Imperial College London

DMT 04: Biofuel Internal Combustion Engine

Heat Recovery and Phase Change Storage

Subgroup 04B:

Isaac Blanc

Jessica Parker

Alexander Schmidt

Julie Leclercq

Matthew Willis

Pages: 60

Date: 8 June 2022

Project director: Antonis Sergis

Project supervisors: Antonis Sergis and Alex Taylor

Executive summary

This Testing and Analysis report follows the later stages of the Design, Make and Test project of Group 04B. It covers the manufacture and testing of a heat storage device containing phase change material, which was created in order to recover waste heat from a Honda CBR600f4i 2001 engine and use it to preheat the engine and cabin.

The device contained a matrix of temperature sensors which aimed to predict the state-of-charge, and had a series of valves which were intended to control the distribution of heat inside the phase change material. Four major tests were conducted to assess the performance of the device, which were a charging, discharging, storage and state-of-charge sensing test.

The key results from testing were that the heat exchanger outperforms the PDS requirements in relation to storage capacity, charge time and discharge time. The PCM unit stored the required capacity of 1250kJ in 145 seconds, and delivered 1250kJ in 100-150 seconds, outperforming the required time of 30 minutes and 5 minutes respectively. A maximum of 5.5MJ was stored, significantly higher than the required capacity. The unit retained 35% of its energy after 24 hours, with 1750kJ remaining. In order for 75% of energy to be retained over 48 hours, 0.012m of insulation over the entire system is needed.

CFD simulations of phase change behaviour illustrated hysteresis effects, resulting in different heat transfer magnitudes when comparing solidification to melting. These simulations were validated by experimental results, which also showed that adding graphite powder to phase change material improves its melting behaviour and has no adverse affects on solidification.

Other than additional insulation, possible redesign and future work included enhancing the State of Charge system, changing some materials to reduce weight, and significantly reducing the size of the assembly.

Contents

Executive summary **i**

Contents **ii**

Nomenclature **iii**

1 Project Overview **1**
1.1 Super-Project Overview 1
1.2 Sub-Assembly Overview 1
1.3 System Diagram 4

2 Manufacture and Assembly Review **5**
2.1 Manufacture 5
2.2 Assembly 7

3 Preliminary Testing **12**
3.1 Graphite testing 12
3.2 CFD simulation 19

4 Main Tests **27**
4.1 Charge Test 27
4.2 Discharge Test 33
4.3 Storage Test 35
4.4 Localised Heating Test 39
4.5 Size and weight test 41
4.6 Implication of Results 41

5 Redesign **42**
5.1 Insulation 42
5.2 Heat Exchanger Water-Sealing 44
5.3 State of Charge System 44
5.4 Pipe Plates 45
5.5 Manifold Plates 46
5.6 PCM Container 47

6 Discussion **49**
6.1 Design Review 49
6.2 Future Progression 49
6.3 Team Progression 51

7 Conclusion **54**

References **56**

Nomenclature

A	Surface Area [m^2]
c_p	Specific Heat at Constant Pressure [$Jkg^{-1}K^{-1}$]
h	Heat Transfer Coefficient [$Wm^{-2}K^{-1}$]
K	Thermal Conductivity [$Wm^{-1}K^{-1}$]
L	Thickness [m]
$\dot{m}$	Mass Flowrate [kgs^{-1}]
P	Power [W]
S	Standard Deviation
T	Time [s]
T	Temperature [°C]
U	Overall Heat Transfer Coefficient [$WK^{-1}m^{-2}$]
V	Volume [m^{-3}]

Δ	Total Change
ρ	Density [kgm^{-3}]
ϕ	Concentration by Mass [%weight]
θ	Non Dimensional Temperature

eff	Effective
env	Environment
g	Graphite
i	Initial
In	At the Inlet
ins	Insulation
m	Metal
max	Maximum
min	Minimum
o	At the Outlet
p	At a Point
pcm	Phase Change Material
∞	Final

ABS	Acrylonitrile Butadiene Styrene (Plastic)
CFD	Computational Fluid Dynamics
CNC	Computer Numerical Control
EOL	End of Life
FDM	Fused Deposition Modelling
HSS	High Speed Steel
MIG	Metal Inert Gas
PCM	Phase Change Material
PEEK	Polyether Ether Ketone
SS	Stainless Steel
SoC	State of Charge
STW	Student Teaching Workshop
TIG	Tungsten Inert Gas

Project Overview 1

1.1 Super-Project Overview

The Biofuel Internal Combustion Engine Group aimed to improve the sustainability of a Honda CBR600f4i 2001 engine by modifying it to run on E85 biofuel (Figure 1.1). This super-project comprised of 3 subgroups: inlet, control, and exhaust and heat recovery, who agreed to focus on characteristics relevant for a consumer car. The first subgroup focused on the air intake and fuel injection systems. The second subgroup designed a new engine control unit. Finally, the third subgroup developed a novel system for recovering and storing waste heat - the subject of this report.

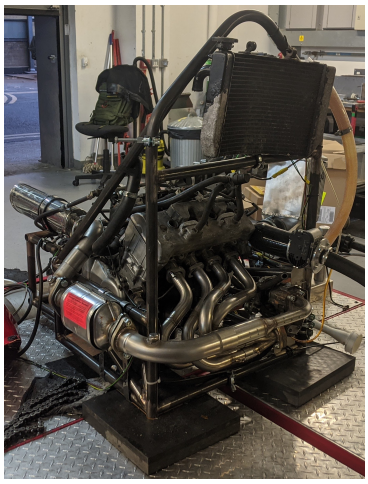

Figure 1.1: The Honda CBR600f4i 2001 engine that was the focus of the super-project.

In addition, the three subgroups worked together to troubleshoot the engine, which would not start. The problem was isolated, and the findings were documented for future groups [1]. This will not be the focus of this report.

1.2 Sub-Assembly Overview

Subgroup 04B was tasked with developing a system that extracts waste heat from the engine and stores it using a phase change material (PCM). This can then be used to pre-heat the engine and cabin.

At the heart of this project, lie phase change materials. These absorb large amounts of heat during melting (charging) and release it again during solidification (discharging), both of which are isothermal. Moreover, latent heat storage offers higher energy density than

sensible heat storage [2]. While any phase transition can be used for latent heat storage, the solid-liquid transition was chosen due to its small volumetric expansion (compared to liquid-vapor) and moderate-to-high latent heat of fusion (compared to solid-solid) [2].

However, there are still limitations and unknowns with PCMs. These materials present low thermal conductivities and hysteresis effects, where phase change varies with respect to time and direction (i.e. melting vs solidifying) [2]. These affect the rates of charging and discharging, limiting the performance of the heat storage system. This was one of the challenges the group was tasked with overcoming.

Furthermore, methods for state of charge (SoC) estimation for latent heat storage are still immature, with little published literature [3, 4]. The obvious approach for estimating the energy stored in an object by sensible heat is to measure its temperature, but this approach is less effective for latent heat, since phase change is isothermal. The subgroup recognised the importance of SoC estimation if the device were to be part of a wider system, such as a car. Therefore, this was another key design challenge.

After thoroughly understanding the challenges, objectives and market landscape, the subgroup created a Product Design Specification detailing the major requirements. Table 1.1 contains a condensed version focusing on the requirements tested and analysed in this report; it should be noted that all other requirements were met during the various development gateways.

Table 1.1: Summary of Product Design Specification

Capacity	≥ 1250kJ
Charge time	≤ 22 mins
Discharge time	≤ 5 mins
Heat retention	≥ 75% over 48 hours
Size	≤ 1×1×0.15m
Mass	≤ 15kg

Figure 1.2: The PCM unit in a partially assembled state, showing the insulation, container and heat exchanger.

The final product is an integrated heat storage and heat exchange system, shown in Figure 1.2. It consists of a heat exchanger inside a container filled with CrodaTherm-74 PCM, henceforth referred to as PCM. The heat exchanger is a network of several independent loops through which engine coolant is directed by valves, shown in Figure 1.3, to transfer heat to and from the PCM. Although originally intended to extract heat from the exhaust system too, it was decided that this was beyond the resources of this year's project; it does remain an open opportunity for future work.

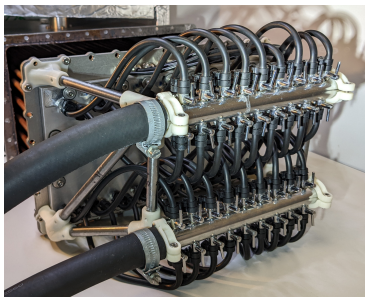

Figure 1.3: The 18-valve inlet and outlet splitters, which are used to direct coolant flow to specific heat transfer loops within the PCM unit.

Software written by the subgroup uses readings from eight temperature sensors to interpolate the 3D temperature field in the container, which it visualises (Figure 1.4) and uses to estimate SoC. This software, combined with the use of valves and independent loops in the heat exchanger, enables one to direct the flow of heat transfer fluid to specific regions of PCM, based on their temperature, in order to accelerate charging/discharging or maximise storage capacity. This would solve the problematic non-uniformity, low thermal conductivities, and hysteresis effects mentioned above. It was intended that this functionality would be automated in future iterations of the design, if it were found to be effective.

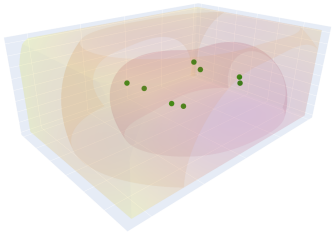

Figure 1.4: 3D temperature field visualisation produced by software written by the subgroup.

In addition, Subgroup 04B played a leading role in troubleshooting the engine, in particular by removing and replacing rusted exhaust studs and isolating an electrical fault in the fuel cut relay. However, it was agreed across subgroups that integration with the engine was not feasible in the given timeframe.

1.3 System Diagram

Table 1.2 provides the names of each key part and assembly. Figure 1.5 shows their locations.

Table 1.2: Names and descriptions of each part and assembly discussed in this report.

No.	Part Name	No.	Part Name
1	Front Manifold Plate	8	Purge Valve
2	Rear Manifold Plate	9	Splitters
3	Front Pipe Plate	10	Temperature Sensor Mounrs
4	Rear Pipe Plate	A1	Heat Exchanger Assembly
5	Copper Pipes	A2	Metal Container
6	Manifold Supports	A3	Splitter Support Structure
7	Fill Port		

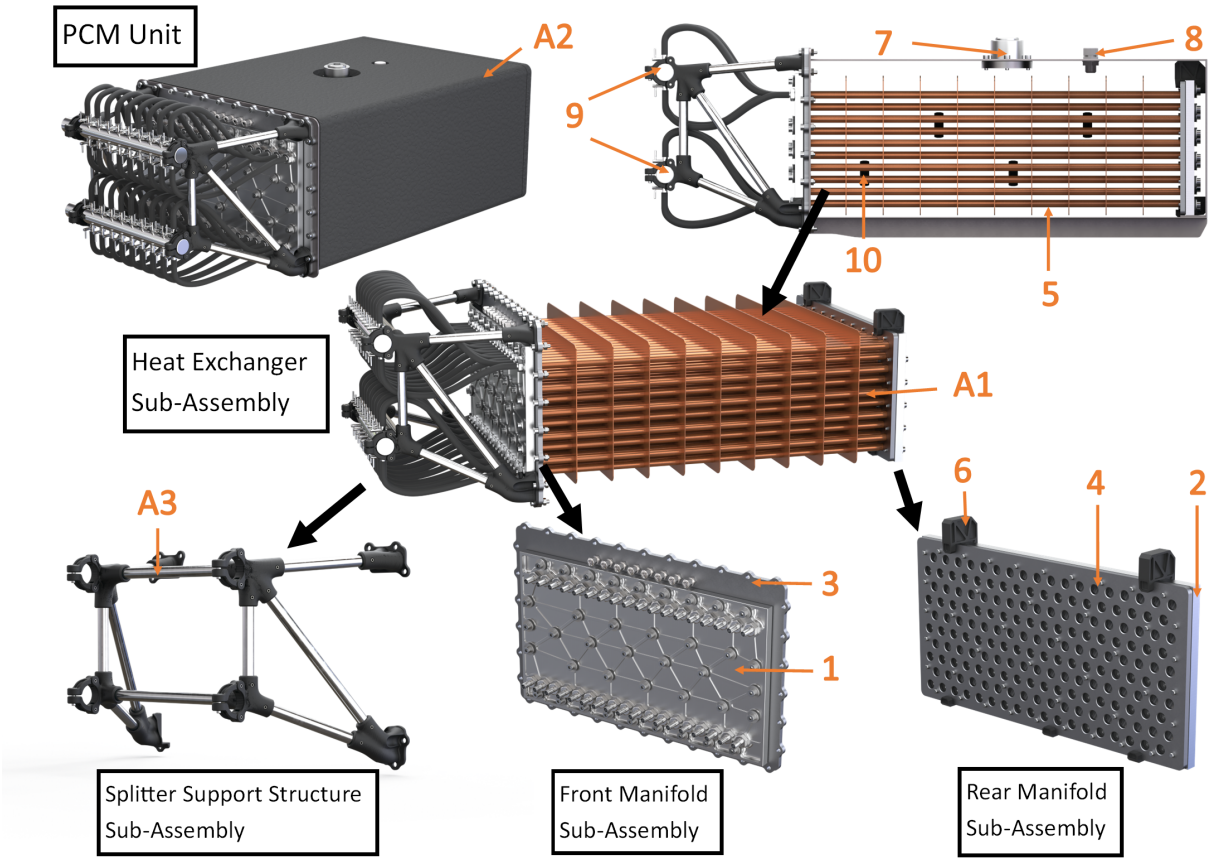

Figure 1.5: The sub-assemblies of the PCM unit

Manufacture and Assembly Review 2

2.1 Manufacture

2.1.1 Manifold plates

Early thermofluids calculations found that the best heat exchanger configuration to be a large number of small diameter copper pipes. This presented a manufacturing challenge, since such a design required hundreds of tight-radius return bends. As shown in Figure 2.1, these could not be fabricated using a standard pipe bender or hand bending, nor could affordable fittings be sourced online. The manifold plate was designed to resolve this.

Figure 2.1: The unsuccessful result of an early pipe bending test.

The components were CNC-milled by the subgroup. Return bend slots were milled while maintaining flatness, as shown in Figure 2.2. However, an axis direction error during setup meant that the weight-saving ribs, shown in Figure 2.3, were misaligned. Consequently, the stepped geometry around the sealing washers' holes would prevent them from working properly. This was remedied by face-milling the ribs away, which had no effect on the plate's function or fit, indicating that the ribs were unnecessary. Overall, the manifold plates were machined and finished to a high dimensional and cosmetic standard, resulting in a good fit with their interfacing components.

Figure 2.2: The rear of the manifold plates, showing the return-bend slots.

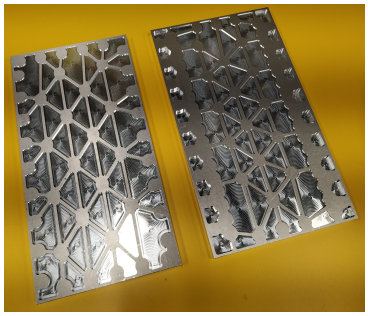

Figure 2.3: The front of the manifold plates, showing the misaligned weight-saving ribs.

2.1.2 Pipe plates

The main purpose of the pipe plates was to ensure alignment of the copper pipes to the manifold plate.

Laser cutting was chosen, since the plate required a large number of accurately positioned holes. Unfortunately, both pipe plates underwent thermal warping due to heat concentrations from the laser. This was likely caused by the high melting point of stainless steel that, the 5mm thickness, and the close proximity of the holes.However, re-programming the laser cutter to allow cooling time between adjacent cuts did not yield significant improvements.

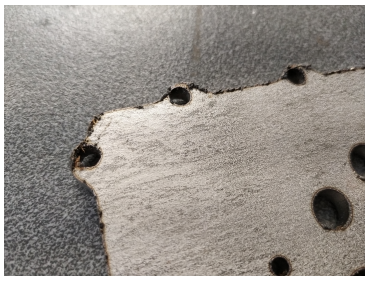

Figure 2.4: Poor results from laser cutting

The rear pipe plate did not warp significantly so the hole positions remained in tolerance. Therefore, 3-point bending was used to flatten the plate enough to form a seal with the rear manifold plate. The front pipe plate warped to a much greater extent, such that holes were badly misaligned shown in Figure 2.4, necessitating the use of a different manufacturing method. Waterjet cutting was chosen because, while it is less precise than laser cutting and exhibits tapering, it does not suffer from thermal issues. This achieved a successful result, shown in Figure 2.5, with the plate remaining flat and the holes being accurately positioned. It was concluded that waterjet cutting should have been the first-choice method.

