and heater were then sized against the top of that range, and the expansion vessel was conceived to fix the saturation condition independently of the heat balance.

The approach worked in the sense that matters. The loop was built, it sustains natural circulation, and Section 8.2 shows that the saturation temperature follows the vessel bath setpoint and is therefore a controlled quantity rather than an outcome. The model also established a limit that no choice of geometry can remove. At a liquid-to-vapor density ratio of roughly $12{,}200$ the liquid superficial velocity is too small for any of the transitions out of stratified flow to be crossed, and all four flow pattern maps place the measurement section in the stratified region across every diameter and every heat input examined, as shown in Section 6.4. Stable circulation and MRI compatibility are achievable together; access to more than one flow regime is not, and that is a consequence of the operating pressure rather than of the design.

2. How does the operating behavior of the developed thermosyphon vary under different operating conditions, and to what extent can this behavior be predicted?

The variable that governs the behavior is not the one the design anticipated. Across eleven days at filling ratios between 0.48 and 0.70, the heat input proved to have far less influence than the starting condition of the fluid. Every run made from a freshly charged loop entered a rapid, strongly intermittent state circulating 170 to 230 mL/min at an outlet quality below 0.015, and every run begun from a loop left to stand overnight failed to reach that state at all. At the lowest filling ratio the heated wall could not be kept wetted above roughly 100 W and the heater overheated, which fixes the lower bound on charge.

The loop then leaves the rapid state on its own. Three independent runs at $FR = 0.65$ transitioned to a slow state of about 12 mL/min within twenty-one minutes of one another, at around 240 min after charging, at different heat inputs and without any change being made, and no adjustment available during operation reversed it. The behavior is therefore reproducible but not controllable.

Predictability follows the same division. The model agrees with the slow state above 250 W, predicting 16, 13 and 12 mL/min against measured values of 13.5, 12.1 and 12.4 mL/min, and fails below it. It does not describe the rapid state at any heat input, lying more than an order of magnitude below the measurements. This is consistent with its formulation for a steady, fully developed two-phase circulation, which describes the slow state and not the rapid one. The transition between them remains unexplained. Three candidate mechanisms are set out in Section 9.2, and the two that rest on liquid inventory both predict the opposite of the observed dependence on filling ratio.

3. To what extent does the developed thermosyphon satisfy the requirements for future MRI-based visualization of two-phase flow, and what modifications are required before MRI measurements can be performed?

The parts of the requirement that concern compatibility are met. The sections that would sit inside the bore are non-magnetic and non-conducting, the saturation condition is controllable and holds near 50 °C, and the rapid state carries flow structures an MRI measurement would capture, a stratified liquid layer crossed by discrete fronts, with film thickness, interface shape and velocity all quantities MRI can resolve and none of which have been reported for water and steam at sub-atmospheric pressure.
The parts that concern the measurement itself are not met. Only one flow regime is reachable, so a sweep across regimes is not available. The two states divide the difficulty between them. The rapid state is the one with something to image, but it drifts continuously, it is not described by the model, and the condition at which any image was taken could not be stated. The slow state is steady and predictable, but no liquid film can be seen in the measurement section at all, only condensate on the wall, and the circulation rate sits below the resolution of the present flow meter. Since the change from one to the other is governed by elapsed time since charging rather than by any setting, the state the loop occupies during a measurement would be imposed rather than chosen.
The changes required before the loop is taken to the scanner are set out in Sections 10.1 and 10.2. The first group is the instrumentation that would resolve the transition, namely a flow meter on the branch to the expansion vessel, a sight glass on the vessel, a temperature sensor upstream of the condenser and a logged pressure transducer. The second covers the loop itself, with the flexible polymer sections at the top replaced by a rigid material on a frame that holds the circuit level, an automatic cutout on the heater, and, if measurement shows it to be necessary, direct current heaters in place of the alternating current cartridges.

A Extended Calibration Data

This appendix gives the setpoints used, the comparison between first and second degree corrections, the test of how many independent measurements each record contains, and the per-sensor results behind the summary in Section 7.1.

A.1 Setpoints

The two sensor groups were calibrated over the ranges in which they operate, since a correction fitted to a narrow range leaves a smaller residual than one stretched across the full range of the facility. TS1 and TS2 measure the cooling water near $20\,^{\circ}\mathrm{C}$; TS3 to TS5 measure the loop, where subcooling in the downcomer and superheat at the heater extend the range to roughly $30\,^{\circ}\mathrm{C}$ to $70\,^{\circ}\mathrm{C}$. The setpoints are listed in Table 11.

Table 11: Calibration and verification setpoints.

Nominal setpoint (°C)	Role	T_{ref} (°C)
Cooling water sensors, TS1 and TS2		
10	calibration	12.202
15	calibration	16.321
20	verification	20.549
25	calibration	24.862
30	calibration	29.591
Loop sensors, TS3 to TS5		
30	calibration	29.578
40	calibration	39.202
45	verification	44.134
50	calibration	49.010
55	verification	53.935
60	calibration	58.822
70	calibration	68.707

The verification points sit near the temperature at which each group actually operates, so the correction is checked where it matters most, while the extreme setpoints were kept as calibration points so the fit covers the whole measured range. The bath controller setting differed from the measured reference temperature by up to about $2\,\mathrm{K}$, which is why T_{ref} and not the nominal setpoint is used in the fit.

Before correction the sensors read between $0.4\,\mathrm{K}$ and $3.5\,\mathrm{K}$ low, and the error is not a fixed offset. TS5, for example, reads $2.85\,\mathrm{K}$ low at $29.6\,^{\circ}\mathrm{C}$ and $3.87\,\mathrm{K}$ low at $68.7\,^{\circ}\mathrm{C}$. The error therefore has an offset and a gain, which is what Equation (66) predicts. Since the sensors are identical and were immersed together in the same bath, a gain that differs this much between them must come from the individual converter boards.

A.2 Choice of polynomial degree

Table 12 compares first and second degree fits.

Table 12: Comparison of first and second degree fits. All values in K.

	Degree 1			Degree 2		
Sensor	$r_{\mathrm{max,cal}}$	$r_{\mathrm{max,ver}}$	S_r	$r_{\mathrm{max,cal}}$	$r_{\mathrm{max,ver}}$	S_r
TS1	0.067	0.008	0.064	0.024	0.060	0.037
TS2	0.085	0.033	0.077	0.038	0.043	0.059
TS3	0.148	0.090	0.117	0.140	0.073	0.141
TS4	0.126	0.091	0.090	0.119	0.078	0.108
TS5	0.089	0.184	0.073	0.066	0.143	0.063

The second degree fit improves $r_{\max,\mathrm{cal}}$ for every sensor, which proves nothing: an extra free parameter always fits the points it was built from more closely. The useful test is the verification points, which no fit has seen. For TS1 and TS2 the quadratic makes them worse, from $0.008\,\mathrm{K}$ to $0.060\,\mathrm{K}$ and from $0.033\,\mathrm{K}$ to $0.043\,\mathrm{K}$. The verification point at $20.5\,^{\circ}\mathrm{C}$ falls in the largest gap between the cooling water setpoints, which is exactly where added curvature is least constrained by data, so the gain at the calibration points is bought at the cost of the region between them. For TS3 to TS5 the verification residuals barely move. Nothing is gained anywhere, and a first degree correction was adopted for all five sensors. This also matches the error model: the converter contributes an offset and a gain, both linear, and the curvature of the Pt100 itself is already removed by Equation (67).