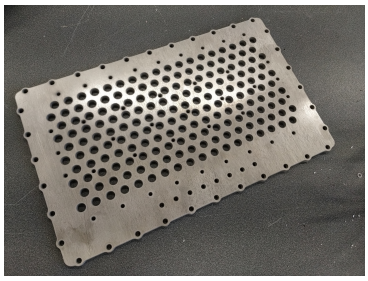

Figure 2.5: The front pipe plate, having been water-jet cut after the first attempt with laser cutting led to warping.

Finally, the holes were chamfered, as shown in Figure 2.6, to guide the flow of solder into each joint when the pipes were brazed to the plates.

Figure 2.6: Chamfered holes on the pipe plates, to guide the flow of solder during brazing.

2.1.3 Gaskets

To seal the manifold subassemblies, the original design used O-ring thread and O-ring grooves in the manifold plates. However, due to the warping of the rear pipe plate, it was thought that fluid might short-

circuit between the return slots, which the O-ring was not designed to prevent. To resolve this, a 1mm thick silicone gasket was laser cut, shown in Figure 2.11. This would seal each return slot individually to prevent short-circuiting.

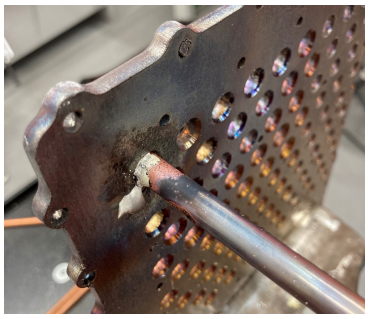

Figure 2.7: The unsuccessful result of a practise run at brazing the copper pipes to the pipe plates.

2.2 Assembly

2.2.1 Heat exchanger

The original plan had been to braze the copper pipes to the pipe plates. However, practice attempts with spare material proved unsuccessful, as shown in Figure 2.7. While copper to stainless steel brazing is possible, the wrong brazing flux had been purchased. In addition, the recent experience with laser cutting led to concerns about the pipe plates warping again, especially with the correct flux's high brazing temperature approaching 1000°C. Furthermore, due to the close proximity of the pipes, brazing one joint may have melted the solder in adjacent joints. It was suggested that the pipes be assembled with solder and flux wrapped around each joint and then heated in a kiln to braze all simultaneously. While this may have solved the issue of pipe proximity, there was still a substantial risk of warping and concern about the feasibility of arranging a new complex assembly procedure late in the manufacturing phase.

Figure 2.8: The plywood assembly jig holding the pipe plates in place while the copper pipes were glued into place.

The group decided to seal the copper pipes to the plates using high temperature epoxy resin[1]. A plywood assembly jig was built to hold the pipes and pipe plates in alignment during gluing, shown in Figure 2.8. The plan was to pour epoxy onto the pipe plates, con-

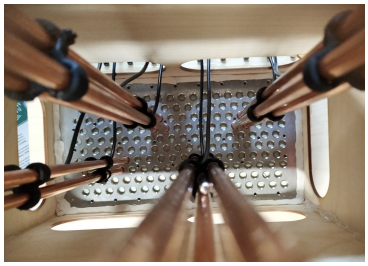

Figure 2.9: The temperature sensors mounted to the copper pipes held by the assembly jig. These would be used by the SoC system.

taining it between plasticine dams. However, upon mixing the epoxy, it was found to be too viscous.

Rather than placing all the remaining copper pipes into position, the group lathered the front pipe plate with epoxy. Then each pipe was lathered with epoxy at one end and slid through the pipe plate, coating the inside of the hole with epoxy. Pipes with temperature sensors mounted to them, shown in Figure 2.9 were inserted first. Overall, this proved to be extremely time consuming, because of the large quantity of pipes and the challenge of getting the pipes into the correct holes. After 48 hours curing time, the result was a stiff structure, shown in Figure 2.10.

Excess epoxy and proud pipes on the rear plate were sanded down with an electric sander. Once this had ensured the surface was smooth, a gasket and manifold plate was fitted to each pipe plate, shown in Figure 2.11.

Then followed a series of iterative leak tests. In the first leak test, pumping cold water through the heat exchanger revealed significant leaking from the rear pipe plate. Upon disassembly, it was evident that it had been caused by the electric sanding, which had removed some of the epoxy from the countersinks. Epoxy was re-applied and re-sanded, but this time with greater care to avoid excess and without the use of power tools.

In the second leak test, the rear pipe plate no longer leaked from around the gasket but water was seen emerging from under the screws on both manifold plates. As a result, screws were fully tightened, and

1: Proset ADV 170 epoxy resin was used with Proset M 2052 Hardener

Figure 2.10: All of the copper pipers glued between the pipe plates, with the jig having been removed after curing.

Figure 2.11: The front manifold plate being attached to the front pipe plate, with the silicone gasket in between.

Loctite was applied between the sealing washer and the screw. Loctite was used in this manner to facilitate the removal of the entire screw and washer assembly if necessary and because water was leaking from the top of the washer rather than the side.

In the third leak test, hot water was pumped through the heat exchanger and a little was seen leaking from the jubilee clips on the front manifold plate. Following discussion with technicians, these were replaced with nylon hose clamps, since they apply more even force to the hose. The success of this solution is shown in Figure 2.12.

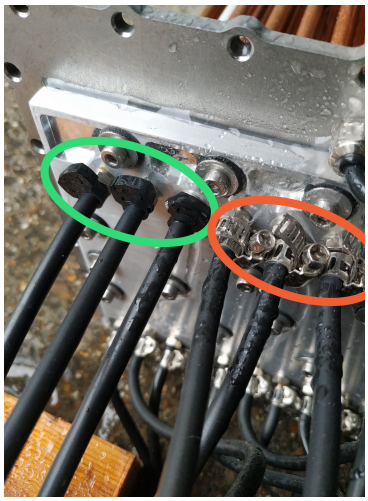

Figure 2.12: The third leak test, where nylon hose clamps (circled in green) were found to perform better than jubilee clips (circled in red).

In the fourth leak test, the heat exchanger was found to be leak-free, indicating a successful assembly of the heat exchanger. However, shown in Figure 2.13, small patches of water were seen on the pipe-side of the rear pipe plate hours later, suggesting there were still some very slow leaks. It was concluded that these leaks were from epoxy failure on the pipe-side, since the other side had already been re-epoxied. Since these epoxy joints were unreachable and the leak was very slow, no attempt was made to fix it.

Figure 2.13: Slow leakage in rear manifold sub-assembly.

2.2.2 PCM container

The PCM container was manufactured from 2mm mild steel sheet, which was laser cut and sanded to remove oxidation. The sheets were mostly TIG welded together, however severe problems with weld contamination were encountered, shown in Figure 2.14. This was not merely the result of an inexperienced student welder; even the technician commented that they struggled to weld it. This may have been caused by

a chemical reaction during the high temperatures of laser cutting. As a result of this, the final seams were MIG welded, since it is less sensitive to contamination issues. The subassembly was rigid but filling it with water revealed minor leaks.

Figure 2.14: A badly contaminated TIG weld, which caused the PCM container to leak.

In response, the internal edges were sealed with epoxy. After curing, another leak test found no leaks, shown in Figure 2.15. However, when test fitted with the heat exchanger, the epoxy prevented full insertion. The subgroup resolved this by chiselling away excess epoxy. It was also found that the flange plate had warped during welding. To rectify this, the holes in this part were redrilled, enabling assembly with the heat exchanger.

Figure 2.15: The PCM container, having been painted and sealed with epoxy, filled with water for a leak test.

2.2.3 Splitter support structure

Initially, the support joints were intended to be printed through MJF printing, using nylon, as this would yield good temperature resistance, high strength, and good dimensional accuracy. However, in the interest of cost and time, FDM printing with ABS was chosen instead. These components, which had complex geometry and overhangs, were printed to a reasonable quality thanks to the dissolvable support material. The only minor cosmetic flaws were some build support material left within the cavities and some pin holes being slightly too small which were resolved using a hand drill.

The rods were cut from 10mm mild steel rod. Following datum errors on the mill, spring pin holes did not align so epoxy was used instead. The resulting structure, shown in Figure 2.16 was rigid but heavy.

Figure 2.16: The splitter support structure, before the splitters were attached.

2.2.4 Summary

While assembly was successful, several aspects would be done differently in future. The subgroup would commit to using epoxy instead of brazing, enabling the pipe plates to be made from aluminium. This would save money on outsourcing by enabling the use of laser cutting, without risking warping, and reduce the tapping time, since the metal is softer. The lack of warping would, in turn, negate the need for laser cut gaskets. The greatest short falling of the assembly was that the rear pipe plate leaked slightly even after four iterative test and improvement cycles. In future, this could be solved by using a less viscous (more pourable) epoxy to ensure thorough application among the pipes. Finally, using a guillotine instead of a laser cutter for the PCM container could have avoided the weld contamination issues. Also, the decision to use MIG rather than TIG welding should have been made sooner after the problem was discovered. This would have prevented leaking and eliminated the need for sealing.

Preliminary Testing 3

3.1 Graphite testing

Although phase change materials have a high latent heat capacity, their performance is limited by their low thermal conductivities, which prevent heat from being dissipated evenly through the material, causing sub-optimal heat transfer. Graphite powder has a very high thermal conductivity of 25-470 W/mK [5], in comparison to the PCM which only has a thermal conductivity of 0.27 W/mK while liquid and 0.29 W/mK while solid [6]. In theory, adding graphite powder to the PCM should increase its thermal conductivity and result in a higher heat transfer rate from conduction. This theory is also supported by the literature [7].

However, there were concerns that increasing the graphite concentration too high would increase the viscosity of the PCM so that there is less heat transfer via convection, or that the graphite would settle to the bottom of the PCM and become useless or even reduce its performance. Therefore, an experiment was conducted to calculate the optimal concentration of graphite powder to maximise heat transfer. Additionally, it has been reported in the literature [8] that, since graphite does not have latent heat within the operating temperature range, adding it to the PCM decreases the latent heat of the mixture.

The expected thermal conductivities of the PCM-graphite mixture with varying graphite concentrations are summarised in Table 3.1. These values were calculated using Maxwell's exact solution for the conductivity of randomly distributed and non-interacting homogeneous spheres in a homogeneous medium, Equation (3.1) [9].

$$k_{\text{eff}} = k_{\text{pcm}} \frac{k_{\text{g}} + 2k_{\text{pcm}} + 2\varphi_{\text{g}}(k_{\text{g}} + k_{\text{pcm}})}{k_{\text{g}} + 2k_{\text{pcm}} - \varphi_{\text{g}}(k_{\text{g}} - k_{\text{pcm}})} \tag{3.1}$$

This model is only accurate when the graphite concentration by weight (φ_{g}) is low and the particles are not touching. It also does not account for the uneven shapes

and sizes of the graphite particles, so is only an estimate used to illustrate the above theory approximately. The thermal conductivity of graphite powder, k_g , was taken to be 25 W/mK which is the lower bound. This value was chosen because the graphite had not been aligned to increase its conductivity [10]. The thermal conductivity of the PCM, k_{pcm},

was taken to be 0.28 W/mK, which is the average of the solid and liquid thermal conductivities. As shown in Table 3.1, increasing the graphite concentration can significantly improve the thermal conductivity of the PCM. This should increase the heat transfer via conduction.

Table 3.1: Expected thermal conductivities of PCM-graphite mixtures.

φ_g (% wt)	k_{eff} (W/mK)
0	0.28
2	0.30
4	0.31
6	0.33

3.1.1 Stability test

A stability test was conducted to assess whether the graphite powder remains homogeneous in the PCM during the liquid and solid phases. Graphite powder was combined with liquid PCM and left on a heated plate for 15 minutes. The graphite powder remained suspended within the PCM for this length of time and showed no signs of settling or inhomogeneity, as shown in Figure 3.1. It is possible that the convection currents caused by the hot plate kept the graphite suspended.

Figure 3.1: The PCM-graphite mixture when liquid.

The PCM-graphite mixture was then left to fully solidify, which took 1 hour and 40 minutes. Some of the graphite remained suspended in the PCM, but a significant amount settled to the bottom of the container, as shown in Figure 3.2. This was a very slow solidification process, since there was no active cooling like there would be in the PCM unit (i.e. cold water flowing

through pipes). Therefore, it is likely that there would be significantly less accumulation of graphite powder at the bottom of the PCM in the actual unit, since there would be less time for it to settle before the PCM is fully solidified.

Figure 3.2: The PCM-graphite mixture when solid, showing the accumulation of graphite at the bottom of the container.

3.1.2 Mixture phase change test

A test was conducted to determine what concentration of graphite powder to put in our PCM in order to maximise heat transfer. A testing rig was manufactured, which was comprised of a copper pipe of the same size as the pipes used in the actual assembly, which was placed inside an aluminium box with the same internal width and height as the pipe spacing in the full assembly. The rig was fitted with k-type thermocouples, shown in Figure 3.4, and insulated to reduce heat loss during testing. It was connected to a boiler and two pumps in parallel, as shown in Figure 3.4, so that hot and cold water could be run through it to melt and solidify the PCM.

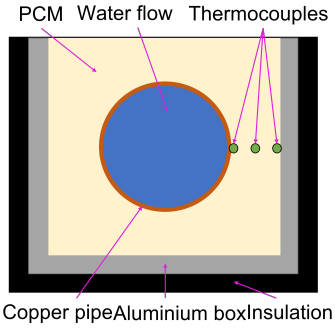

Figure 3.3: Diagram of the three thermocouple positions on the copper pipe, the aluminium wall and centred in between.

Samples of PCM with graphite concentrations of 0%, 2%, 4% and 6% wt were poured into a testing rig and hot and cold water was run through in order to assess the melting and solidification behaviour for each graphite concentration.

Due to inconsistencies in the test, the temperature data was non-dimensionalised. For example, the boiler was not able maintain a perfectly constant temperature, so the flow temperature T_∞ varied between tests. Additionally, the ambient conditions varied, meaning that the PCM initial temperature, T_i, also varied between tests. The temperature readings were non-dimensionalised using Equation (3.2), where T is the recorded temperature and θ is the non-dimensional temperature.

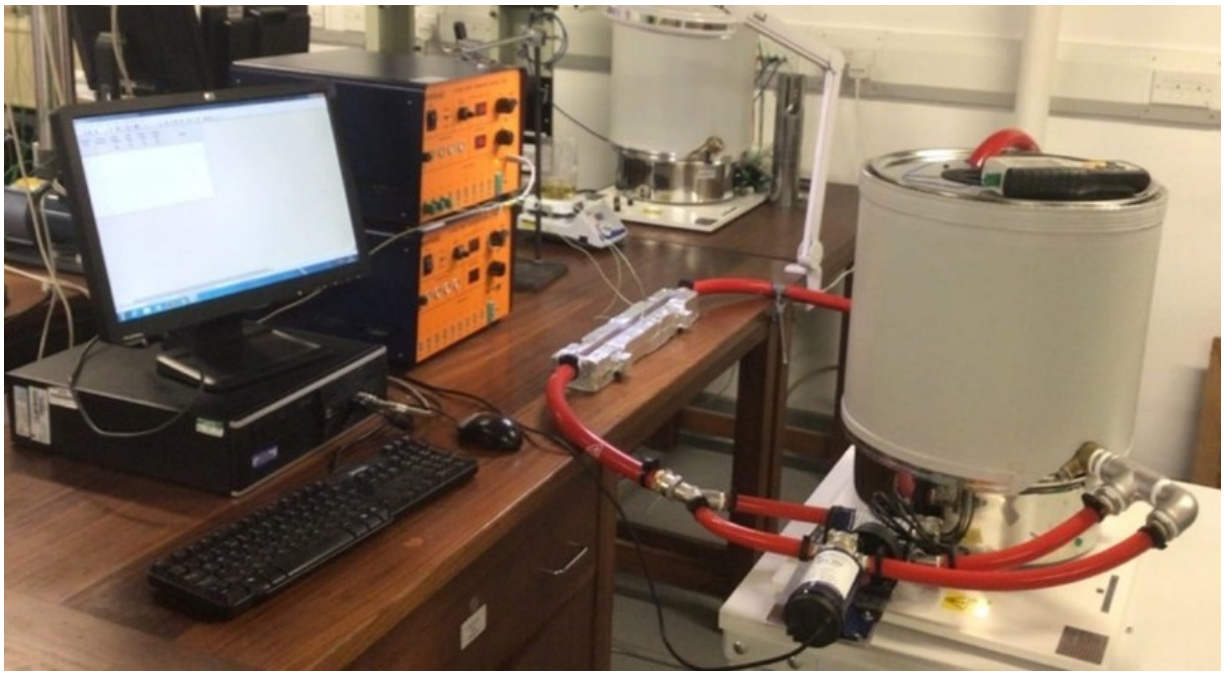

Figure 3.4: The apparatus set up for the graphite-PCM mixture phase change test.

$$\theta = \frac{T - T_\infty}{T_i - T_\infty} \tag{3.2}$$

3.1.3 Error analysis

The main source of error was the ± 2.2°C uncertainty in k-type thermocouples. Errors were calculated by assigning worst case values of θ for the greatest underestimate possible, and the greatest overestimate. For example, the maximum overestimate increased the value of T by 2.2°C, and decreased the value of T_i by the same amount. These upper and lower bounds for error are represented by shaded regions on the graphs.