A.3 Independence of the samples

Figure 32 shows how long each channel stays correlated with itself, measured from a continuous $29\,\mathrm{min}$ record at the $50\,^{\circ}\mathrm{C}$ setpoint. The reference channel becomes effectively independent after about $30\,\mathrm{s}$, which matches the thermal response of the probe. The Arduino channels take between $80\,\mathrm{s}$ and $220\,\mathrm{s}$, so successive samples ten seconds apart carry almost the same information. Table 13 gives the resulting number of independent measurements in a forty sample record.

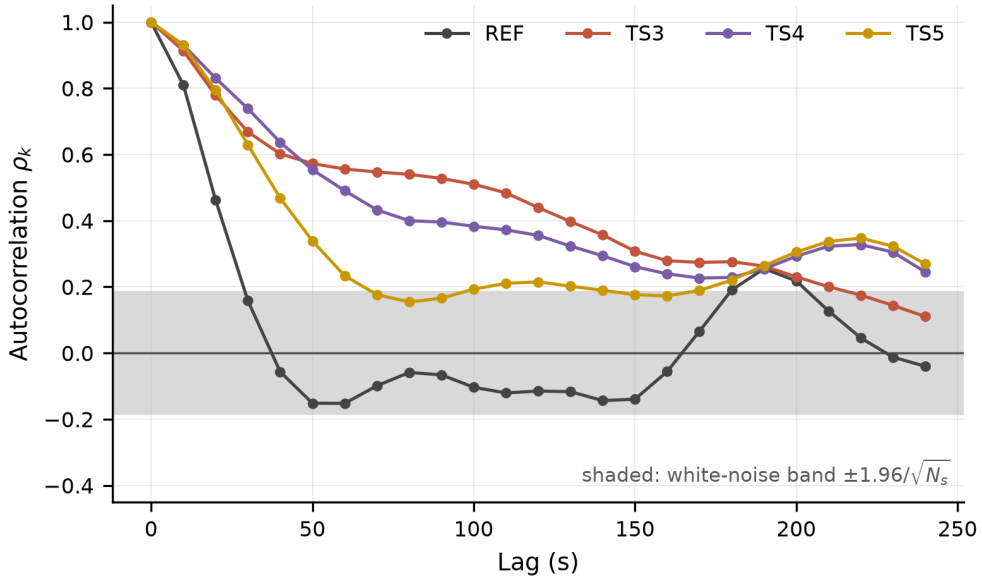


Figure 32: Autocorrelation of each channel at the $50\,^{\circ}\mathrm{C}$ setpoint.

Table 13: Independent measurements contained in a 40 sample record at the $50\,^{\circ}\mathrm{C}$ setpoint.

Channel	Decorrelation time (s)	τ (samples)	N_{eff}
Reference	30	3.5	11.3
TS3	220	20.5	2.0
TS4	210	21.1	1.9
TS5	80	7.8	5.1

A.4 Results

The two cooling water sensors meet the acceptance criterion everywhere, with expanded uncertainties of $0.134\,\mathrm{K}$ and $0.159\,\mathrm{K}$ at $20\,^{\circ}\mathrm{C}$.

Table 14: Calibration results, first degree fit, $T_{\mathrm{corr}} = c_1 T_{\mathrm{raw}} + c_0$. U_{95} is evaluated at 20 °C for TS1 and TS2 and 50 °C for TS3 to TS5.

Sensor	c_1	c_0 (°C)	$r_{\mathrm{max,cal}}$ (K)	$r_{\mathrm{max,ver}}$ (K)	S_r (K)	U_{95} (K)
TS1	1.002220	0.356938	0.067	0.008	0.064	0.134
TS2	1.014593	1.731703	0.085	0.033	0.077	0.159
TS3	1.010154	-0.041747	0.148	0.090	0.117	0.241
TS4	1.004282	-0.885147	0.126	0.091	0.090	0.188
TS5	1.026178	2.219544	0.089	0.184	0.073	0.156

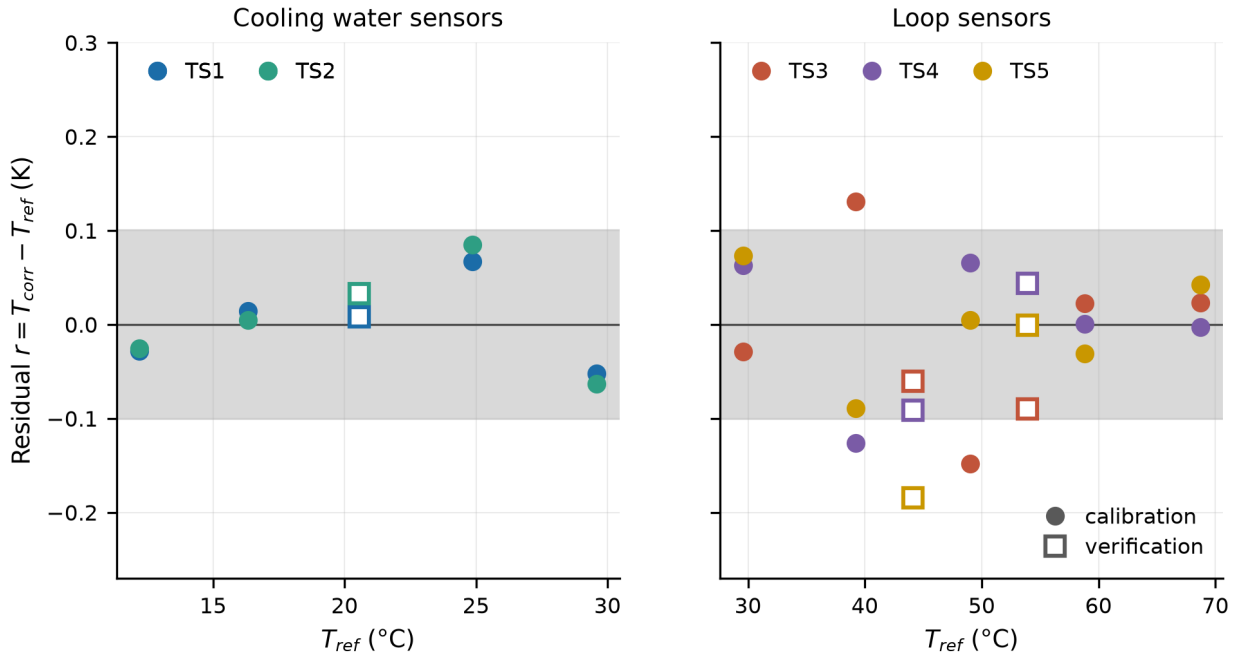


Figure 33: Residuals after the first degree correction. Shaded band marks the ±0.1 K acceptance criterion.