3.1.4 Results

For the 0% graphite melting test, the temperature did not follow the same trend as the other tests. It was believed that the thermocouple broke on this test, since the raw data shows that the temperature reading did not reach above 30°C. Therefore, the mid-point temperature reading was rejected from the data.

Figure 3.5 was obtained by smoothing the non-dimensionalised temperature data using a non-weighted 51 point moving average to reduce the noise, and then the

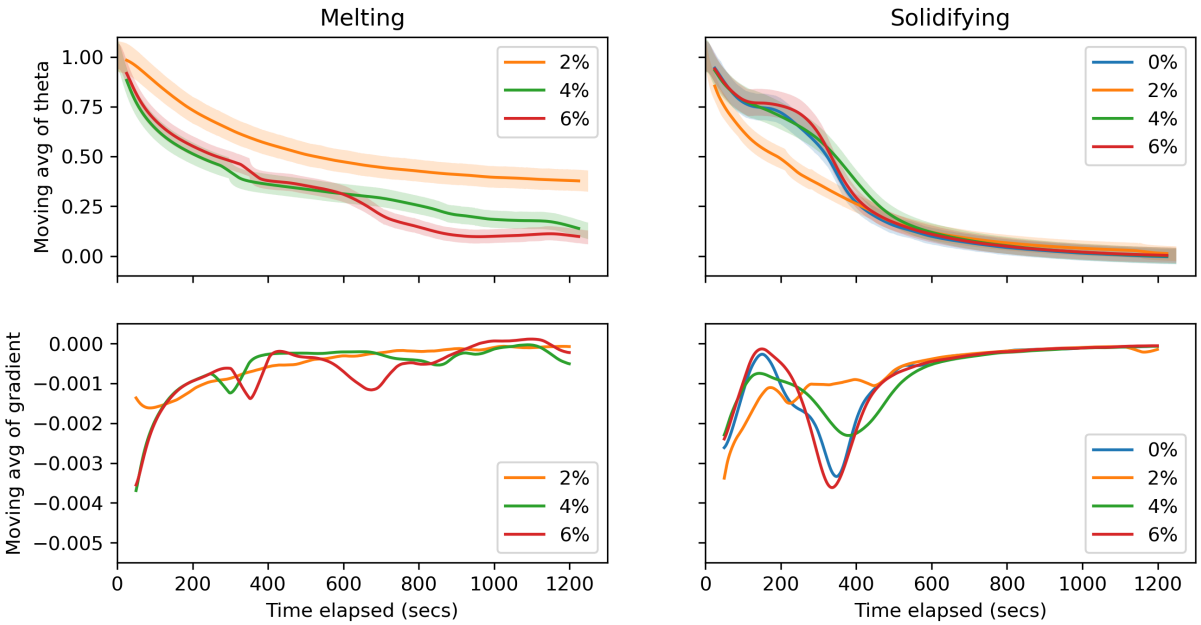

Figure 3.5: Moving averages of θ and $\frac{d\theta}{dt}$ against time for the mid thermocouples in every melting and solidifying test.

gradient was calculated using the second order central difference method. The mid thermocouple data was used for this, since the temperature change of the PCM was of greater interest, and the wall temperature data was too greatly affected by heat transfer to/from to the aluminium box.

The gradient was calculated to show trends in the rate of change of temperature during different stages with varying amounts of convection. For melting, it was expected that three difference mechanisms were going to affect the magnitude of the gradient. The first mechanism related to the amount of convection occurring. As the test progressed, it was expected that there would be an increasing gradient magnitude because more PCM would be liquid, allowing more convection and therefore more heat transfer. The second mechanism related to the temperature difference between the pipe and PCM. At the beginning of the test, the temperature difference is very large, resulting in a peak heat transfer rate. As the heat is transferred to the PCM, the PCM temperature gets closer to the pipe temperature, resulting in a decreasing heat transfer rate. This would reduce the gradient magnitude over time. The third mechanism related to the temperature difference between the PCM and the atmosphere. As the PCM temperature increases, this temperature difference increases also, resulting in larger heat losses to the atmo-

sphere. This would reduce the gradient magnitude because the PCM would heat up slower.

In this particular test, the mechanisms which reduce the gradient magnitude seem to balance with the mechanism which increases the gradient magnitude to give a reasonably constant gradient for the later part of the test. In the earlier part of the test, the temperature difference between the pipe and the PCM results in a very large reduction in the temperature gradient magnitude, which is expected.

For solidification, it is expected that the same mechanism related to the temperature difference between the pipe and PCM, would reduce the gradient magnitude over time. The mechanism related to the temperature difference between the PCM and the atmosphere would still decrease the gradient magnitude because as the PCM temperature reduces it gets closer to the atmospheric temperature, so heat losses reduce and the PCM cools slower. However, for solidification the convection mechanism has the opposite effect, since the amount of convection decreases over time so the gradient magnitude would also decrease over time. All three mechanisms therefore decrease the gradient magnitude, which corresponds to the final 700 seconds of Figure 3.5. However, the first 500 seconds show a sinusoidal shape, which is when the phase change occurred. Phase change happens at a constant temperature, so it makes sense for the rate of temperature change to decrease when the PCM reaches the phase change temperature and then increase again once it has changed phase. This trend should also be present in the melting results, but is less clear here.

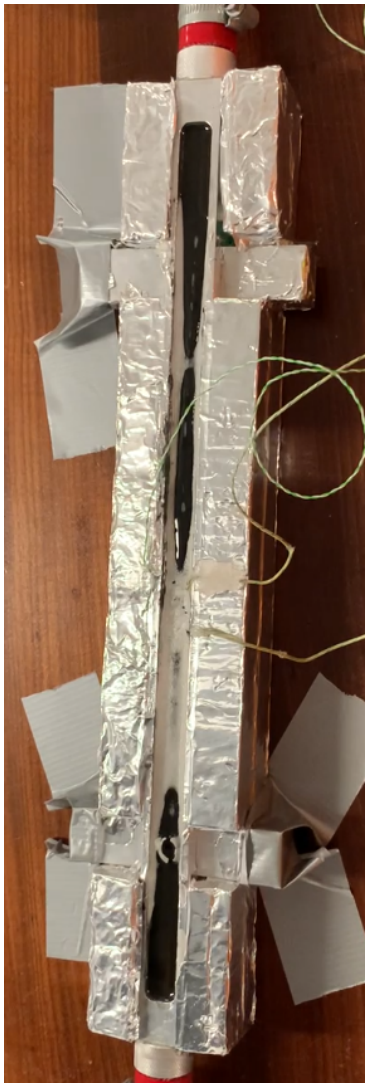

Figure 3.6: Midway through a melting test.

According to the theory presented earlier, one would expect to see the non-dimensionalised temperature for 6% graphite to decrease and stabilise the fastest. As shown in Fig-

ure 3.5, this is true for melting, but there is no clear trend for solidification. However, the graphite powder does not seem to adversely affect the heat transfer of the PCM for solidification, so it is still preferable to add it. Additionally, the literature supports this conclusion [11].

One explanation as to why there was no significant improvement in heat transfer for solidification was that the graphite powder may not have been mixed in properly to the PCM, because it was very difficult to do so properly in the test rig and it was too messy to keep pouring the PCM in and out of the rig to mix it in a beaker, since a significant amount of PCM/graphite would have been spilt and the concentrations would no longer be accurate. However, if a lack of mixing was the culprit, this should have also affected the melting results in the same way.

The decision was made to add 6% graphite, since it showed the best results for melting, no adverse results for solidification, and this conclusion was supported by the theory and literature [11]. A higher concentration was not added than this because it was not tested for, and Equation (3.1) is not valid for high concentrations. Additionally, there was concern about the graphite settling to the bottom of the container. However, it had previously been calculated that a maximum of 22.3% graphite could be added before the accumulated powder would touch the bottom pipes of the heat exchanger and prevent heat transfer via convection, so this was not a concern for the amount of graphite that was added.

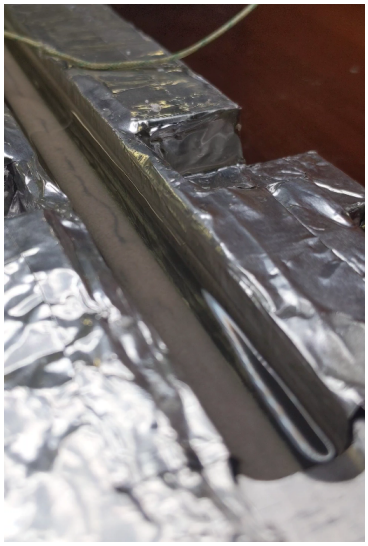

Figure 3.7: An early stage in a solidification test, where the PCM is solidifying around the copper pipe.

It was expected that the PCM would begin melting at the inlet and finish at the outlet, since the water temperature should be hotter at the inlet and cool down as it transfers heat to the PCM. As shown in Figure 3.6, the melting process starts at the ends of the test rig and propagates inwards towards the middle. This may be because the inlet and outlet of the test rig have abrupt changes in pipe diameter and is not a smooth transition, so will trigger turbulence and therefore increase the heat transfer rate in these areas. Additionally, the increased melt time at the centre may be

caused by the aluminium tape used to hold down the pipe thermocouple interfering with the heat transfer.

As shown in Figure 3.7, the PCM solidifies around the pipe as expected. However, it solidifies at the top surface before the underneath part has fully solidified, as shown in Figure 3.8. This is due to heat losses to the atmosphere at the top surface.

Figure 3.8: A later stage in the same solidification test, where the PCM is solidifying on the top surface, due to heat losses to the atmosphere.

3.2 CFD simulation

CFD analysis was performed using Ansys Fluent in order to examine the melting and solidification behaviour inside the PCM. If this simulation is validated by experimental results, it can then be used to provide a better understanding of what is happening inside the PCM unit in areas that cannot be observed visually.

A hexagonal domain with symmetry boundary conditions was chosen, with the height of the hexagon equal to the pipe spacing. The circular cut-out in the middle represents the pipe, and a constant temperature wall condition was put onto this boundary. For solidification, this was 15°C, and for melting it was 90°C. The thickness of this domain was 1mm, whereas the actual length of the pipe is 410mm.

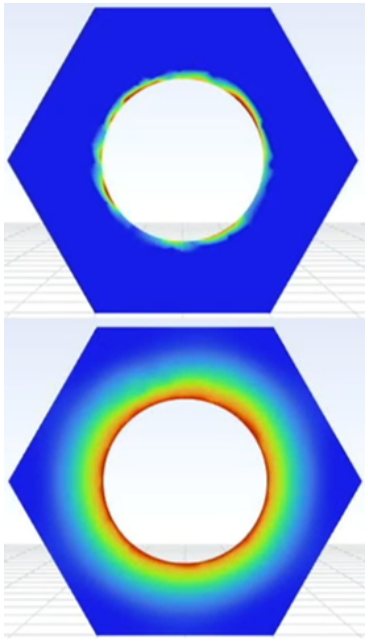

Figure 3.9: Liquid fraction (above) and temperature (below) plots across the CFD domain, 10 seconds into the melting simulation.

The heat transfer rate from the pipe was calculated in order to compare the heat transfer for melting versus solidification. As seen by comparing Figures 3.10 and 3.12, the heat transfer rate is lower for solidification than for melting for the latter part of the test. This is because, as the PCM solidifies around the pipe, convection is stopped so heat can only be transferred by conduction. Upon further inspection of the liquid

fraction contour plot over time, it is clear that as the PCM solidifies around the pipe, the heat transfer rate drops significantly. Once the pipe has been fully encased by solid PCM, the heat transfer rate continues to decrease but at a slower rate, since the thickness of the solid PCM is still increasing but slowly.

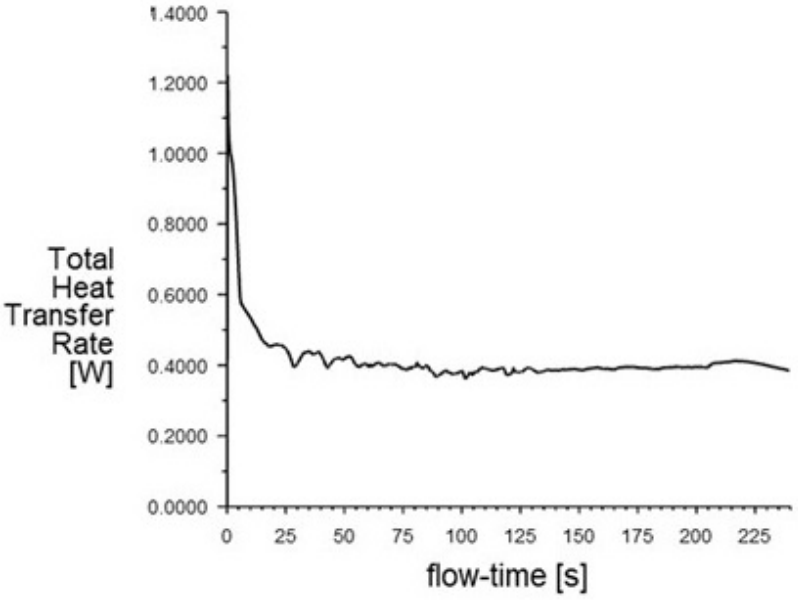

Figure 3.10: Heat transfer rate, predicted by the CFD, during melting of the PCM.

For melting, the heat transfer rate drops significantly in the first few seconds due to the large temperature difference between the pipe and PCM (65°C), but once this temperature difference decreases to 5°C, the heat transfer rate stays relatively constant at 0.4W. This is because for almost the entirety of the melting process, convection can occur around the pipe. In contrast, the heat transfer rate for solidification drops as low as 0.08W as the solidification front increases, decreasing the effect of convection.

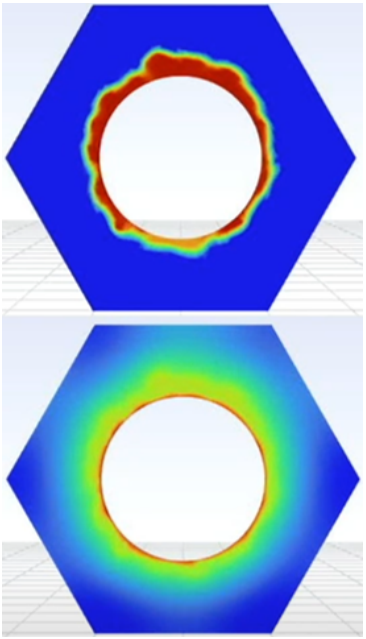

Figure 3.11: Liquid fraction (above) and temperature (below) plots across the CFD domain, 45 seconds into the melting simulation.

The heat transfer results scaled in order to calculate the heat transfer for whole system, as shown in Table 3.2. The discharging heat transfer rate is higher because these are the values at the beginning of the test, when the PCM being charged is solid and the PCM

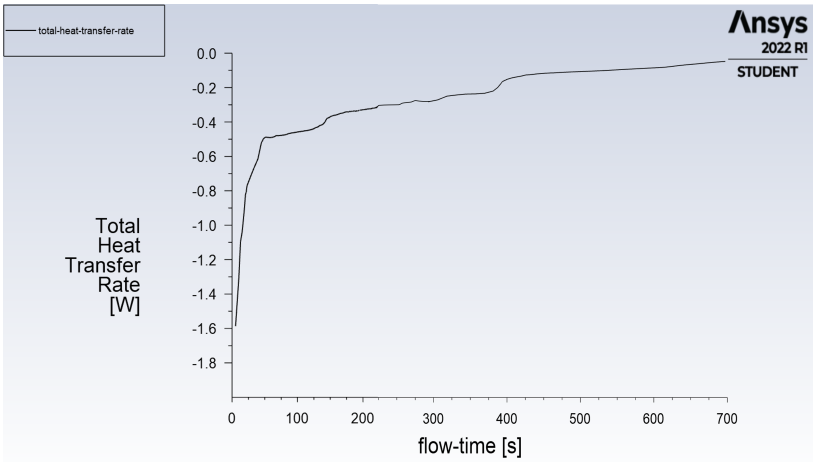

Figure 3.12: Heat transfer rate, predicted by the CFD, during solidification of the PCM.

being discharged is liquid, so there is more convection occurring in discharging. As mentioned previously, this lessens significantly as the test progresses. These variations in heat transfer rate for solidification versus melting are an example of the hysteresis effect, where the phase change behaviour in one direction is different to the other direction.

Table 3.2: Heat transfer results from CFD, scaled for the length and number of copper pipes.

	Domain (1mm) (W)	**Pipe** (410mm) (W)	**System** (180 pipes) (kW)
Charge			
Peak	1.00	410	74
Sustained	0.45	190	33
Discharge			
Peak	1.60	660	120
Sustained	0.45	210	37

As seen in Figure 3.13, the PCM liquefaction front has an oval shape, with more melting occurring at the top than the bottom due to convection. This asymmetric melting front means that the symmetry boundary condition imposed around the edge of the hexagon is no longer entirely accurate. A plane of symmetry along the top face would imply that the hexagon above the domain is flipped upside down, with convection and gravity acting in the opposite direction, which is clearly not

the case.

The result of this inaccurate condition is that the asymmetric nature of the melting front is exaggerated, since the hotter part in the upper half of the domain would be next to another hotter part in the reflected domain above, when in reality it would be next to a cooler part. This could be considered a ‘worse case scenario’ of uneven melting, and it can be assumed that in reality the melting front would be more symmetrical than what these CFD results show.

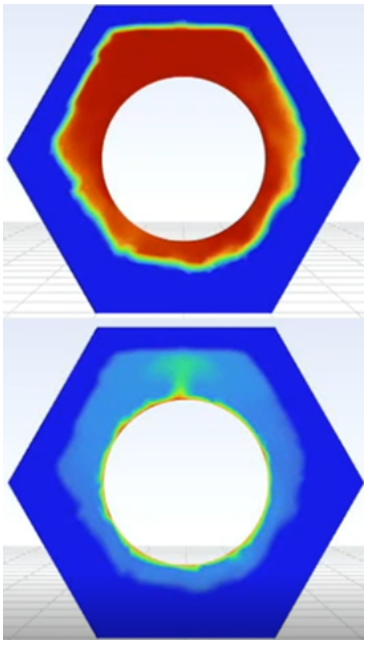

Figure 3.13: Liquid fraction (above) and temperature (below) plots across the CFD domain, 140 seconds into the melting simulation.

Instead, an ‘antisymmetry’ boundary condition would be more suitable, since the geometry above is symmetrical but the thermal loading is in the opposite direction. To our knowledge, this type of boundary condition was not available in the software used for this analysis, so an insulated wall could be used instead. However, an insulted wall would not accurately portray the situation either, as it would ignore the effect of a lower hexagon heating up the bottom of a hexagon above it.

The solidification front exhibits a similar asymmetric nature due to convection. The mushroom type shape of the solidification front in Figure 3.15 is caused by convection. As the warm PCM below the pipe rises, it heats up the solid PCM above it, limiting the thickness of the solidification front. Above the PCM, heat rises away from the solidification front, so the thickness is greater here. As the cooling continues, the warmer PCM rises to the top of the domain, meaning the bottom part begins to solidify first. A very interesting observation from these results is that the ring of solid

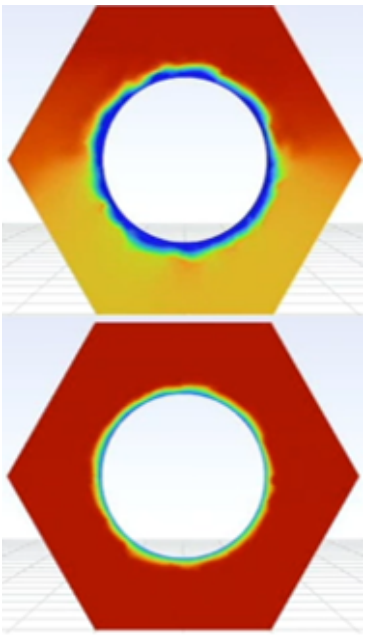

Figure 3.14: Liquid fraction (above) and temperature (below) plots across the CFD domain, 60 seconds into the solidification simulation.

PCM stops growing and the remaining liquid PCM all begins to change phase together as a whole, with a gradient liquid fraction from top to bottom.

The slight asymmetry along the z axis is likely caused by the mesh not being perfectly symmetrical. This could be improved by refining the mesh to be symmetrical about the z axis. In fact, only half the hexagon needs to be modelled in order to get the necessary results, which would reduce simulation time.

At the point that the melting front reached the edge of the domain, the melting front would protrude into other hexagons and PCM would flow between hexagons. At this point, bulk convection would begin to occur, which we would expect to increase the heat transfer rate and melting rate significantly. This CFD analysis using this small domain only provides useful information up to this point, as it is not able to simulate bulk convection without modelling the entire unit.

Similarly, when the solidification front reaches the edge of the domain, it would begin to protrude into other hexagons, which the simulation does not model, so the CFD only gives useful results up to this point.

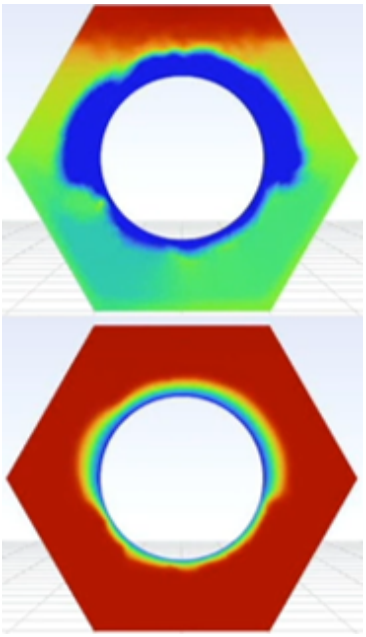

Figure 3.15: Liquid fraction (above) and temperature (below) plots across the CFD domain, 275 seconds into the solidification simulation.

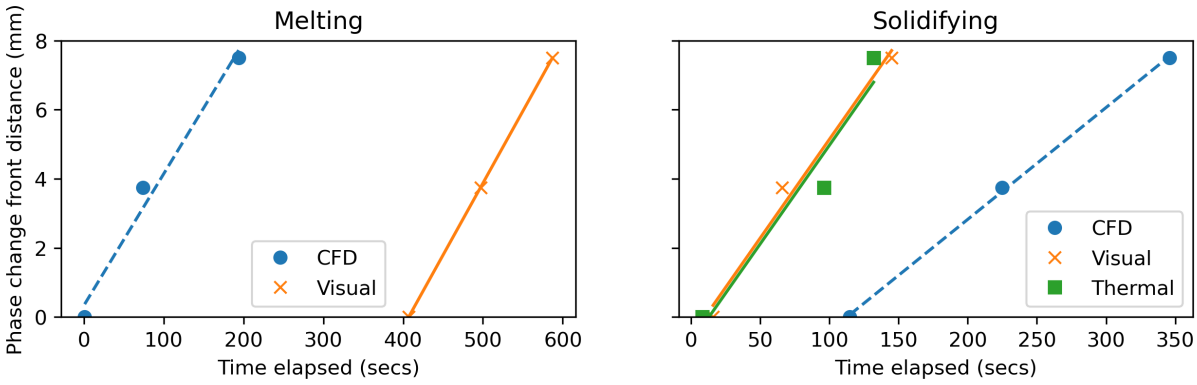

Figure 3.16: Phase change front location against time, as predicted by the CFD, compared to experimental results both from thermocouple measurements and visual inspection.

Figure 3.16 compares the phase change front movement results from the CFD to the 0% graphite experimental results. The legend labels refer to the following: 'CFD' means phase change time according to the CFD, 'Visual' means phase change time according to visual observations of the experiment, and 'Thermal' means phase change time according to thermocouple measurements (i.e. when temperature reaches 74°C). There was not enough data to make a trendline for the melting phase change time from thermocouple readings, since the mid-thermocouple broke, so this data was omitted from the graph.

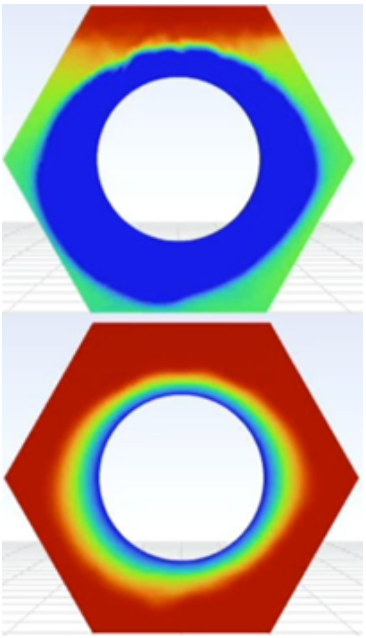

Figure 3.17: Liquid fraction (above) and temperature (below) plots across the CFD domain, 488 seconds into the solidification simulation.

The gradients of the CFD results are very close to the experimental results and the visual data matches the thermal data very well for solidifying. However, there are very significant offset errors for both solidification and melting. There are three possible reasons for this. Firstly, in the solidification test, there is a significant amount of heat lost to the atmosphere, whereas this is not the case in the CFD, so we would expect solidification to start earlier in the experiment. Secondly, the experimental results for melting are very error-prone because the opaque surface layer of solid PCM means we could not see what was happening at the pipe. The times were approximated by watching the section of PCM above the pipe darken. Similarly, it was very difficult to say when the PCM between the pipe and wall melts, because it happens before the surface has melted and turned transparent. Finally, the CFD starts with the pipe at 90 deg C, whereas the pipe has to heat up in the experiment. Granted, this happens pretty quickly, but it would still contribute slightly to the delay.

In Figure 3.18, the lines represent the experimental results and the dots represent the CFD results. Since the 0% graphite mid thermocouple broke during the experiment, it has not been included.

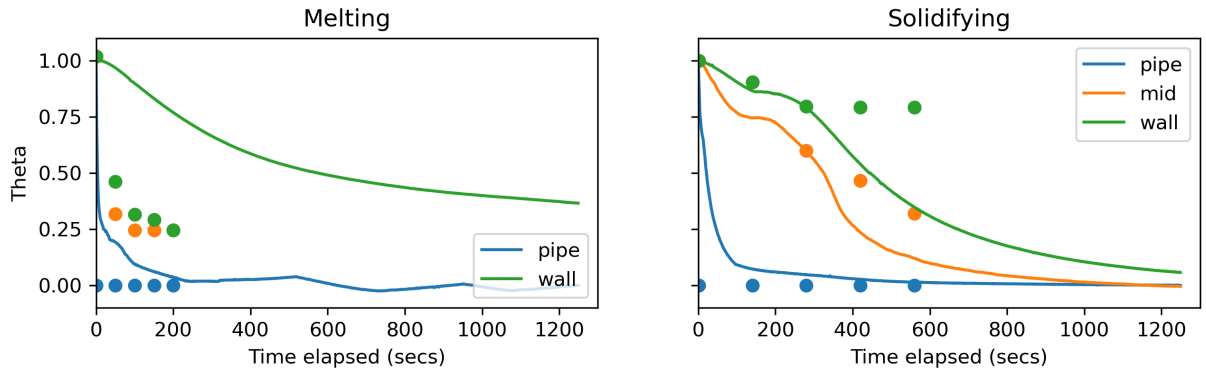

Figure 3.18: Comparison of the 0% graphite test temperatures with those predicted by the CFD, against time.

As shown in Figure 3.18, the melting process happens much quicker in the CFD compared to the experimental results. This is likely due to the significant heat losses to the atmosphere during the experiment, since the top of the test rig was uninsulated so that the melting process could be observed. Another factor leading to this discrepancy was that the value used for the thermal conductivity of the PCM in the CFD may not be accurate. The PCM supplier did not have a full dataset for the PCM properties, including its thermal conductivity, so they also provided the properties of a 'similar' PCM. The similar PCM is unlikely to have the exact same properties. Additionally, in the experiment the pipe takes about 200 seconds to fully heat up, which would delay the melting process, whereas in the CFD the pipe is hot immediately due to the prescribed conditions.

The CFD and experimental results for solidifying results match up much more accurately. However, they main difference is present in the wall temperature. For the experimental results, the wall temperature does not cool down to the extend it does in the CFD. This is likely because the wall is made of aluminium which acts as a heat sink, so it takes a significantly longer time to cool down than in the CFD. This also affects the experimental mid temperature and prevented it from cooling down as quickly as in the CFD. Similarly to the melting results, the pipe took about 100 seconds to cool in the experiment, but cooled immediately in the CFD, so this would have delayed the experimental solidification. However, the heat losses in the experiment would have spend up the solidification process.

Another interesting point is that in the graphite experiment, solidification took

significantly less time than melting, which is the opposite of what we would have expected and what the CFD shows. The CFD showed that heat transfer was slower for solidification than for melting due to the prevention of convection around the pipe. This is discrepancy is likely due to the heat losses to the atmosphere in the experiment, which spend up solidification and slowed down melting.

Main Tests 4

4.1 Charge Test

As specified by the PDS, the PCM unit must recover and store **1250kJ** of heat **within 22 minutes**, the length of an average UK car journey [12]. Furthermore, the SoC system must accurately **monitor the energy stored** and the temperature distribution of the PCM. These three requirements were tested as follows.

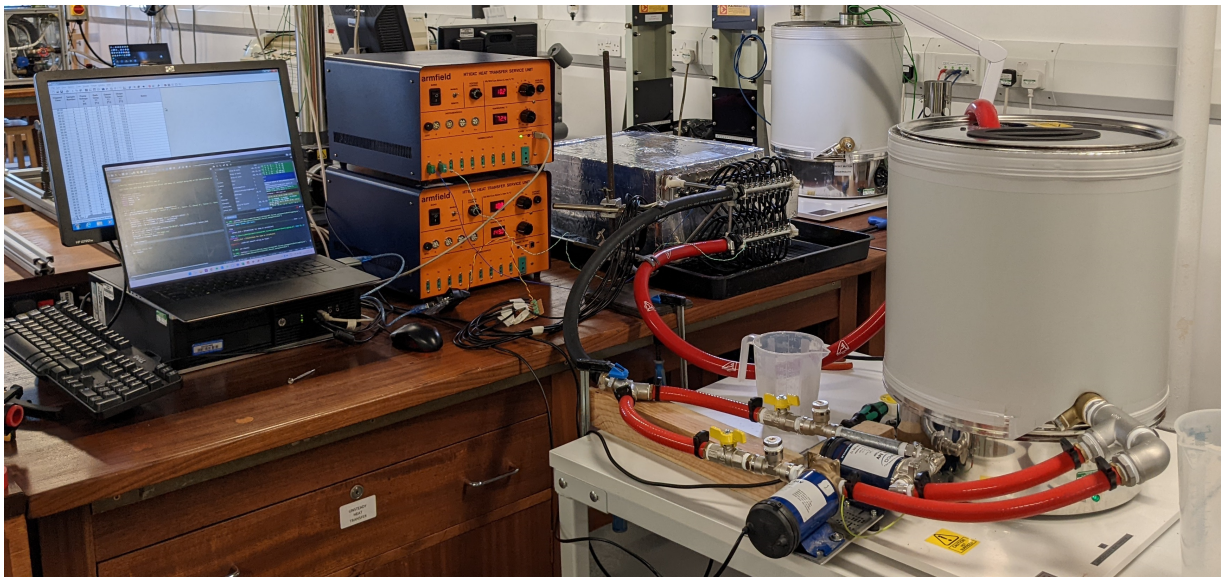

Figure 4.1: The apparatus used for the charge, discharge, storage and localised heating tests, all discussed in this section.

The PCM unit was connected a boiler and two pumps in parallel, as shown in Figure 4.1. Water was heated in the boiler to approximately 95°C in order to simulate engine coolant and then pumped through the PCM unit before recirculating. Meanwhile, water temperatures at the inlet and outlet of the PCM unit were recorded at a sampling rate of 1 Hz by thermocouples and an Armfield HT10X Heat Transfer Service Unit. Using this data, the heat flux, or power P, into the PCM unit could be calculated using Equation (4.1).

$$P = \dot{m} c_p \Delta T \tag{4.1}$$

where $\dot{m}$ is the mass flow rate, c_p is the specific heat at constant pressure of water, and ΔT is the difference in temperature between the inlet and outlet. It was expected that the heat flux would eventually stabilise to a constant value, representing heat loss from the PCM unit. At this point, it would no longer be charging

so the test would be stopped. When post-processing the data, the power would be numerically integrated over time to calculate the total energy stored in the PCM unit. Henceforth, this calculation is referred to as the SFEE (steady flow energy equation) estimate of stored energy.

In early analysis and design, we had intended for a coolant flowrate of 16kg/min, based on the expected flowrate of the engine coolant pump. However, when it became apparent that the engine would not run in time for testing, we were unable to afford a suitably powerful hot water pump. Therefore, we ran the two borrowed pumps at as high a power as possible: 20W each at 10V, which provided 5kg/min. This was determined by lifting the outlet hose from the boiler (while the water was cold) and measuring the time taken, t, for a 2 litre jug to be filled with water. Measurements were repeated five times in order to quantify error. Then Equation (4.2) was used:

$$\dot{m} = \frac{\rho V}{t} \tag{4.2}$$

Where ρ is the density of water, assumed to be 1kg/litre, and V is the volume of water collected, which was 2 litres. All joints were secured with nylon hose clamps to prevent leaking. This was important from a safety perspective as well as to ensure that the flowrate measurements were accurate.

4.1.1 Error analysis

As in the graphite testing discussed previously, the main source of experimental error was due to the uncertainty of the K-type thermocouples, of 2.2°C. Additionally, the uncertainties in the power depend also on the uncertainty in the measurement of the volumetric flowrate. We have assumed no error in the values used for the density and specific heat at constant pressure of the water.