The loop sensors do not, and how each one fails is worth noting. TS3 and TS4 exceed the criterion at the calibration setpoints, at 0.148 K and 0.126 K, but meet it at both verification setpoints. A correction that predicts two withheld temperatures better than it reproduces the points used to build it means the line itself is right and the individual setpoint means are scattered around it. That is the short-record problem of Section 7.1.1, and the sizes match: the mean of a forty sample record was separately measured to move by up to 0.234 K for TS3 and 0.128 K for TS4.

TS5 fails the other way round, at one verification point only. Its residual is −0.184 K at 44.1 °C and −0.001 K at 53.9 °C. A correction that was genuinely wrong would miss both points, and a real curvature would miss them in opposite directions. Neither happens, and the size of the one bad residual matches the 0.191 K drift measured for that channel, so the 45 °C record simply began at an extreme of the slow drift.

Even where the criterion is missed, the calibration reduces the loop measurement uncertainty from ±0.55 K at 50 °C to between ±0.16 K and ±0.24 K, and the cooling water measurement from ±0.40 K to ±0.16 K or better. The improvement is largest for the loop sensors, because the Class B tolerance grows with temperature through its $0.005\,|T|$ term while S_r does not.

B Description of the Experimental Setup

This appendix describes how each subsystem of the loop shown in Figure 19 and Figure 20 was built. It covers the materials and dimensions chosen for each section, the hardware installed, and the reasoning behind the choices made during assembly. The complete component list, with suppliers and part numbers, is given in Section 6.9.

B.1 Piping and connectors

The loop is constructed from two materials, divided according to where thermal contact matters and where MRI compatibility is required. The bottom horizontal pipe, the riser and the downcomer are made from soft copper tube of 13 mm inner and 15 mm outer diameter, supplied by Wildkamp. Copper is used in these sections because they carry the heating and condensation zones, where the high thermal conductivity of the wall is required to transfer heat efficiently between the working fluid and the heater block and cooling water respectively.

The three horizontal sections at the top of the loop, which are the sections intended to pass through the scanner bore, are made from PVC-U, which is neither magnetic nor electrically conductive and therefore satisfies the constraints imposed by the magnetic resonance environment. Two grades were used, both supplied by Eriks. The left pipe is fully PVC-U of 12.4 mm inner and 16 mm outer diameter, while for the rest of the horizontal section a transparent PVC-U section of 17 mm inner and 20 mm outer diameter is used, so that the flow pattern can be inspected directly by eye during startup and operation, independently of the MRI measurement.

The corners of the loop and the branches to the expansion vessel, pressure gauge and charging port are formed with brass Bonfix T-junctions and elbows on the copper sections, also supplied by Wildkamp. The transitions between copper and polymer are made with PVC-U three-part union couplings, shown in Figure 34, which allow the polymer sections to be disconnected and removed without disturbing the remainder of the loop. This was a practical requirement, since the polymer sections are the parts that must be transported and installed within the scanner.



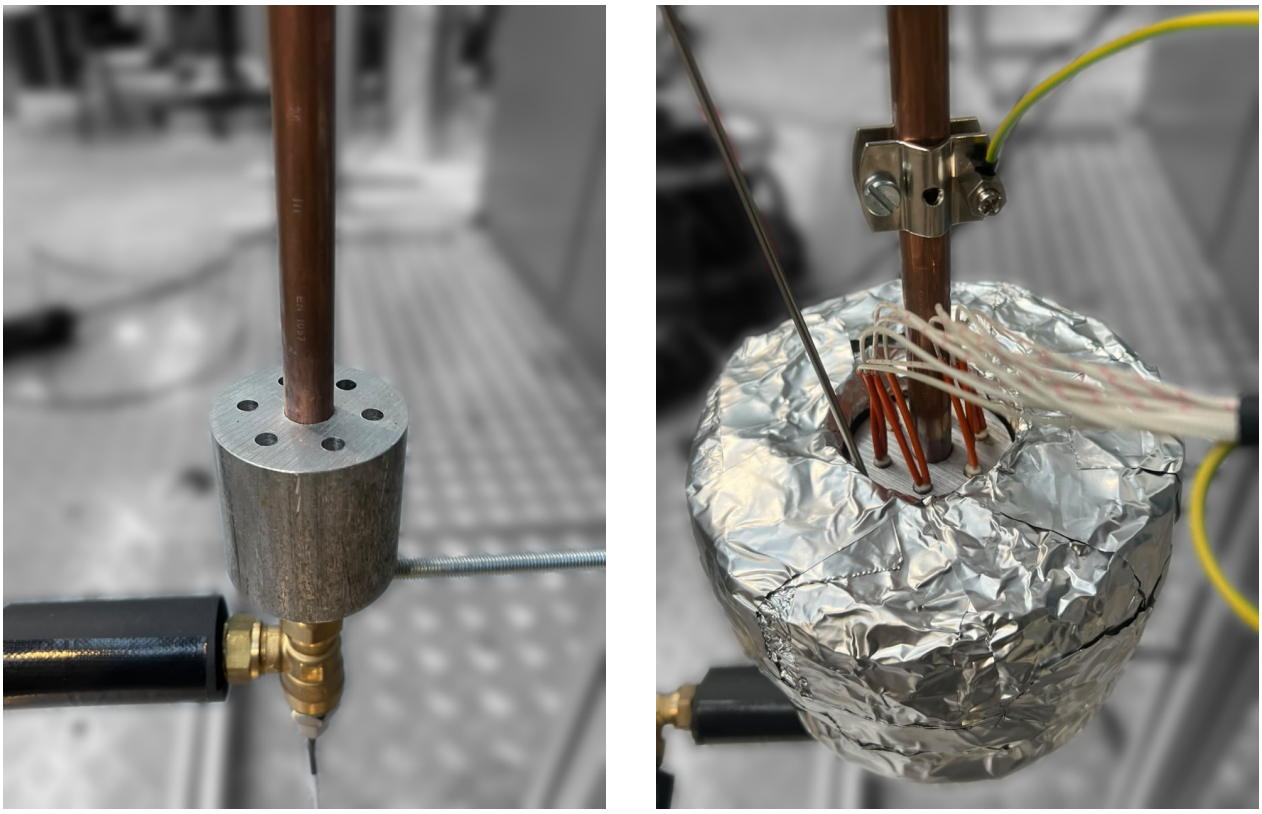
Figure 34: Three-part union coupling ($Cu \rightarrow$ PVC-U)

The complete set of piping and fitting components, with suppliers and part numbers, is listed in Table 2 in Section 6.9.

B.2 Heating System and Control

The heater consists of six cylindrical AC cartridge heaters rated at 180 W each, giving a nominal total of 1080 W. The cartridges are embedded in a cylindrical aluminum block clamped around the copper riser, so that heat is conducted through the block and into the pipe wall homogeneously and radially. The block is wrapped in rock wool insulation, rated to 500 °C, to limit heat loss to the surroundings and to keep the heat

flux directed into the working fluid. The rock wool itself is wrapped with aluminum foil so that it can stay in place. The bare block and the assembled arrangement with the cartridges installed and the insulation fitted are shown in Figure 35.