As done previously, the graphs of inlet and outlet temperature over time contain a shaded region, which represents the maximum and minimum possible Temperature readings. In order to translate this to the contribution of this uncertainty to the uncertainty in power, a maximum and minimum temperature difference was calculated with Equations (4.3) and (4.4):

$$\Delta T_{\text{max}} = (T_{\text{inlet}} + 2.2) - (T_{\text{outlet}} - 2.2) \tag{4.3}$$

$$\Delta T_{\text{min}} = (T_{\text{inlet}} - 2.2) - (T_{\text{outlet}} + 2.2) \tag{4.4}$$

The uncertainty in the measurement of the volumetric flowrate was calculated simply by calculating the standard deviation of the 5 measurements that were taken. These were then used to generate similar upper and lower bounds for volumetric flowrates. These were all added together for each value of power, to produce a shaded area to visualise the uncertainty in the calculated power.

To calculate the error in the total energy transfer, a slightly more complicated method was required. While it was possible to simply integrate the upper and lower bound power curves, this would result in a very large upper and lower bound for energy, as the errors would accumulate over the range of the integration. Instead, the errors were calculated by calculating the variance of the integral – this was possible because the integral was evaluated using a trapezoidal approximation, which is a sum. The variance of a weighted sum of uncorrelated variables, X and Y, is given by Equation (4.5):

$$\text{Var}(aX + bY) = a^2\text{Var}(X) + b^2\text{Var}(Y) \tag{4.5}$$

Where Var represent the Variance operator, and a and b are constants. Thus, the approximate variance for a trapezoidal integral approximation, s_I, is given by Equation 4.6.

$$s_I = \frac{\Delta t^2}{4}\left[s_0 + s_n + 4\sum_{i=1}^{n-1} s_i\right] \tag{4.6}$$

Where Δt is the time step of the interpolation, and s_0 to s_n are the variances at each point. In our calculations, the variance at each point was set to the total calculated uncertainty in power, and the uncertainty in energy stored was set as the standard deviation of the integral value, equal to the square root of s_I.

4.1.2 Expected results

During the design process, thermal circuit theory and flow correlations were used to model the temperature distribution along the copper pipes. By combining these results with the expected engine coolant pump flowrate and the specific heat of water, a **peak charging power of 10.8kW** was predicted. This was expected to occur at the start of the test, followed by a reduction in power as the temperature

difference between the water and PCM decreased. However, this process could not be predicted analytically, since it was thought the prevalence of natural convection would increase as the PCM melted, re-boosting heat transfer. This analysis required CFD.

As shown in Table 3.2, the CFD predicted a much greater **sustained charging power of 33kW**. This was expected to be an overestimate, since it modelled the pipe as isothermal. In contrast, the aforementioned analytic method accounted for the reduction in pipe temperature along its length, which was expected to occur as heat diffused from the coolant into the PCM. This CFD result predicted the PCM unit would **store 1250kJ in 38 seconds**. Containing 10.5kg of PCM, the unit was expected to have a **maximum capacity of 2.4MJ**. This was more than the 1250kJ target, since the exploratory nature of the project had necessitated a large safety factor.

The **SoC estimation was expected to be accurate before and after phase change, but not during**. It uses radial basis function (RBF) interpolation to generate a 173000 point 3D grid of temperatures using measurements from eight sensors at known locations in the PCM, which are subsequently used to estimate the SoC. This mathematical approach has not been developed to describe heat transfer in PCM. The decision to use it was largely driven by its ability to extrapolate as well as interpolate, since the sensors are located within the domain, and its availability in the SciPy Python library. SoC estimation for latent heat storage is an emerging field where temperature field interpolation is but one method currently being trialled [3]. With little published literature on this approach, we had low expectations.

4.1.3 Results

The fill port was left open during the test, to mitigate the risk of a pressure build up due to thermal expansion in the event that the purge valve failed. After 20 minutes, water was observed leaking from here at about 50mL/min. Previous leak tests indicated it came from the rear pipe plate; the only area with remaining leaks. After 30 minutes, liquid PCM began to emerge from the fill port, likely being displaced by the water. Once the PCM unit cooled, this leaking stopped.

The **peak power observed was 22.8 ± 1.2kW**; significantly higher than expected. However, at the beginning of the test, the pipes were filled with residual cold water. For the period between hot water reaching the inlet and reaching the outlet, the thermocouples recorded an extremely large temperature difference. The large apparent power calculated from this does not reflect heat flux from the water to the PCM, so it is unclear what the true peak power was. After the initial spike subsided at around 300 seconds, the system charged at 1.5kW. This gradually decreased, as expected and explained previously.

According to the SFEE estimation, the PCM unit **stored the required 1250kJ in just 145 seconds (2.4 minutes)**. On the one hand, this method is likely to overestimate the energy stored, since it doesn't account for heat loss. Furthermore, the power spike discussed earlier may have falsely increased the energy estimate. On the other hand, 2.4 minutes is much faster than the required 22 minutes, so we can be confident we met the PDS. This result is reinforced by the SoC estimation, but this system may also have overestimated the energy stored, as discussed later.

At the end of the test, the SFEE estimate for **stored energy was 4.2 ± 0.1 MJ**, exceeding the PDS total capacity requirement more than threefold. Upon reflection, this was due to three reasons.

Firstly, we only accounted for heating the PCM to 74°C, because that was the minimum needed to achieve latent heat storage. In the test, our heat exchanger outperformed expectations and heated the PCM to a mean temperature of 85°C, storing more heat. Secondly, we only accounted for latent heat when calculating the required volume of PCM. This was a serious oversight. Even at 74°C, one would expect the 12L of PCM to store 1.3MJ of sensible heat and 2.3MJ of latent heat, meaning the contribution from sensible heat should not have been neglected. Thirdly, due to the exploratory nature of the project, we had lacked confidence in our calculations. To ensure it met the specification, we had therefore applied a large safety factor. This posed no problem in this test, since the requirements were exceeded, but it was one of the reasons why the weight test was failed.

The SoC estimate aligned with the SFEE estimate in the early and late stages of

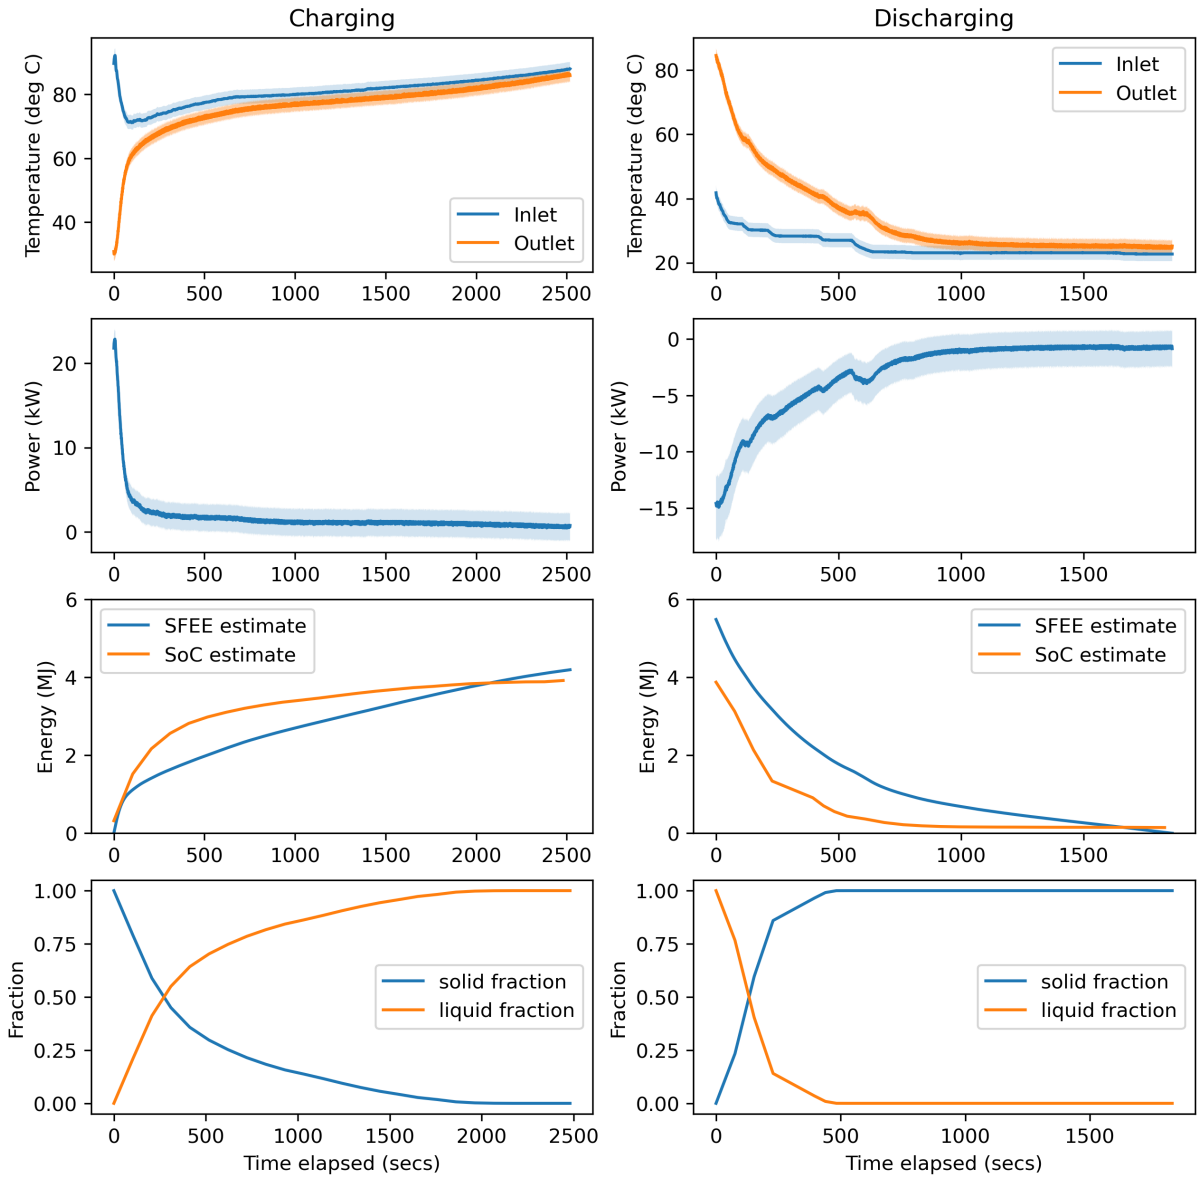

Figure 4.2: From the top: inlet and outlet temperatures, calculated power, calculated energy and calculated phase fractions over time for the charge and discharge tests.

the test, but exceeded it during the mid stages. This matched the expectation that the interpolation would perform worst during phase change.

From 2000 seconds, the SoC estimate stabilised to 3.7MJ. This suggests that it is accurately measuring the PCM unit state of charge, since we would expect the stored energy to stop increasing once it reaches capacity. At 2250 seconds, the SFEE estimate exceeds the SoC one, but this reflects error in that method rather than the SoC one. The power we were still putting into the PCM (0.8kW) was likely being lost as waste heat to atmosphere, which the SFEE estimate doesn't account

for.

As shown in 4.2, SoC system estimates that phase change (I.e. reduction in solid fraction) starts to happen immediately, but takes until 2000 seconds to complete. The phase change happens slower in the melting test than in the solidifying test. This matches the graphite test, but not the CFD, where melting happens faster due to increased convection around the pipe. This discrepancy is likely due to the significant heat losses for the experimental results, which are not present in the CFD. The alignment of SoC system phase fraction estimates with the graphite test results suggest this aspect performed well, even during phase change. While this did exceed expectations, later tests found opposing evidence.

4.2 Discharge Test

As specified by the PDS, the PCM unit must discharge **1250kJ** of heat **within 5 minutes**, so that the driver can use the energy to preheat the engine (from 10°C to 70°C) or the cabin without having to wait long.

These three requirements were tested immediately after the charge test, using the same set up, shown in Figure 4.1, except the hot water was drained and switched for cold water.

4.2.1 Expected results

The analytical and error estimation methods as described for the charge test were used to estimate **a peak discharge power of 7.8kW**. Indeed, discharge rate was the focus of these early thermal calculations, since it was expected to be the harder requirement to meet. As with the charge test, this was a peak power prediction and was expected to occur at the start of the test, followed by gradual reduction to zero. Thereupon, we would conclude that the PCM unit was fully discharged.

As shown in Table 3.2. the CFD predicted **a sustained discharge power of 14.8kW**, which is lower than the melting prediction due to the solidification front reducing heat transfer by preventing convection. For the same reason as with the charging CFD, this was expected to be an overestimate. It predicted that the PCM unit **would discharge 1250kJ in 85 seconds**.

4.2.2 Results

No leaking was observed during this test, since the PCM was solidifying around the pipes, which blocked the leaks.

Figure 4.2 shows that a **peak power of 15 ± 2.8kW** was calculated, which is 7.5kW lower than the peak power during charging. Since the pipes were full of hot water before being flushed out by cold water, the problem of overestimating peak power, explained for the charging test, also applies here. At 1000 seconds, the power reached a **steady value of 1kW**.

The inlet temperature took a while to reach a steady temperature of 23°C. This was because of practical difficulties when switching the hot water out for cold water. We had to keep a small amount of hot water in the boiler when we began adding cold water in order to prevent the pumps from running dry, which would damage them. Additionally, residual heat in the boiler and pipes kept the inlet temperature warm for longer than intended. This means that the time taken to discharge has been overestimated. Additionally, it means the peak power has been underestimated and that it took longer to reach a steady power.

The outlet temperature followed an exponential-like decrease from 80°C to roughly 30°C at the steady state. It is difficult to determine how much of the increased temperature at the beginning is due to the remaining effects from the hot water used in the charging test, but the results suggest that roughly half of the total energy was pumped out at an inlet temperature of >50°C, while the rest is below that. Considering the intended uses of the stored heat, this suggests that there is less capacity to output high grade heat, to preheat, for example, the engine. The steady state output temperature of below 30°C has limited use, but may be used for car cabin HVAC systems.

The required energy of **1250kJ was delivered in 100 seconds** according to the SFEE estimate and in **150 seconds** according to the SoC estimate, both of which significantly outperform the PDS requirement.

As shown in Figure 4.2, the SFEE energy estimate at the end of charging does not match the SFEE energy estimate at the beginning of discharging. This is because

there was a period of time between stopping data recording for the charging test and beginning the discharge test where the pumps and boiler were still running, so more energy was provided to the PCM than recorded, and our assumption that we ended the test when the heat being added was equal to the heat being lost to the surroundings must have been incorrect. To calculate the actual amount of energy stored, we vertically offset the SFEE energy trace so that it was equal to zero at the end of the test. This actually provided a more conservative estimate of total capacity, since heat losses cause the SFEE to underestimate power discharged from the PCM. Despite this, it was found that the PCM unit had been **charged to 5.5MJ** before the discharge test.

4.3 Storage Test

As specified by the PDS, the PCM unit **must retain 75% of stored energy over 48 hours**. Furthermore, the SoC system must be able to **monitor the PCM** during this period, in particular the movement of solidification fronts, to determine the best valve configuration for later discharging the unit.

This was tested with the same set up and method as the charge and discharge test. The key difference was that once the PCM unit had been charged, the boiler and pumps were turned off and it was left for 48 hours before being discharged, rather than being discharged immediately.

4.3.1 Expected results

The original intention was to predict the heat loss from an imperfectly-insulated container over 48 hours. However, this transient problem was too complex to solve analytically and would have required CFD simulation of the entire system – impossible given the computing resources available. Instead, one layer of 25mm thick insulation was chosen, with a thermal conductivity of 0.022 W/mK. The insulation covered 5 faces of the box, excluding the front manifold plate since this was difficult to insulate due to the splitters. The experimental

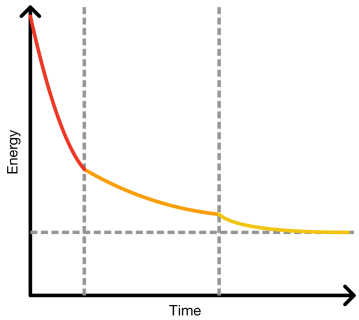

Figure 4.3: Expected graph for energy storage decline as heat was lost.

results would then be used to determine the insulation thickness needed to meet the PDS requirement. Therefore, the PCM unit was not expected to pass this test.

As shown in Figure 4.4, the energy stored was expected to reduce in three stages. It would decrease rapidly in the first stage, since convection can occur in liquid PCM and temperature difference with surroundings would be large. In the second phase, PCM was expected to solidify at the walls, insulating the rest of the liquid and reducing heat loss. In the final stage, full solidification would prevent convection and the temperature difference would be small, further reducing heat loss.