(a) (b)

Figure 35: Heater assembly. (a) The machined aluminum block. (b) The block with the six cartridge heaters installed and the rock wool insulation fitted around it.

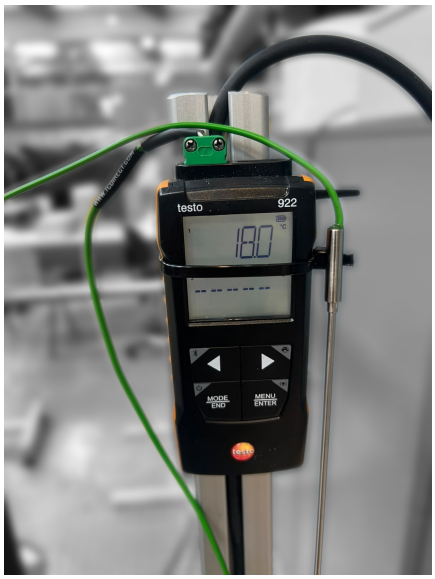
The cartridges are wired in parallel and supplied through a variable autotransformer, which regulates the supply voltage and thereby the delivered power. The setting is entered on a digital display as a percentage of the total installed capacity, so that a given percentage corresponds to the same fraction of the 1080 W nominal rating. The control box housing, the autotransformer and its display were built by Martijn Karsten, the electrical technician of the Process and Energy department at TU Delft (its internal wiring and control logic were not specified as part of this work).
A Type K thermocouple is inserted between the aluminum block and the surrounding insulation to monitor the block temperature during operation, and is read on a separate display independent of the heater control box. This measurement is not used in any of the derived quantities, but serves as a protective indication: if the liquid inventory becomes insufficient to keep the heated wall wetted at the applied heat flux, the wall dries out, the heat is transferred to vapor alone, and the block temperature rises rapidly. Monitoring this temperature therefore allows the onset of dryout to be identified and the power withdrawn before the block or the cartridges are damaged. The control box and the thermocouple readout are shown in Figure 36.
The exact heating-system components, with suppliers and part numbers, are listed in Table 3 in Section 6.9.

B.3 Cooling and Condenser System

The condenser follows the tube-in-tube arrangement of Section 6.6, with the outer cooling-water jacket formed from soft copper tube (22 mm OD / 19.8 mm ID) joined to the inner condensing tube with brass reducing T-pieces. Cooling water is supplied from a recirculating water bath at a fixed flow rate of 0.069



(a)



(b)

Figure 36: Heater power control and block temperature monitoring. (a) Control box containing the variable autotransformer. (b) Type K thermocouple used to monitor the aluminum block temperature, with its separate display.

kg/s, but a needle valve on the cooling-water line allows this to be throttled manually if needed (no separate flow meter was installed on this line). The heat rejected in the condenser is therefore obtained from an energy balance on the cooling water, using this fixed flow rate together with the inlet and outlet temperatures measured by TS01 and TS02 (Section B.5). The exact condenser and cooling-system components, with suppliers and part numbers, are listed in Table 4 in Section 6.9.



(a) Condenser pipe setup



(b) Lauda Eco Gold RE 620 Water bath.

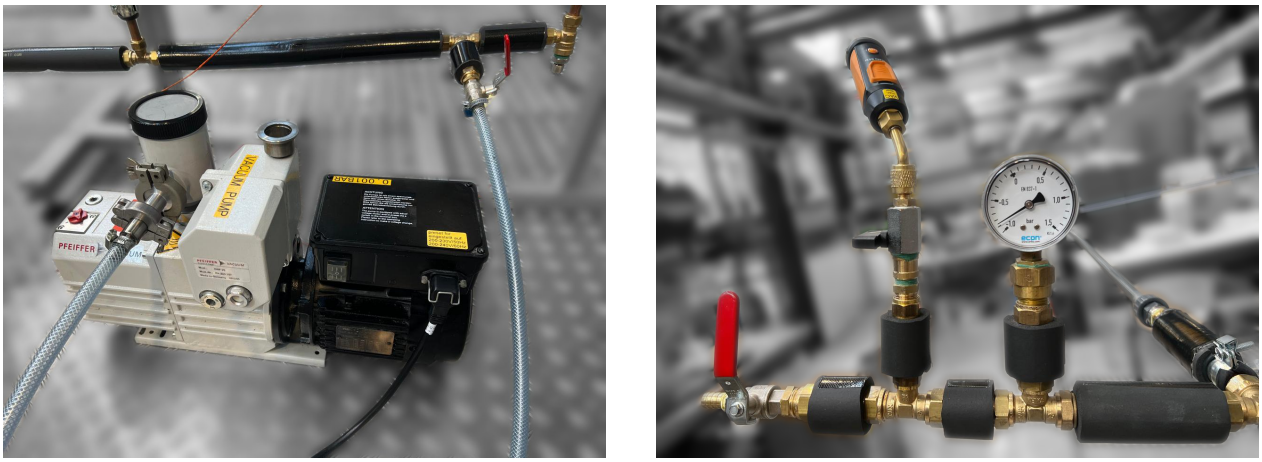
Figure 37: Tube-in-tube condenser and water bath for controlling inlet temperature and flow rate.

B.4 Pressure Control and Expansion System

Before the loop can be charged it must be evacuated, and the standard that evacuation is held to matters more here than in a system operating near atmospheric pressure. Any air remaining in the loop is a non-condensable gas: it does not participate in the phase change, but it does contribute its own partial pressure, so that the total pressure in the system is the sum of the saturation pressure of the water and the partial pressure of whatever gas is left behind. At an operating pressure of order 0.12 bar even a small residual quantity represents a significant fraction of the total, and the loop would then operate at a saturation temperature different from the one the measured pressure implies.

The consequences extend beyond a shifted operating point. He et al. [40] show that non-condensable gas accumulates preferentially in the condenser of a loop heat pipe, where it is swept along by the vapor and left behind as the vapor condenses, progressively blanketing the condensing surface and reducing the area available for heat transfer. Song et al. [41] report the same mechanism in a two-phase closed thermosyphon and quantify the resulting degradation in heat transfer performance. In both cases the gas collects at the coldest part of the circuit and is not removed by continued operation, so the effect is cumulative rather than transient. Evacuating the loop thoroughly before charging is therefore not merely a matter of setting the correct initial pressure, but a precondition for the condenser behaving as designed. On the basis of these studies a target base pressure of the order of a few pascal is adopted before charging.