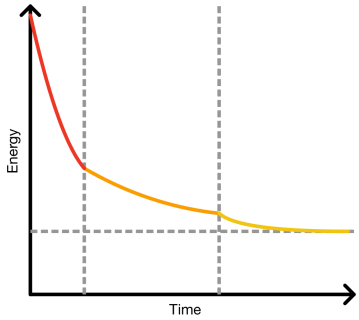

Figure 4.4: Expected graph for energy storage decline as heat was lost.

4.3.2 Results

The SFEE estimated that 5 ± 0.1MJ was stored in the PCM unit during charging, before the pumps were turned off and the inlet/outlet valves shut. It was discharged only 24 hours later, due to time constraints and the observation that it felt cold. Upon discharging, the SFEE energy estimate was **1750 ± 60kJ**. This was **35% of the originally stored energy**, indicating that the PCM unit did not meet the PDS requirement of retaining 75% of initial stored heat. The heat was initially discharged at 34°C, which fell to 25°C over discharging the first 690kJ. Such low grade heat would not be useful for pre-heating an engine and would only be of limited use for cabin heating. Upon reflection, the PDS requirement should have specified the temperature at which the heat needed to be discharged. The SoC estimation graph (Figure 4.5) did not match the expected curve, shown in Figure 4.4.The main difference was the ‘hump’ during the time when the PCM is solidifying. To further investigate, contour plots of the interpolated temperature field, sliced at mid-height, were produced (Figure 4.6). These reveal that the RBF was interpolating an unphysical temperature field, where the centre was cooling faster than the exterior regions. Furthermore, it was found to estimate temperatures (130°C) exceeding that of the water used in charging, which was also impossible. This confirmed the

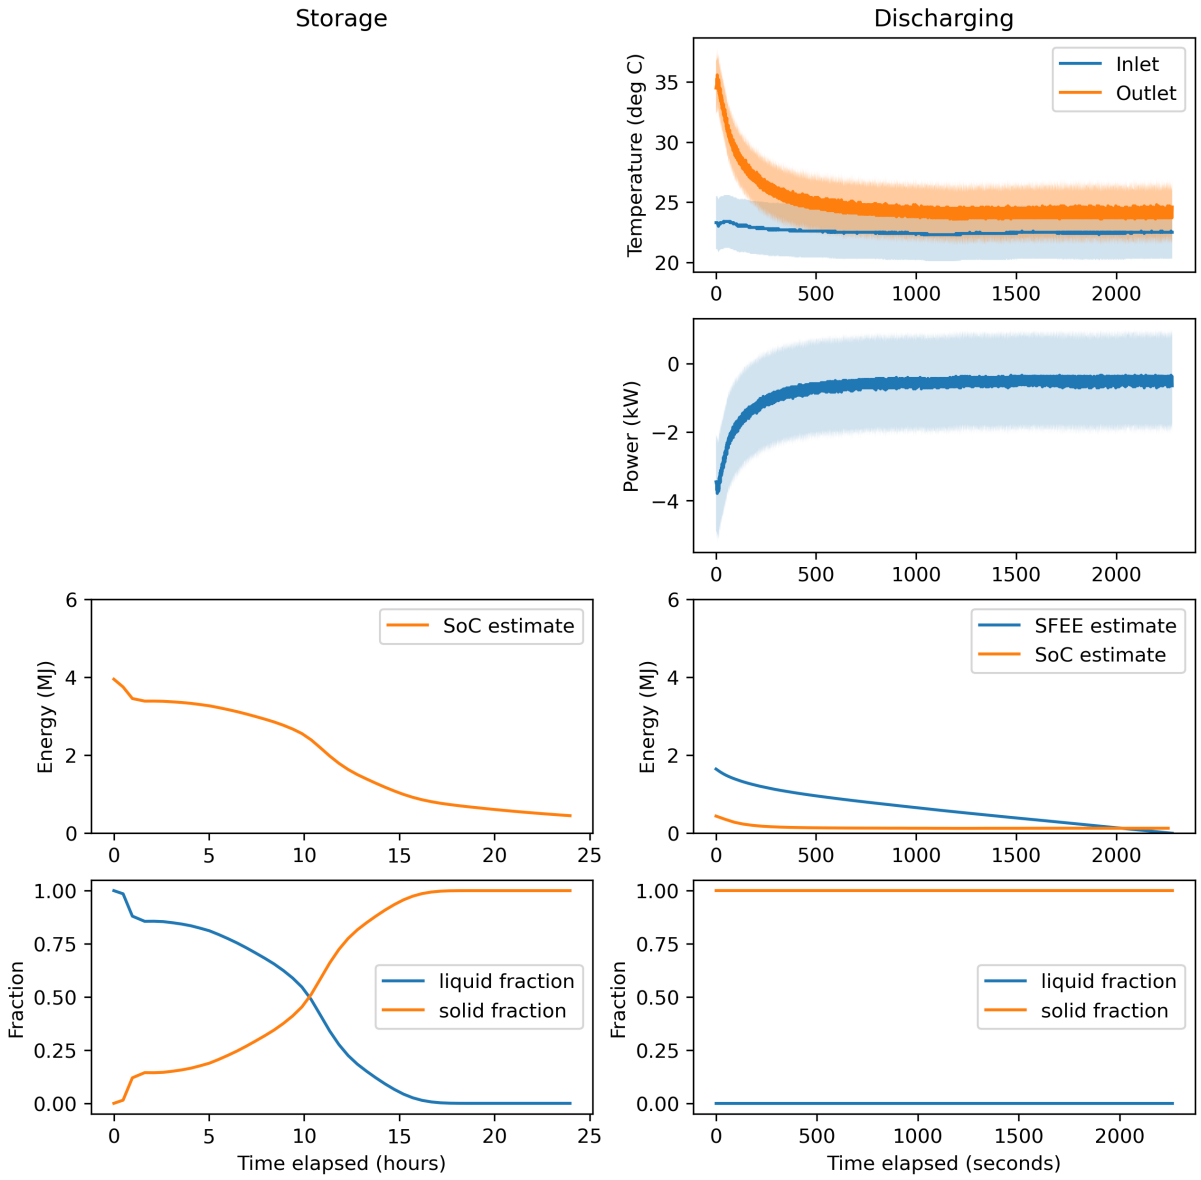

The charging plots are omitted because they are comparable to those in Figure 4.2.

Figure 4.5: From the top: inlet and outlet temperatures, calculated power, calculated energy and calculated phase fractions over time for the storage test. There are no SFEE estimates for power or energy during the 24 hour period before discharging, because water was not flowing through the PCM unit.

expectation that the method would not perform well during phase change, meaning that the **SoC system did not meet the PDS**. Upon inspection of the raw un-interpolated results, it was observed that the temperature sensors had cooled in pairs, as shown in Figure 4.7. For example, sensors 3 and 7 recorded the first instance of temperatures below 74°C after 3 hours, followed by sensors 4 and 8 after 5.5 hours. Comparing these results with the temperature sensor positions, shown

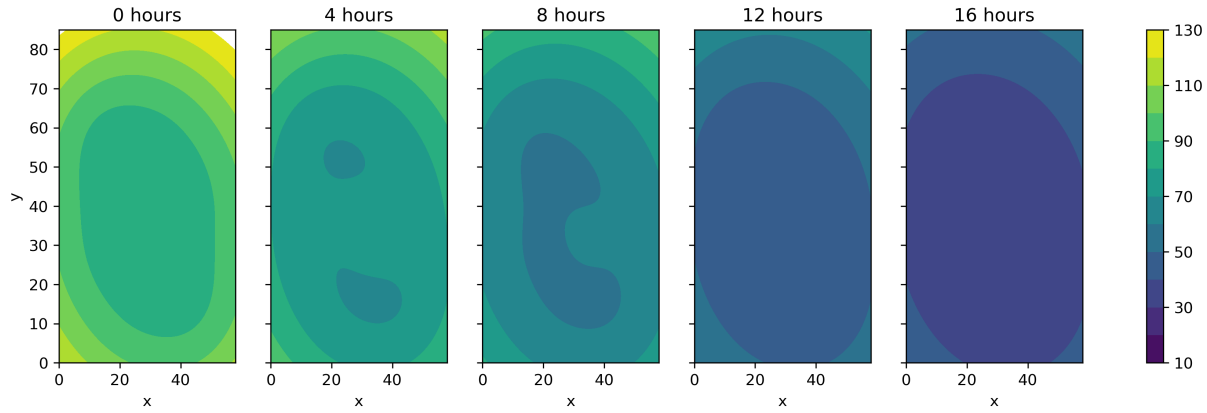

Figure 4.6: 2D slices through the mid-height of the 3D temperature field interpolated by the SoC system at different times during the storage test. Temperatures are in °C.

in Figure 4.9, it appears there are two mechanisms moving the solidification front through the PCM.

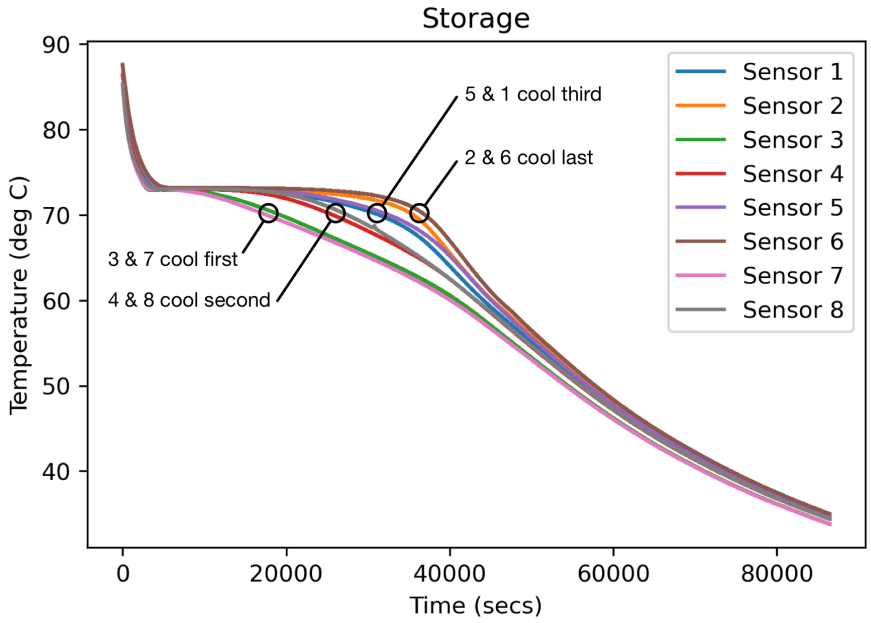

Figure 4.7: Individual temperature traces from the eight sensors in the PCM, intended for SoC estimation.

Firstly, heat loss to the atmosphere occurred primarily through the front pipe and manifold plates, since they were not insulated. This was confirmed by thermal images (Figure 4.8) taken during the test, which show those components were much hotter than the rest of the PCM unit. It is thought this mechanism caused a

solidification front to move through the PCM from front to back, which is why sensors 3 and 7 (closest to the front) cooled first (shown in Figure 4.9).

Secondly, it is thought that natural convection within the remaining liquid PCM promoted solidification at the bottom of the unit, since the fluid's density would have increased as it cooled. This caused a second solidification front to move from the bottom upwards, which is why sensors 4 and 8 (close to the front and closest to the bottom) cooled second (shown in Figure 4.9). This theory was supported by the CFD, which predicted liquid PCM would accumulate at the top of the domain during solidification, as shown in Figure 3.17.

The order in which the temperature sensors cooled, combined with the above understanding of the two mechanisms, could be used to track the solidification front movement as the SoC system had been intended to do. This is fully demonstrated in Figures 4.7 and 4.9, which explain why the remaining temperature sensors cooled in the order that was observed.

The similarity in results between sensors on either side of the PCM (but equal distances from the front) shows that the temperature field was symmetric. This suggests the number of sensors required to capture the entire temperature field could be made smaller.

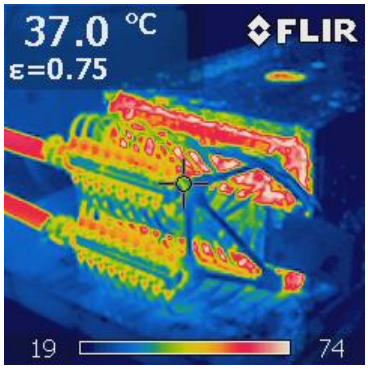

Figure 4.8: Thermal image showing that most of the heat loss was through the front manifold plate.

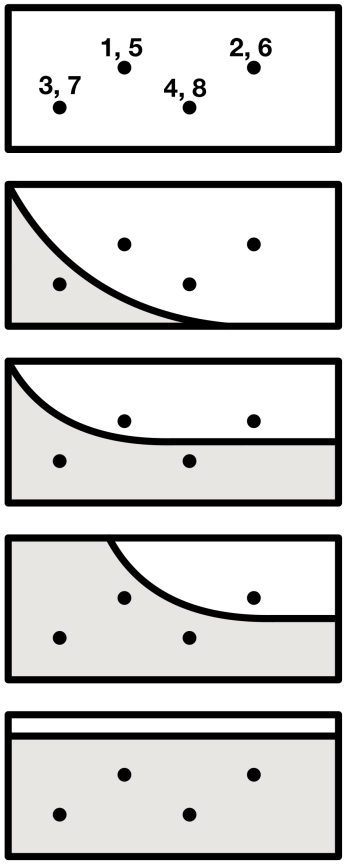

Figure 4.9: Side view diagram showing new theory for movement of solidification front. Solid regions shaded grey. Front of PCM unit on the left.

4.4 Localised Heating Test

As described in the sub-assembly overview, it was thought that directing heat to specific regions of the PCM would improve performance. This is why the PCM unit features 18-valve splitters and separate heat

transfer loops. Time restraints did not allow for quantitative testing of this functionality, but a brief qualitative test was conducted to determine whether this is a worthy avenue for further development.

4.4.1 Expected results

Based on the assumption that the inner region of the PCM is likely to heat up faster than the exterior regions, the inlet valves to the inner six heat exchanger loops were turned off. This configuration was maintained for the first 10 minutes of charging, whereupon they were re-opened. It was expected that this would accelerate the charging process by helping to maintain a uniform temperature field in the PCM, which would otherwise be hotter in the middle than on the outside. Predicting the extent of this improvement would have required CFD simulation of the entire PCM domain, which was impossible with the available computing resources. Qualitatively, it was expected that re-opening these valves would cause a spike in the charging power, by increasing the average temperature difference between the water and PCM.

4.4.2 Results

The test was inconclusive as to whether this valve configuration accelerated charging. Nonetheless, a 2.5 ± 1.4kW spike in charging power was observed when the valves were shut, shown in Figure 4.10. This was a 150% increase in power, compared to immediately beforehand, indicating that directing heat transfer to specific regions of the PCM can have a strong effect.

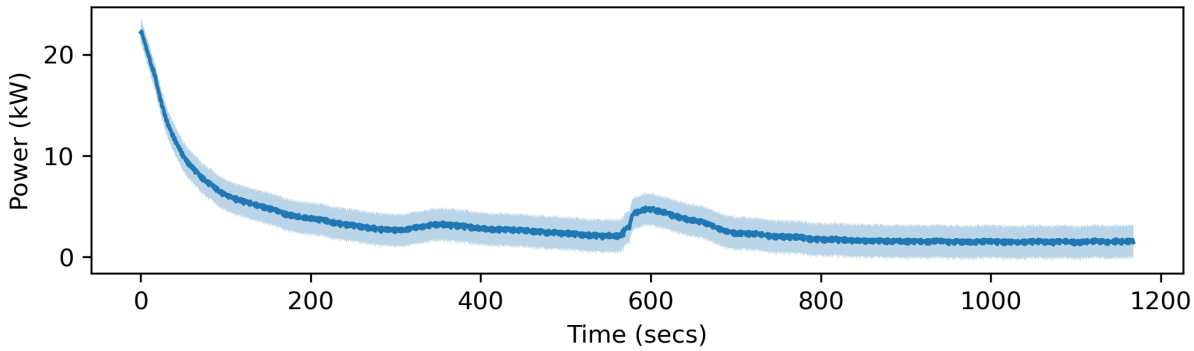

Figure 4.10: Charging power against time for this valve configuration. Note the power spike at 600 seconds, when the middle valves were re-opened.

This final test demonstrated that use of splitter valves and independent heat transfer loops can be used to control the temperature field and, thereby, the charging power. The failure to reduce the overall charge time should not be viewed as conclusive and further research into optimal valve configurations is encouraged. In particular, similar tests should be conducted for discharging.

4.5 Size and weight test

The **final size of the assembly was 348mm x 240mm x 630mm**, with a volume of $0.05m^3$. While the volume of the assembly is significantly less than the volume of a typical EV battery of $0.15m^3$, the dimensions of the assembly mean that it could not fit within the 1m x 1m x 150mm space that EV batteries usually take up. The final weight of the assembly with and without the PCM and water was measured to be **25kg and 37kg**. The majority of this weight was due to the PCM, copper pipes, and PCM container. This is much **larger than the target weight of 15kg**. However, since the PCM unit outperformed the PDS in terms of capacity, charge and discharge time, its size and weight could be significantly reduced in order to meet the size and weight requirements.