Evacuation is performed with a Pfeiffer Vacuum rotary vane pump, available in the laboratory, connected to the loop through a dedicated manual valve and a hose pillar. The pressure during evacuation is monitored with a Testo 552i digital vacuum gauge, designated as PS02 in the P&ID (19), and shown in Figure 38(b). The instrument measures from 0 to 26.66 mbar $(0$ to 2666 Pa$)$ with a resolution of one micron of mercury, equivalent to 0.133 Pa, over the lower part of its range. This resolution is what makes the approach to base pressure observable at all: an ordinary Bourdon gauge cannot distinguish between a few pascal and a few hundred. The gauge is isolated from the loop by its own valve once evacuation is complete, both to protect it from the operating pressure and because it has no useful resolution at the pressures reached during operation. The pump and gauge are shown in Figure 38.



(a) (b)

Figure 38: Vaccum pump, Manometer and vaccum gauge

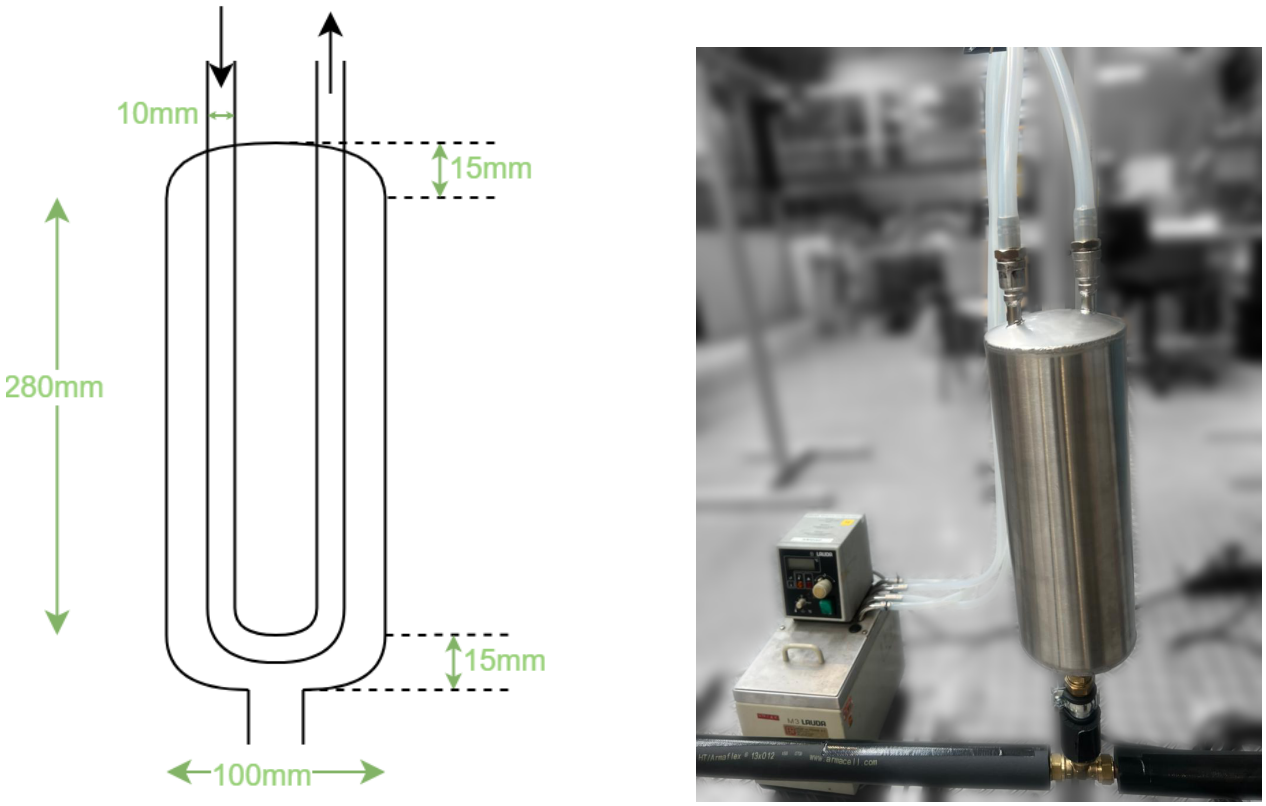
During operation the system pressure is instead read from a Bourdon-tube manometer, designated as PS01, which remains permanently connected to the loop and is shown in Figure 38(b). Its resolution is coarse compared with the vacuum gauge, but it covers the operating range and, unlike the Testo, it can be left in place indefinitely.

The loop therefore carries two independent indications of its operating pressure, and the comparison between them is diagnostically useful. The first is the direct reading of PS01. The second is inferred from

temperature: the fluid immediately downstream of the heater block is at saturation, so the temperature measured there fixes the saturation pressure through the saturation curve, evaluated here using CoolProp [39]. In a system containing only water the two must agree. If non-condensable gas is present the manometer reads higher than the saturation pressure inferred from the temperature, by the partial pressure of the gas, so that

$$p_{NCG} = p_{PS01} - p_{sat}(T_{05}) \tag{65}$$

A discrepancy large enough to be read off the manometer is therefore direct evidence of residual gas, and the two readings together provide a continuous check on the integrity of the evacuation throughout a run.
The expansion vessel, whose function in fixing the operating pressure is developed in Section 6.8, was fabricated in the laboratory from a stainless-steel tube closed by two welded end caps, and is shown in Figure 39b. It connects to the loop through the 10mm diameter copper pipe. A 10mm stainless-steel tube passes through the interior of the vessel, entering and leaving through the shell, and carries water from a dedicated temperature-controlled bath. Circulating water at a set temperature through this tube holds the vessel, and hence the liquid–vapor interface it contains, at that temperature. Because the interface is in equilibrium, the saturation pressure of the entire connected system follows, and the loop operates at the pressure corresponding to the bath setpoint rather than at a pressure determined by the heat balance. The vessel geometry and the routing of the internal tube are given in Figure 39a.
The exact pressure-control and vessel components, with suppliers and part numbers, are listed in Table 5 in Section 6.9.



(a) (b)

Figure 39: Expansion Tank

B.5 Instrumentation and Data Acquisition

Five Pt100 resistance thermometers are installed at fixed locations around the loop. Each is fitted into a stainless-steel thermowell entering through a dedicated brass T-piece, so that a sensor can be withdrawn without breaking the pressure boundary of the loop. TS01 and TS02 measure the cooling-water inlet and outlet temperatures at the condenser (Section B.3); TS03 measures the sub-cooled liquid leaving the condenser; and TS04 and TS05 measure the process fluid immediately before and after the heater block, giving the degree of subcooling at the heater inlet and the saturation or superheat condition at its outlet.

The elements are Class B, for which IEC 60751 specifies a tolerance of $\pm(0.3 + 0.005\,|T|)\,\mathrm{K}$, or $0.55\,\mathrm{K}$ at $50\,^{\circ}\mathrm{C}$. They are wired in a three-wire configuration, in which a third conductor allows the converter to measure and subtract the resistance of the leads themselves; without it the lead resistance would appear as an apparent temperature offset that varies with cable length and ambient conditions.

Each element is read by its own MAX31865 resistance-to-digital converter. The converter is ratiometric: it passes the same excitation current through the sensor and through an on-board reference resistor R_{ref} of nominal value $430\,\Omega$, and digitizes the ratio of the two as a 15-bit number, so that

$$R_{\mathrm{RTD}} = \frac{\mathrm{ADC}}{2^{15}}\,R_{\mathrm{ref}}. \tag{66}$$

The measurement therefore depends on the reference resistor rather than on any internal voltage reference, and the $0.1\,\%$ tolerance of that resistor propagates directly into the result, contributing between $0.28\,\mathrm{K}$ and $0.33\,\mathrm{K}$ over the range $20\,^{\circ}\mathrm{C}$ to $70\,^{\circ}\mathrm{C}$. Each converter board additionally carries a small gain error of its own. Both are systematic and repeatable, which is what makes them removable by calibration, but only so long as every element stays permanently paired with the same converter board, cable and chip-select line once calibrated.

Two converter settings are used throughout. The internal notch filter is set to $50\,\mathrm{Hz}$ to reject mains-frequency interference picked up on the sensor leads, and each reported value is the mean of eight consecutive conversions, which reduces random noise by a factor of $\sqrt{8}$ at the cost of measurement time.

The resistance returned by Equation (66) is converted to temperature on the host computer using the Callendar–Van Dusen relation as standardized in IEC 60751,

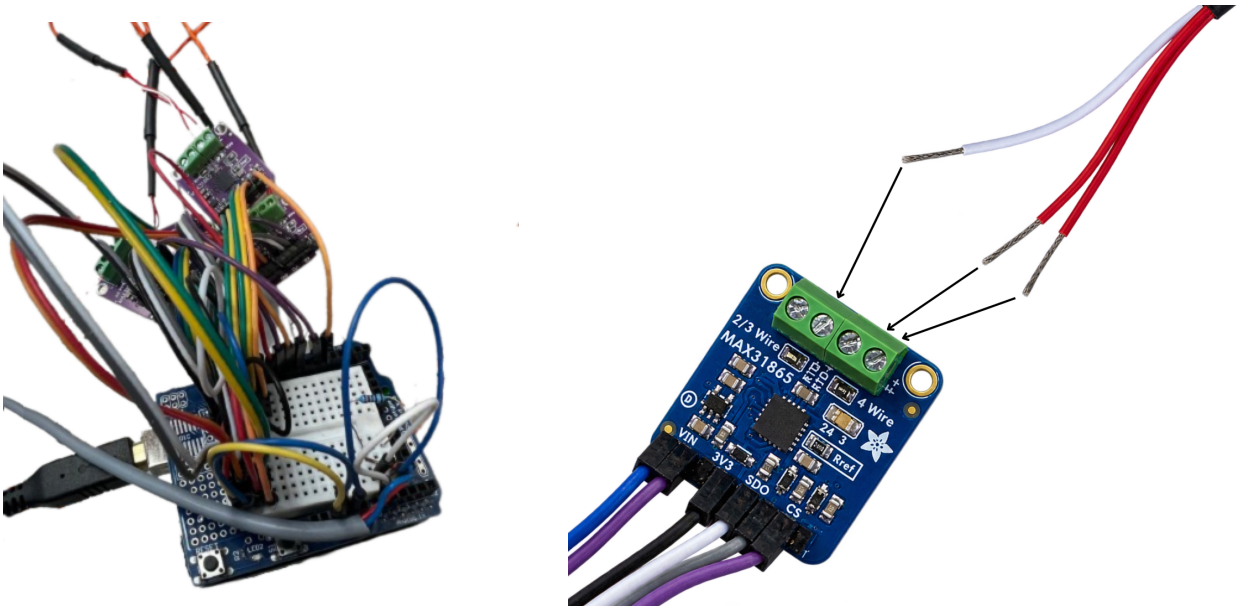
$$R(T) = R_0\left(1 + AT + BT^2\right), \qquad T \geq 0\,^{\circ}\mathrm{C}, \tag{67}$$

with $R_0 = 100\,\Omega$, $A = 3.9083 \times 10^{-3}\,\mathrm{K}^{-1}$ and $B = -5.775 \times 10^{-7}\,\mathrm{K}^{-2}$. Since the loop never operates below $0\,^{\circ}\mathrm{C}$ the third coefficient of the full standard is never active, and the inverse of Equation (67) reduces to the closed-form quadratic solution. The resulting temperature is then corrected by the per-sensor calibration polynomial derived in Section 7.1, where the complete uncertainty budget for this chain is also given.

The five converter boards are read by a single Arduino Uno over a shared SPI bus, shown in Figure 40. The clock, data-in and data-out lines are common to all five boards, as are the supply and ground connections; each board is distinguished only by its own chip-select line routed to a separate digital pin, so that the microcontroller addresses one converter at a time while the remainder stay electrically inactive on the bus. The boards are configured for three-wire operation by their on-board jumper.

One further temperature measurement is made outside this chain. A Type K thermocouple is inserted between the aluminum heater block and its surrounding insulation, and is read on a separate handheld display rather than through the Arduino, so its value is not logged automatically and is recorded by hand when required. It is not used in any derived quantity; its purpose is protective, as described in Section B.2, since a rise in block temperature is the first indication that the heated wall has dried out.

The circulating mass flow rate of the loop is measured with a Krohne AF-E 400 electromagnetic flow meter installed on the bottom pipe, immediately downstream of TS03, where the fluid is single-phase liquid. The meter is shown in Figure 41. It is powered from its own external supply unit and its current output is scaled



(a) Arduino Uno.

(b) MAX31865 resistance-to-digital converter board.

Figure 40: Data acquisition hardware

so that the $4\,\mathrm{mA}$ to $20\,\mathrm{mA}$ span corresponds to $0\,\mathrm{L\,min^{-1}}$ to $5\,\mathrm{L\,min^{-1}}$.
The meter carries its own display, which shows the instantaneous flow rate and is used for direct observation at the setup, but the value is also logged. Because the Arduino has no current input, the output loop is closed through a $220\,\Omega$ shunt resistor and the voltage across that resistor is read on analogue input A0; at full scale this develops $4.4\,\mathrm{V}$, which remains inside the $0\,\mathrm{V}$ to $5\,\mathrm{V}$ input range of the converter. The current is recovered from the measured voltage by Ohm's law and then converted to a flow rate using the configured span.
This arrangement sets the resolution of the logged flow signal. The Arduino's converter is 10-bit against a $5\,\mathrm{V}$ reference, so one count corresponds to $4.89\,\mathrm{mV}$, or $0.022\,\mathrm{mA}$ across the shunt, which is equivalent to approximately $7\,\mathrm{mL\,min^{-1}}$. This is a fixed resolution independent of flow rate, and is therefore a larger relative uncertainty at the low circulation rates characteristic of steady operation than at high ones.