4.6 Implication of Results

The PCM unit outperforms the storage capacity, discharge rate and charge rate PDS requirements, so could be reduced in scale significantly. This would reduce the cost of the project, and make the unit more appropriate for use in cars, since it would be smaller and lighter. The leak during the charging test means that modifications to the design manufacturing process are needed in order to prevent leaking. The thickness of insulation used in the storage test was not sufficient to retain enough heat over a 48 hour period, which is as expected. Therefore, numerical calculations were conducted to determine how much insulation is actually needed. These implications of results will be discussed further in the following section, which focuses on redesigning the PCM unit in order to improve its performance.

Redesign 5

5.1 Insulation

Based on the storage test results, the design's insulation needed to be improved so that it could retain 75% of its initial energy over 48 hours. The ideal insulation thickness was calculated with Equation 5.1.

$$P = U_1 A_1 (T_{\text{PCM}} - T_{\text{ENV}}) + U_2 A_2 (T_{\text{PCM}} - T_{\text{ENV}}) \qquad (5.1)$$

Where P is the heat transfer rate, U is the overall heat transfer coefficient, A is the surface area, T_{PCM} is the average temperature of the PCM throughout the 48 hours and T_{ENV} is the temperature of the surrounding air, defined as 15°C. Subscript 1 refers to the container sides, and subscript 2, the front manifold side, shaded in orange and red, respectively in Figure 5.1. The surface area A_1 and A_2 were calculated to be 0.45m^2 and 0.05m^2, respectively. The overall heat transfer coefficient was calculated using Equation 5.2

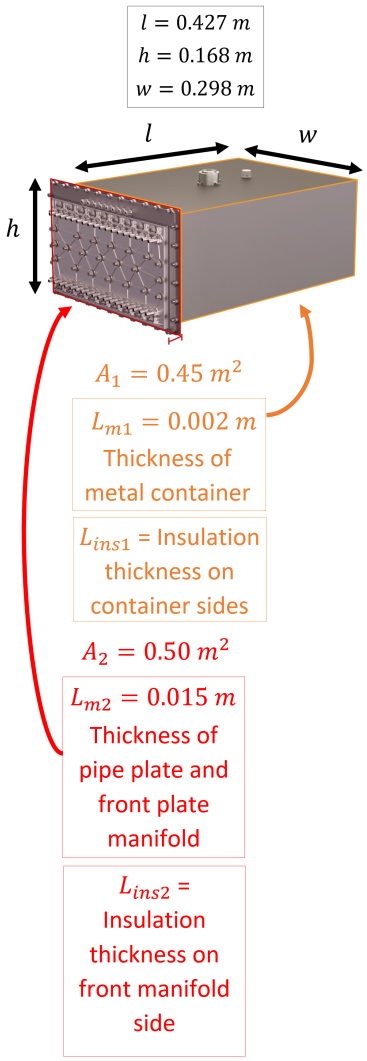

Figure 5.1: Diagram showing insulation dimensions.

$$U_\alpha = \frac{1}{\frac{L_{m\alpha}}{k_m} + \frac{L_{ins\alpha}}{k_{ins}} + \frac{1}{h}} \qquad (5.2)$$

Where the index subscript α has range (1,2), the subscript m is for the metal, the subscript i is for the insulation, L denotes the thickness and k, the thermal conductivity. On the container side (shaded in orange), the thickness of the metal, L_{m1} is 0.002m, and the insulation thickness, L_{ins1} is variable being solved for, while for the front manifold side (shared in red), the thickness of the metal, L_{m2} is 0.015m, and the insulation thickness, L_{ins2} is zero since no insulation was applied to this

surface – see Figure 5.1.

According to the storage test results, the initial energy of the PCM unit was 5MJ. To retain 75% of this initial energy, i.e. 3.75MJ, a maximum of 1.25MJ can be released by the system over 48 hours, resulting in an average heat transfer rate of $P = \frac{\text{Energy}}{\text{Time}} = \frac{1.25\text{MJ}}{48\text{hours}} = 7.23W$. The average temperature of the PCM between the start of the test and the time when 75% of the energy was remaining was 70.9°C, so this was the value used for T_{PCM}. This averaging method was an approximation used to reduce the complexity of this calculation.

Combining Equation 5.1 and 5.2 and rearranging to solve for L_{ins} allowed the minimum insulation thickness to be calculated. The results revealed that no matter the thickness of the insulation, the requirement could not be met if the front manifold was left uncovered. Therefore, the calculations were repeated, but with insulation covering the whole box so that $L_{\text{ins1}} = L_{\text{ins2}} = L_{\text{ins}}$. This showed that only 0.012m of insulation is needed over the entire system for 75% of the initial energy of the PCM unit to be retained over 48 hours – this includes covering the exposed fill port and pressure safety valve.

Since the standard thickness of one layer of insulation is 0.025m, which is greater than our required thickness of insulation, the proposed redesign consists of one layer of insulation around the entire box, including the front manifold plate. The excess thickness should account for any errors due to the approximations used in the calculations.

The insulation redesign can be seen in Figure 5.2. The entirety of the assembly is covered in insulation with every hole having a cover except the ones for the inlet and outlet hoses. In order for the valves to be covered, there must be an air gap between the front piece of insulation and front manifold plate. The volume of this air gap is approximately 9000cm^3, which is only 10.8g of air, assuming a constant density of 1.2kg/m^3. It would only take 567J of energy to heat this to from 15° to 90°C, assuming a heat capacity of 700J/kgK. This amount of energy is negligible in comparison to the actual capacity of the PCM unit, so should not make a significant difference to the storage ability of the unit over 48 hours.

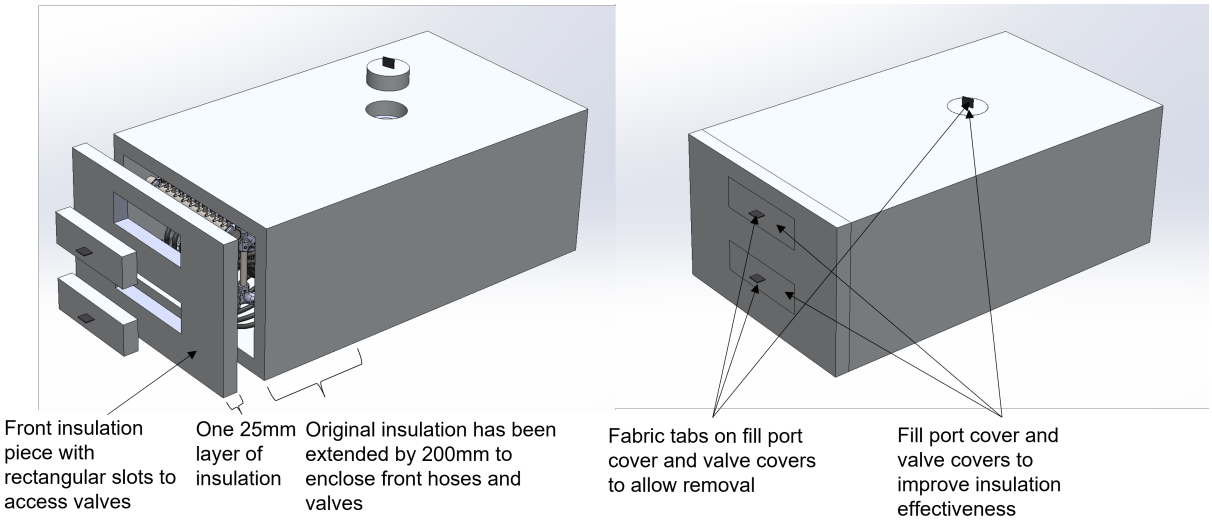

Figure 5.2: Insulation redesign

5.2 Heat Exchanger Water-Sealing

A key failure experienced during experimental testing was the slow leakage of water from a small number of copper pipe joints in the rear manifold. The manufacturing process mentioned in Section 2.2.1 for attaching the pipes with epoxy, would be adjusted where epoxy would be lathered on the underside of the rear pipe plate in the assembly jig, which would ensure an excellent sealing layer of epoxy. Alternatively, a less viscous, pourable epoxy resin could be procured and follow the team's initial plan for heat exchanger assembly.

5.3 State of Charge System

Experimental testing showed that the SoC system (i.e. the network of thermal sensors) was unable to correctly interpolate the temperature field in order to determine the state of charge, especially during phase-change.

However, as concluded from the storage test, the temperature field of the PCM can be assumed to be symmetric. This result could be used to reduce the overall number of sensors required by only having sensors in one half of the PCM volume. The other benefit of this is that the eight existing thermal sensors can much more effectively capture the phase distribution in the three sections of the PCM, and showing where these sections may be charged or discharged differently.

Thermocouples should be added to the inside surface of the wall, so that the tem-

perature field does not have to be extrapolated and a more appropriate function can be used for interpolation. Despite the extra cost of the ADCs needed for the thermocouples, the compact size of thermocouples is more suitable than the cheaper thermal sensors for this application.

The new SoC system would consist of 8 thermal sensors in the PCM volume, and 2 thermocouples, with placements shown in Figure 5.3. The cost of additional thermocouples with ADCs would cost approximately up to £35 – well within budget.

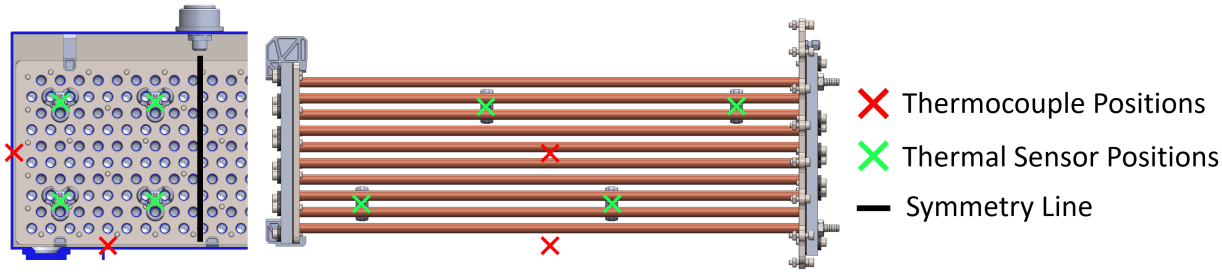

Figure 5.3: Revised Thermocouple and Thermal Sensor Locations

5.4 Pipe Plates

For the pipe plates, there are two options for its redesign, which are dependent on the joining method of the copper pipes and the pipe plates. The original design of the heat exchanger specified the use of brazing to join the pipes to the plate. However, due to budget and complications during manufacture, this was not possible with stainless steel pipe plates. In a redesign, these plates could instead be manufactured from brass which has a significantly lower brazing temperature. Additionally, brazing would be more suitable for long term use and frequent cycling. Yet, the high cost of plate brass makes this option infeasible. This leaves the option of continuing to use epoxy to join the pipes to the plates, which does not require stainless steel, so aluminium would be an appropriate replacement material. Aluminium would provide sufficient strength and rigidity to the structure whilst being more lightweight than stainless steel, and would present a significant cost saving compared to the original use of stainless steel.

Aside from revised material choice, no adjustments to pipe plate design are necessary.

5.5 Manifold Plates

The front manifold plate would instead be CNC milled from a suitable polymer. The benefit of this is the weight saved, and the improved machineability. Additionally, the polymer construction would act as an insulator – this is particularly important since aluminium is an excellent conductor, the current front manifold plate acts as a large heatsink, contributing to the failure of the 24 hour test.

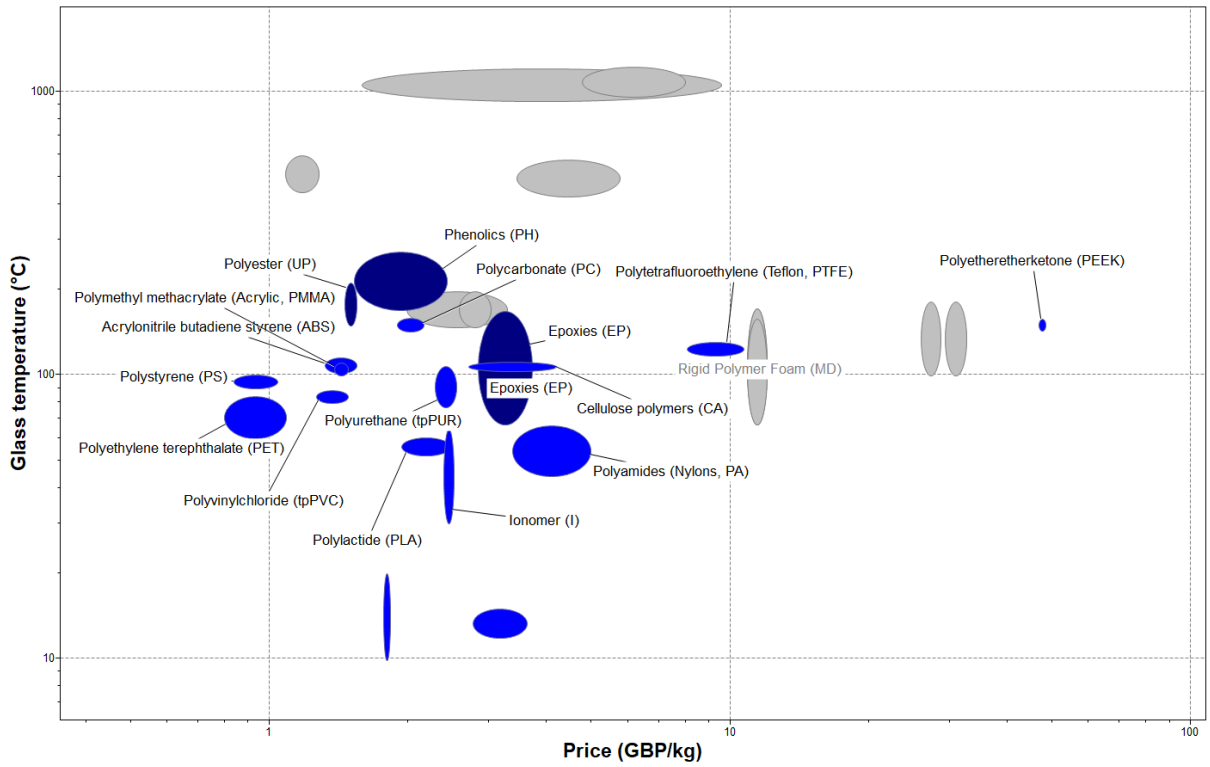

Figure 5.4: Ashby plot of polymers for material selection

For polymer selection, the front manifold plate must be able to operate in the presence of boiling water, so a glass transition temperature greater than 100°C is required. Whilst the component is not structural and does not bear loads, it must be a rigid polymer. Cost is another essential factor for material selection due to the limited budget available. Cambridge Engineering Selector (CES) software was utilized by the team to select a suitable polymer. Figure 5.4 shows an Ashby plot comparing different polymers against glass transition temperature and cost per kg – these are the most important factors for selection. From this, phenolics, polyester and polycarbonate were shortlisted for their combination of excellent temperature resistance and low cost. Polyester was decided against as this is typically used as

a fibre rather than a rigid sheet. The PDS specifies the use of recyclable materials and thus this polymer was rejected. Polycarbonate was found to be the most suitable material for its low cost, availability, strength and recyclability. Additionally, polycarbonate is transparent, allowing for visual inspection of the coolant flow through the manifold.

The rear manifold plate would remain as aluminium, and whilst aluminium is typically at least twice as dense as polymer, the rear manifold is fully submerged in PCM, so the manifold acting as a heatsink is beneficial in this case.

5.6 PCM Container

Our calculations revealed that the PCM container constituted 20.8% of the system weight. To reduce this, the current mild steel construction would be replaced with aluminium, reducing container weight by up to a third, but making welding more difficult. High-temperature polymer was considered as a replacement material for the container as it would greatly reduce mass, simplify manufacture and improve insulation. However, this would be excessively costly for the scope of this project and would only be suitable for mass-manufacture, where processes such as injection moulding or rotational moulding can be utilized.

Due to the all-aluminium construction, the fill port would instead be welded to the container, improving sealing. Additionally, the fill port's flange diameter and thickness would be reduced as fasteners are no longer necessary, further reducing weight.

A problem identified with the current design was it's poor repairability due to permanent fixtures (welding, epoxy, RTV silicone etc). Whilst a modular design was partly achieved with detachable manifold plates to provide access to copper pipe joints and allow replacement of gaskets, it was not possible to remove the heat exchanger from the container. Additionally, the rear manifold subassembly is submerged in the PCM so in order to repair this subassembly, the heat exchanger must be removed from the PCM container. The permanent RTV silicone gasket could be replaced by a removeable 2mm silicone gasket sheet shown in Figure 5.5. This would provide sealing and allow for the assembly to be taken apart.

It is currently very difficult to drain the PCM, since the fill-port is on the topside. A possible modification would be to have a drain port on the underside of the box - flush to the inside bottom surface of the container. This drain port would be placed in the corner to allow tilting of the container to fully drain the PCM. Both the drain port (shown in Figure 5.6) and the new fill port share the same revised design. The container rib was increased in depth to account for the height of the new drain port.

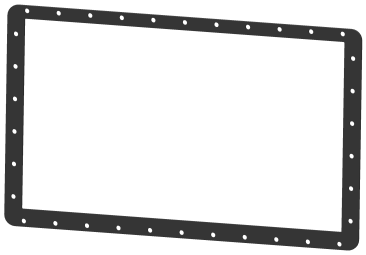

Figure 5.5: 2mm silicone gasket for sealing PCM Container

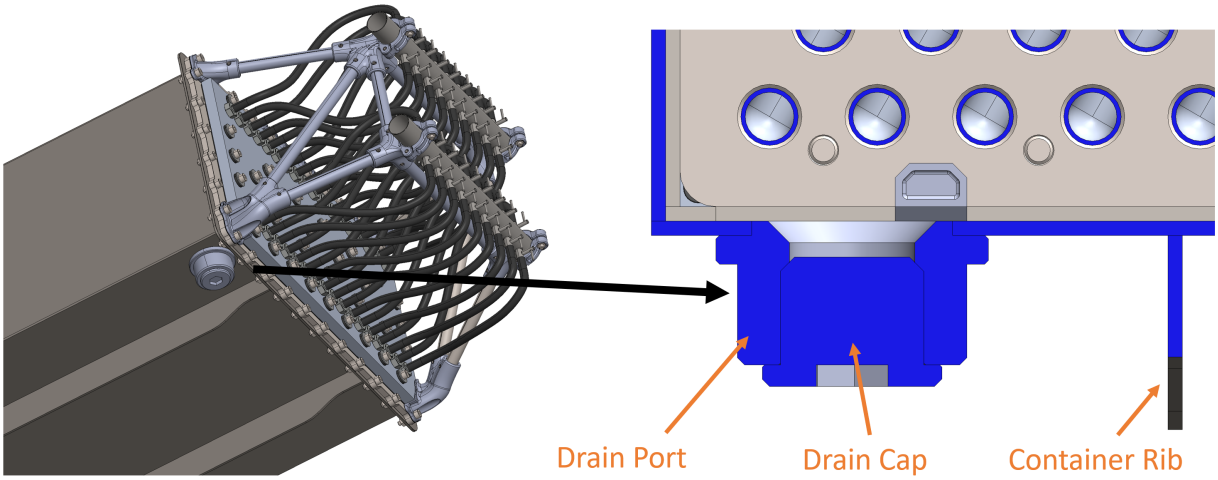

Figure 5.6: New addition of drain port to underside of PCM container

The redesign of the PCM container would have no extra cost to switch to the all-aluminium construction, with the only slight extra costs going towards the stock material needed for the drain port and silicone gasket.

Discussion 6

6.1 Design Review

Our design exceeded our expectations by outperforming the PDS requirements for charging, discharging and storage capacity. The uniqueness of our design, which, to the best of our knowledge, does not yet exist, is its ability to predict the energy content and state of the contained PCM using real-time data from the thermocouples. Based on this, the charging and discharging processes could then be optimised by turning the valves on and off.

Despite being functional and complete, there is still room for improvement. Leaks mean that better sealing is needed in order to be use-able in a vehicle. The insulation design needed to be improved to cover all surfaces. It should be made easier to drain out the PCM, and the overall size and weight of the product should be investigated in order to be reduced.

6.2 Future Progression

6.2.1 Improving existing design

The first stage towards improving the existing design would require partially disassembling the system. Once the heat exchanger has been removed from the PCM container, components could be replaced with the redesigned parts outlined in Chapter 5, then reassembled. These redesigned components would save weight, ease testing and improve heat retention.

The leaking rear manifold assembly was a manufacturing failure. This leak could be fixed by opening up the rear manifold assembly and using a pourable high temperature epoxy resin. This, combined with plasticine dams, could be used to create an consistent epoxy layer, sealing the pipe joints.

Owing to the insufficient time available, the team was unable to thoroughly test different strategies for the operation of the valve network. Experiments should be carried out to see how the valves of the splitter can be manipulated to make better

use of the available heat.

6.2.2 Testing with engine

Since the engine was never fixed, the system could not be integrated onto it. Future teams should aim to bring the engine back to full functionality. This would provide a more realistic testing experience of the system due to the engine's higher output flowrate, use of glycol-water mixture coolant and potentially higher fluid temperature. This would provide a better understanding of how much time is required to warm up the engine by discharging the stored heat – one of the key purposes of this PCM system.

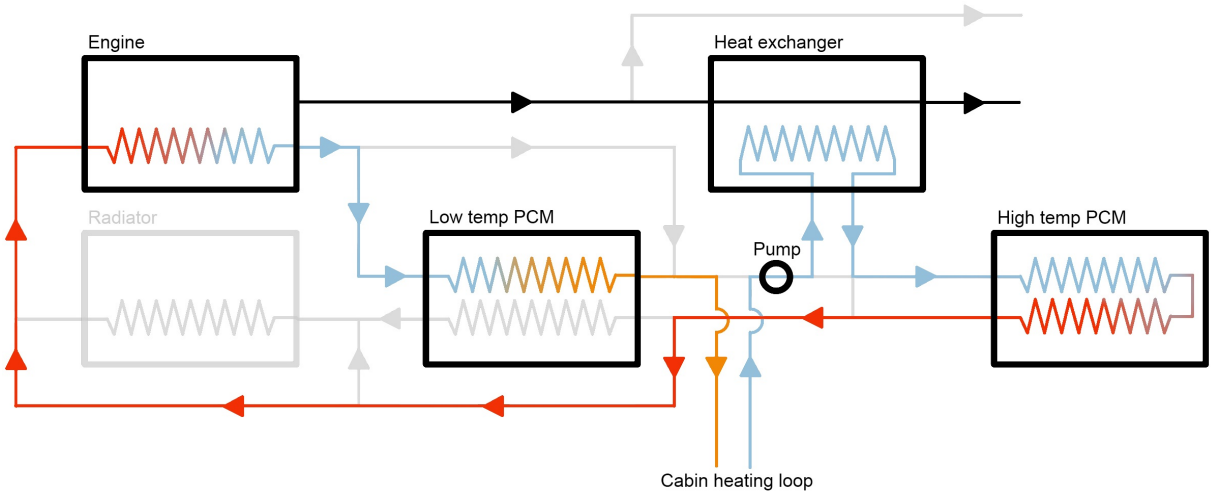

Figure 6.1: Proposed implementation of the PCM unit within a wider vehicle thermal management system.

An electronic valve system could be implemented to control the flow of hot and cold coolant to the different car components requiring thermal management (eg. cabin, engine, radiator, etc.) as well as to control the flow to zoned sections of the PCM system. Early in the project, the subgroup created a flow management schematic, which could be used to create the valve network, shown in Figure 6.1.

6.2.3 Redesigning full system

Since the PCM unit stores significantly more energy than is necessary, a full system redesign should focus on reducing the volume of PCM, which in turn this would reduce the size of the heat exchanger and container. This size reduction would also combat the weight issue as these sections of the system are also the heaviest, as

well as reduce cost and size.

The addition of copper fins in the heat exchanger (Which would require redesign of the full system) would aid thermal conduction throughout the PCM volume, lowering the required number of pipes required and reducing cost and weight, too.

A separate sub-team could design an electronic heat management system - as mentioned above. This could include the greater valve network as well as a compact, custom manifold system with electronic valves to control the charging and discharging of the sections of the PCM.

6.3 Team Progression

6.3.1 Project management

Overall, the project was managed very effectively, aided by consistently good communication. When the project began in October, the group started swiftly and began to make significant progress. The project manager scheduled regular meetings, including internal meetings, meetings with our supervisor, and supergroup meetings with group leaders.

For each sub-group meeting with our supervisor, a PowerPoint presentation containing a progress summary, points for discussion, and actions for the coming week was presented. This enabled extremely effective communication between group and supervisor, and was also a very useful reference later in the project. During the initial design phase, a Gantt chart was created by the PM; this was updated each project phase. It detailed actions to be taken, their deadlines and relative priorities. After most meetings, a task list was created and jobs assigned to group members. To-do lists were created on Excel and uploaded to Teams, but were occasionally disorganized and progress trackers were not kept up to date.

Mid-way through the project we had a group feedback session. Overall, feedback was well-received by the group members and significant improvement was seen in group performance after this. For example, after being criticized for not delegating enough, the PM began delegating tasks more frequently, which aided in faster project progression as the workload was shared. Our sub-assembly did not directly

connect to any of the other sub-assemblies, so less communication was needed compared to other DMT projects. There were no particular communication issues with other sub-groups.

6.3.2 Manufacturing

The manufacturing phase began soon after the manufacturing gateway – the team a was well prepared to progress at this stage, and there were no major changes that needed to be made. Most of manufacturing was done manually by the team, or using STW and Hackspace resources for additive manufacturing and laser cutting. The majority of the manufacturing ran as expected. We also used the pit garage for some more specialist work that prevented us from using the STW – mainly for our initial attempts at brazing, and also for our epoxy work towards the end of the project.

The most significant issue during manufacturing was with the stainless-steel pipe plates. It was initially planned for these to be laser cut in the STW. Unfortunately, the initial attempt resulted in the plate being warped after the process. After consulting with our supervisor, we attempted to re-laser cut the pate with a modified pathing, but this proved unsuccessful once more. As a result, we had to outsource the manufacturing to a company that used waterjet cutting, which resulted in a significant cost increase.

6.3.3 Budget

Budget was a significant constraint throughout this project. The group was initially given a budget of 1000 pounds, but the magnitude of the project quickly revealed the need for a budget extension. During the budget extension request, the sub-group presented three different possible designs with associated budgets and demonstrated the effort made to reduce the budget throughout the design (e.g., using a manifold plate instead of bending pipes) and material procurement (e.g., bargaining lower prices for our most significant costs, the PCM) while meeting the system’s requirements for efficient heat transfer (necessitating a large length of small pipes) and heat storage (requiring a large container). In addition to the 1000 pounds granted at our request, 315 pounds were compensated due to inadequacy

in the production of parts that were guaranteed to be delivered correctly. Overall, 2158.17 pounds of the 2315 pounds given was spent in the creation of this novel product. After deducting the costs incurred by the group as a result of unexpected setbacks, the reproduction of such an assembly is expected to cost only 1700 pounds.

6.3.4 Procurement

Due to the large number of different items that needed to be purchased, each member was assigned a number of suppliers to avoid multiple communications with the same supplier. An excel sheet was used to concisely summaries all procurement-related information including the items purchased from each supplier, their cost, including delivery costs, the lead time, whether or not they were ordered and received and the member responsible for that procurement. One team member was assigned the responsibility of keeping the table up to date and reviewing all purchases. This method ensured that there were very few communication errors throughout the process.

In general, for most suppliers, the purchasing and delivery processes went smoothly and within the estimated lead times. However, many issues were faced for the procurement of our most critical component: the PCM. Because of the specifications for this material, only two suppliers, one German and one English, provided what the group required. The group chose the English company because of the lower costs and shorter delivery times they offered, as well as because industry experts recommended it. Nonetheless, the time required to obtain the PCM from this supplier was extremely long; from negotiating the price of a 30kg sample (considered very small in comparison to the usual size they sell) to getting registered in their system to waiting for the payment to go through before it could finally be dispatched and arrive three days later. The poor performance of the supplier’s administrative team was largely to blame for the PCM procurement delay.

7 Conclusion

The PCM unit was able to **charge and discharge within the required times**. It stored 1250kJ in 145 seconds, outperforming the PDS requirement of 22 minutes, and discharged 1250kJ in 100 seconds, faster than the PDS requirement of 5 minutes. The PCM unit also **met the total capacity requirement**, with the discharge test revealing that 5.5MJ had been stored, significantly more than the PDS-required 1250kJ.

The PCM unit only retained 35% of its energy after 24 hours, **failing to meet the heat retention requirement** of 75%. Analysis revealed this was due to the lack of insulation on the front manifold plate and around the fill port. Indeed, it was predicted that the insulation thickness could actually be reduced to 12mm if if these regions were covered.

The PCM unit **did not meet the size and weight requirements**. However, these could likely be reduced while still meeting the charge/discharge time and total capacity requirements, since it significantly outperformed on those. Weight could be further reduced by manufacturing certain parts from aluminium or polymers.

The SoC system estimated the energy stored in the PCM unit with reasonable accuracy when the PCM was fully solid or fully liquid, but under-performed during phase change. Though disappointing, this was not surprising, given the current immaturity of latent heat state of charge estimation. Improvements to the SoC system were proposed, such as making use of symmetry observed in the temperature field to improve resolution. Furthermore, it was proposed that thermocouples be attached to the container walls to measure boundary conditions and negate the requirement for extrapolation. It was found that the use of splitter valves and independent heat transfer loops in the design enabled local control over the temperature field. Crucially, it was shown that this could control the charging power.

The stability test found that powdered graphite remains suspended in liquid PCM but exhibits significant settling during solidification. This is thought to partially

negate the thermal conductivity increase associated with adding graphite. PCM containing 6% wt graphite powder melted the fastest, but the graphite concentration was found to have no significant impact on the solidification behaviour.

The CFD showed that the heat transfer rate differs for solidification and melting, due to the solidification front limiting convective heat transfer. This illustrated a hysteresis effect which is important to consider when predicting heat transfer with phase change. The rate of phase change predicted by the CFD aligned well with the graphite test experimental results, validating the CFD. This was also true for the solidification CFD temperature profile, but not for melting. The discrepancy for the melting is likely due to heat losses in the experimental results, so the CFD results are still reasonable.

It was concluded that the key failures of the design were insufficient insulation, minor leaking of the heat exchanger and SoC system inaccuracy. The redesign addressed these issues with the an improvement of the heat exchanger assembly process, and the introduction of a new insulation solution which left all hot metal surfaces insulated. Furthermore, the use of temperature field symmetry and addition of wall-positioned thermocouples was proposed to improve the SoC system. The redesign also included other improvements throughout the system to reduce weight, cost and ease future work.

References

[1] Isaac Blanc, Lola Marcos Gamez, and Archie Platt. *Engine Fixing*. Apr. 2022 (cited on page 1).

[2] Yulong Ding, ed. *Thermal Energy Storage : Materials, Devices, Systems and Applications*. 1st ed. Royal Society of Chemistry, 2021 (cited on page 2).

[3] Tilman Barz et al. 'State and state of charge estimation for a latent heat storage'. In: *Control Engineering Practice* 72 (Mar. 2018), pp. 151–166. DOI: 10.1016/j.conengprac.2017.11.006 (cited on pages 2, 30).

[4] Gerald Steinmaurer, Michael Krupa, and Patrick Kefer. 'Development of sensors for measuring the enthalpy of PCM storage systems'. In: vol. 48. Elsevier Ltd, 2014, pp. 440–446. DOI: 10.1016/j.egypro.2014.02.052 (cited on page 2).

[5] Teboho Clement Mokhena et al. *Thermal Conductivity of Graphite-Based Polymer Composites*. Sept. 2018. DOI: 10.5772/intechopen.75676 (cited on page 12).

[6] *Crodatherm 74 Data sheet* (cited on page 12).

[7] R. Sathiyaraj et al. 'Enhancement of heat transfer in phase change material using graphite-paraffin composites'. In: vol. 172. EDP Sciences, June 2018. DOI: 10.1051/matecconf/201817202001 (cited on page 12).

[8] Wei Li et al. 'Preparation and performance analysis of graphite additive/-paraffin composite phase change materials'. In: *Processes* 7 (7 July 2019), pp. 1–12. DOI: 10.3390/pr7070447 (cited on page 12).

[9] Alok Agrawal and Alok Satapathy. 'Development of a heat conduction model and investigation on thermal conductivity enhancement of AlN/epoxy composites'. In: vol. 51. Elsevier Ltd, 2013, pp. 573–578. DOI: 10.1016/j.proeng.2013.01.081 (cited on page 12).

[10] Ziming Wang et al. 'Vertically aligned and interconnected graphite and graphene oxide networks leading to enhanced thermal conductivity of polymer composites'. In: *Polymers* 12 (5 May 2020). DOI: 10.3390/POLYM12051121 (cited on page 13).

[11] Muhammad Usman et al. *Improvement of thermal conductivity of paraffin wax, a phase change material with graphite powder; Improvement of thermal conductivity of paraffin wax, a phase change material with graphite powder*. 2020 (cited on page 18).

[12] Erin Yurday. *Average Car Journeys in the UK*. NimbleFins. Oct. 4, 2021. URL: https://www.nimblefins.co.uk/largest-car-insurance-companies/average-car-journey-uk (visited on 11/04/2021) (cited on page 27).