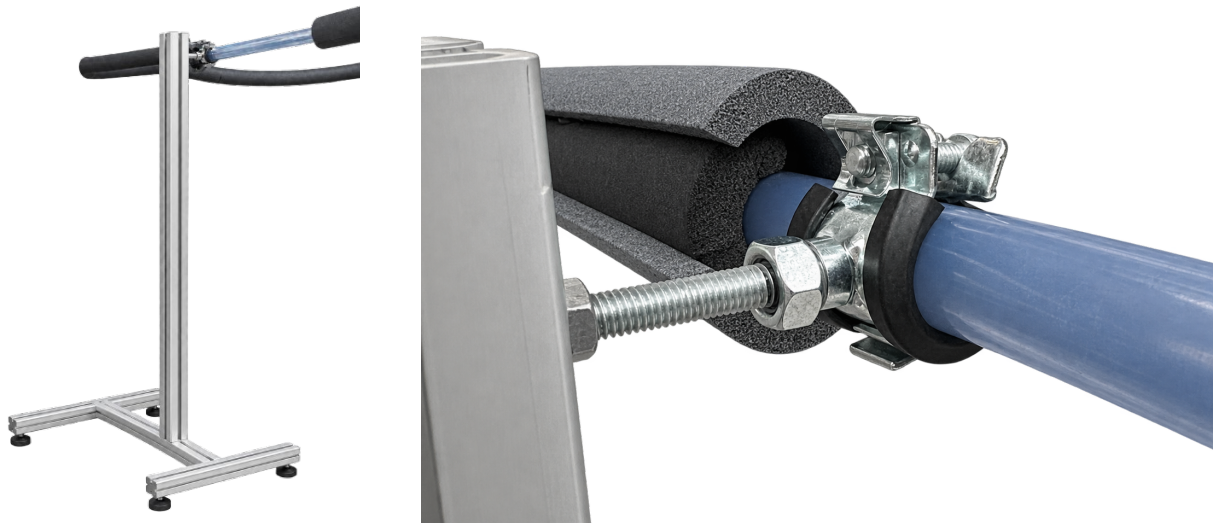
Figure 41: Krohne AF-E 400 electromagnetic flow meter installed on the bottom pipe.

The flow indicated by an electromagnetic meter is instantaneous and, in this loop, strongly intermittent. The acquisition software therefore records not only the instantaneous value but also running means over $10\,\mathrm{s}$, $30\,\mathrm{s}$, $1\,\mathrm{min}$ and $5\,\mathrm{min}$, computed from the recorded timestamps rather than from a fixed sample count. The

five-minute mean is used in all derived quantities.
The sampling interval of the system as a whole is set by the sensor chain rather than by the software. Each conversion on the MAX31865 requires of the order of 75 ms including bias settling, and with eight conversions averaged across five sensors a complete cycle takes approximately 3 s. The host requests a reading, the Arduino returns the five resistances and the flow meter voltage in a single line over the serial connection, and the software performs the conversion, correction and derived calculations before writing the record.
The location of each sensor on the loop itself can be seen in figure 20, while the exact instrumentation components, with suppliers and part numbers can be seen in Table 6 in Section 6.9.

B.6 Mechanical Structure and Mounting

The loop is supported on a frame of mounting profiles, with pipe clamps securing each straight section so that the pipe runs can be aligned with the MRI bore and adjusted during assembly. The mounting profiles and adjustable supports were drawn from existing lab stock. The pipe clamps are listed in Table 7 in Section 6.9. The main components of the mounting system are shown in Figure 42, with the aluminum profile structure shown in Figure 42a and the pipe-clamp arrangement in Figure 42b.



(a) Aluminum mounting profile.

(b) Pipe clamp securing the loop.

Figure 42: Mechanical support components used for mounting the thermosyphon loop.

B.7 Safety Systems

The main safety measure on the setup is the Type K thermocouple between the aluminum block and its insulation, described in Appendix B.2. The operator watches its reading throughout each run and switches the heater off by hand if the block temperature approaches 300 °C, which keeps it well below the 500 °C rating of the rock wool. During the dryout tests of Section 8.2 the block reached 330 °C before the heater was switched off, which shows the limits of manual monitoring, and an automatic cutout is therefore recommended in Chapter 10.
The low operating pressure also helps. Because the loop runs below atmospheric pressure, a leak draws air into the system instead of releasing hot water or vapor.

References

[1] J.G. Collier and J.R. Thome. *Convective Boiling and Condensation*. Clarendon Press, 1994.

[2] Tao Ding, Han Wen Cao, Zhi Guang He, and Zhen Li. Visualization experiment on boiling heat transfer and flow characteristics in separated heat pipe system. *Experimental Thermal and Fluid Science*, 91:423–431, February 2018.

[3] M.P Dudukovic. Opaque multiphase flows: experiments and modeling. *Experimental Thermal and Fluid Science*, 26(6-7):747–761, August 2002.

[4] Christian Poelma. Measurement in opaque flows: a review of measurement techniques for dispersed multiphase flows. *Acta Mechanica*, 231(6):2089–2111, June 2020.

[5] Maitreyi Sangal, Maria Anikeeva, Frank Godenschweger, Jan-Bernd Hövener, and Oliver Speck. Comparative Study of Gas and Liquid Flow in Packed Beds Using Phase-Contrast MRI. *Chemie Ingenieur Technik*, page e70154, August 2026.

[6] J.N Reyes, Jr, A.Y Lafi, and D Saloner. The use of MRI to quantify multi-phase flow patterns and transitions: an application to horizontal slug flow. *Nuclear Engineering and Design*, 184(2-3):213–228, August 1998.

[7] Emanuele Zanetti. Introduction to two-phase gas-liquid flow and heat transfer. Lecture slides, Equipment for Heat and Mass Transfer (ME45165), 2025.

[8] Lynn F. Gladden. Recent Advances in MRI Studies of Chemical Reactors: Ultrafast Imaging of Multiphase Flows. *Topics in Catalysis*, 24(1-4):19–28, October 2003.

[9] Louay Abd Al-Azez Mahdi and Ahmed Yusef Lateif. Prediction of Heat Transfer Coefficient and Pressure Drop in Wire Heat Exchanger Working with R-134a and R-600a. *Journal of Engineering*, 25(11):1–20, November 2019.

[10] J.M. Mandhane, G.A. Gregory, and K. Aziz. A flow pattern map for gas—liquid flow in horizontal pipes. *International Journal of Multiphase Flow*, 1(4):537–553, October 1974.

[11] Ran Kong and Seungjin Kim. Characterization of horizontal air–water two-phase flow. *Nuclear Engineering and Design*, 312:266–276, February 2017.

[12] Afshin J. Ghajar and Clement C. Tang. Heat Transfer Measurements, Flow Pattern Maps, and Flow Visualization for Non-Boiling Two-Phase Flow in Horizontal and Slightly Inclined Pipe. *Heat Transfer Engineering*, 28(6):525–540, June 2007.

[13] Yemada Taitel and A. E. Dukler. A model for predicting flow regime transitions in horizontal and near horizontal gas-liquid flow. *AIChE Journal*, 22(1):47–55, January 1976.

[14] John R Thome and Andrea Cioncolini. Two-Phase Flow Pattern Maps for Macrochannels. pages 5–45. October 2015.

[15] Sven Scharnowski and Christian J. Kähler. Particle image velocimetry - Classical operating rules from today's perspective. *Optics and Lasers in Engineering*, 135:106185, December 2020.

[16] Alberto Aliseda and Theodore J. Heindel. X-Ray Flow Visualization in Multiphase Flows. *Annual Review of Fluid Mechanics*, 53(1):543–567, January 2021.

[17] Min Chan Kim, Sin Kim, Heon Ju Lee, Yoon Joon Lee, Kyung Youn Kim, and Samim Anghaie. An experimental study of electrical impedance tomography for the two-phase flow visualization. *International Communications in Heat and Mass Transfer*, 29(2):193–202, February 2002.

[18] Jian-Nan Liu, Bo-Xiong Wang, Yuan-Yuan Cui, and Hao-Yuan Wang. Ultrasonic tomographic velocimeter for visualization of axial flow fields in pipes. *Flow Measurement and Instrumentation*, 41:57–66, March 2015.

[19] John M. Cimbala, Daniel E. Hughes, Samuel H. Levine, and Dhushy Sathianathan. Application of Neutron Radiography for Fluid Flow Visualization. *Nuclear Technology*, 81(3):435–445, 1988. _eprint: https://doi.org/10.13182/NT88-A16065.

[20] Daniel A. Clarke, Willian Hogendoorn, Alexander Penn, and Maria Raquel Serial. Magnetic resonance imaging in granular flows: An overview of recent advances. *Particuology*, 101:18–32, June 2025.

[21] M. Deimling, E. Mueller, G. Lenz, K. Barth, P. Fritschy, M. Seiderer, and E. R. Reinhardt. Description of Flow Phenomena in Magnetic Resonance Imaging. *Diagnostic Imaging in Clinical Medicine*, 55(1-2):37–51, 1986.

[22] R. R. Edelman, H. P. Mattle, J. Kleefield, and M. S. Silver. Quantification of Blood Flow with Dynamic MR Imaging and Presaturation Bolus Tracking. *Radiology*, 171(2):551–556, 1989.

[23] D. Saloner and C. M. Anderson. Flow Velocity Quantitation Using Inversion Tagging. *Magnetic Resonance in Medicine*, 16(2):269–279, 1990.

[24] Songhua Chen, Fangfang Qin, Kyung-Hoe Kim, and A. Ted Watson. NMR imaging of multiphase flow in porous media. *AIChE Journal*, 39(6):925–934, 1993. _eprint: https://aiche.onlinelibrary.wiley.com/doi/pdf/10.1002/aic.690390602.

[25] Katsumi Kose. RAPID COMMUNICATION: NMR imaging of turbulent structure in a transitional pipe flow. *Journal of Physics D*, 1990.

[26] Katsumi Kose. One-shot velocity mapping using multiple spin-echo EPI and its application to turbulent flow. *Journal of Magnetic Resonance (1969)*, 92(3):631–635, May 1991.

[27] J Blackwell, M J Kraśny, A O'Brien, K Ashkan, J Galligan, and N Colgan. Proton Resonance Shift Thermometry: A review of modern clinical practices.

[28] Sikai Zou, Quan Zhang, Chang Yue, and Sheng Du. Effect of servers' arrangement on the performance of a loop thermosyphon system used in data center. *Applied Thermal Engineering*, 192:116955, June 2021.

[29] A. Jamar, Z.A.A. Majid, W.H. Azmi, M. Norhafana, and A.A. Razak. A review of water heating system for solar energy applications. *International Communications in Heat and Mass Transfer*, 76:178–187, August 2016.

[30] Hanwen Cao, Haibin Li, Tao Ding, Zhiguang He, and Zhen Li. Research on the refrigerant distribution and distinctive flow phenomena in a two-phase loop thermosyphon through a holistic visualization method. *International Journal of Refrigeration*, 125:64–74, May 2021.

[31] Yong Bai and Qiang Bai. Hydraulics. In *Subsea Engineering Handbook*, pages 349–399. Elsevier, 2010.

[32] Stuart W. Churchill. FRICTION-FACTOR EQUATION SPANS ALL FLUID-FLOW REGIMES. *Chemical Engineering (New York)*, 84(24):91 – 92, 1977. Type: Article.

[33] Gherhardt Ribatski, Leszek Wojtan, and John R. Thome. An analysis of experimental data and prediction methods for two-phase frictional pressure drop and flow boiling heat transfer in micro-scale channels. *Experimental Thermal and Fluid Science*, 31(1):1–19, October 2006.

[34] S.Z Rouhani and E Axelsson. Calculation of void volume fraction in the subcooled and quality boiling regions. *International Journal of Heat and Mass Transfer*, 13(2):383–393, February 1970.

[35] B. Chexal, G. Lellouche, J. Horowitz, and J. Healzer. A void fraction correlation for generalized applications. *Progress in Nuclear Energy*, 27(4):255–295, January 1992.

[36] Aleksander Paliwoda. Generalized method of pressure drop calculation across pipe components containing two-phase flow of refrigerants. *International Journal of Refrigeration*, 15(2):119–125, January 1992.

[37] T. L. Bergman and Frank P. Incropera, editors. *Fundamentals of heat and mass transfer*. Wiley, Hoboken, NJ, 7th ed edition, 2011.

[38] Rubin Battino, Timothy R. Rettich, and Toshihiro Tominaga. The Solubility of Nitrogen and Air in Liquids. *Journal of Physical and Chemical Reference Data*, 13(2):563–600, April 1984.

[39] Ian H. Bell, Jorrit Wronski, Sylvain Quoilin, and Vincent Lemort. Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library coolprop. *Industrial & Engineering Chemistry Research*, 53(6):2498–2508, 2014.

[40] Jiang He, Guiping Lin, Lizhan Bai, Jianyin Miao, and Hongxing Zhang. Effect of non-condensable gas on the operation of a loop heat pipe. *International Journal of Heat and Mass Transfer*, 70:449–462, March 2014.

[41] Binglin Song, Guoying Meng, Wei Huang, Aiming Wang, Xiaohan Cheng, and Jie Yang. Experimental Investigation on Heat Transfer in Two-Phase Closed Thermosyphon Containing Non-Condensable Gas. *Energies*, 17(18):4693, September 2024.

Statement on the Use of AI Tools

AI tools were used in a supportive role during the preparation of this assignment, both for the writing and for the supporting code, but never as a substitute for my own thinking, analysis, or implementation.

Writing: AI tools were used to help me assess the overall structure of the report, to skim candidate papers and decide whether they were relevant to the research topic, to organize the points I wanted to make before writing, and in the final stages to suggest grammar and phrasing improvements. They were not used to generate academic content, arguments, analysis or conclusions, all of which are my own. Anything an AI tool provided was checked against the original academic sources and corrected where necessary, and no text was copied directly from AI output into the final submission.

Code development: AI tools were also used to support the development of the MATLAB code, mainly for code structure, debugging and the plotting routines, which is the part I am least comfortable with. The choice of equations, the logic of the steady-state model and the construction of the flow pattern map were all decided by me, and AI was used only to turn those choices into working code. Every suggestion was read, tested and checked against the underlying equations and against reference results from the literature before being used, and was changed wherever it was wrong or did not do what the script required.