OPTIMIZATION OF FASP FOR LI-ION CATHODE MATERIAL PRODUCTION

by

Connor Hickey

Submitted to the Department of Mechanical Engineering in partial fulfillment of the requirements for the degree(s) of

BACHELOR OF SCIENCE IN MECHANICAL ENGINEERING

at the

MASSACHUSETTES INSTITUTE OF TECHNOLOGY

May 2024

Authored by: Connor Hickey

Department of Mechanical Engineering

May 10, 2024

Certified by: Sili Deng, PhD

Class of 1954 Career Development Professor, Assistant Professor of Mechanical Engineering, Thesis Supervisor

Accepted by: Kenneth Kamrin

Professor of Mechanical Engineering

Undergraduate Officer, Department of Mechanical Engineering

OPTIMIZATION OF FASP FOR LI-ION CATHODE MATERIAL PRODUCTION

By

Connor Hickey

Submitted to the Department of Mechanical Engineering on May 10, 2024, in Partial Fulfillment of the Requirements for the Degree of

BACHELOR OF SCIENCE IN MECHANICAL ENGINEERING

ABSTRACT

The need to improve battery technology is higher than ever given the projected increase in battery consumption in the next decade and beyond. One key limiting factor of batteries is their cost, and one major way to reduce battery costs is by decreasing the time needed to produce them. The FASP system uses Flame-Assisted Spray Pyrolysis, and the principles of combustion to speed the process of creating the materials for lithium-ion cathodes, more specifically the NCM-811 variation of the lithium-ion batteries. This study aims to identify key factors in improving the powder production of the FASP system. One way it aims to do this is by creating a CFD simulation, within ANSYS, to create an accurate picture of the behaviour of the fluids within the pipe flow. Another way it aims to optimize FASP is by conducting various experiments to test the simulations and try to find areas of disagreement to find a direction to improve the CFD model. The final way this paper aims to optimize FASP is by conducting several powder-producing experiments and testing various variables to find the best combination.

Thesis supervisor: Sili Deng, PhD

Title: Class of 1954 Career Development Professor, Assistant Professor of Mechanical Engineering

Acknowledgments

Special thanks also go to the members of Deng Energy and Nanotechnology Group, especially Chuwei Zhang for supervising of experiments.

Contents

List of Figures

List of Tables

1. Introduction

There are many types of batteries such as Lead-Acid Battery, Nickel-Cadmium Batteries, Nickel-Metal Hydride Batteries, Lithium Polymer Batteries, and Lithium Batteries. Lithium-ion batteries are currently the most used type of battery, with their expected share of the market set to increase from around 45% to over 60% in the next 10 years [1]. The number of Lithium-ion batteries is set to increase by over 300% between 2024 and 2030. Their popularity is due to their great energy density and voltage capabilities. Lithium-ion batteries have one of the highest energy densities available for commercial use, with a capacity of almost 300Wh/kg compared to 75Wh/kg for other battery types [2], while their voltage of about 3.6V is 1.5 to 3 times that of other battery technologies. The high energy allows for more energy to be stored in the battery at the same weight, whilst the high voltage improves the efficiency of the battery. Since Li-ion batteries are the highest-performing battery and most widely used battery, this study is focused on producing the materials for Li-ion materials, more specifically the NCM811, which has a particularly strong energy density of up to a potential 300Wh/kg [3] due to the chemistry. The large energy density of this configuration comes from the increased nickel concentration within the configuration, leading to increased material activity which produces a greater energy density. The cobalt stabilizes the laminar structure of the material, which helps increase the discharging capacity and the manganese provides necessary stability during charging and discharging [4]. The NCM811 has a chemical formula of $LiNi_{0.8}Co_{0.1}Mn_{0.1}O_2$.

Given the projected increase of lithium-ion batteries within the next decade [2], manufacturing will have to increase to meet that demand. One way to increase production is by decreasing the time it takes to make each lithium-ion battery. The current production process for lithium-ion batteries is capital-intensive, contributing significantly to the overall cost of the batteries. By shortening the production time, manufacturers can reduce labour and operational costs, ultimately lowering the cost of producing batteries and making the end products more affordable for consumers.

The current process involves many stages. The first stage is called precipitation. Metal salts including the metals of nickel, manganese, and cobalt are added to water and are thoroughly mixed inside a chemical reactor. The pH of the liquid is altered throughout the process to help the precipitation occur; this is done with the help of various chemicals. Once the metal particles have precipitated out of the solution, they are removed from the solution, dried, ground up, and put through a sieve. The second stage involves adding solid lithium to the precipitated metals. This is done until the lithium permeates the other metal particles. This mixture is then heated at 750°C to ensure complete mixing. The last step is to put the particles through a deagglomerator. This separates particles that have come together. This process takes about 34 hours – 10 hours for the coprecipitation, 13 hours for the heating and 11 hours to slowly taper off the heat to prevent cracking of the particles [5]. The process takes time and therefore contributes to the high cost of these batteries.

To help tackle the problem of increasing the speed of lithium-ion cathode material production, a novel technique was created called FASP (Flame-Assisted Spray Pyrolysis). The FASP system was able to reduce the calcination time from 20 hours to 20 minutes, while still maintaining a layered structure with cation mixing at a low level and producing promising electrochemical performance [6]. This initial study introduced the idea of combustion in order to speed up the process of making cation

materials. The following study examined the correct temperature for heating the particles before the combustion flame. The morphology is crucial for the quality of the li-ion particles produced by the FASP. It was found that no preheating to the droplets created a broken particle, as its first exposure to heat would have been a powerful flame. This would cause the moisture to be removed too quickly from the droplet, causing the particles to break [7], and for an inferior lithium-ion cathode to be made. The preheating section was increased in steps of 50°C, from 25°C to 375°C. 175°C was found to have the best particle morphology, temperatures past that started causing deformations and hollow shells. A further study was done analysing multicomponent droplets and how each of the components in a lithium-ion NCM cathode interacted with each other. It helped bring an understanding of how different components in the droplet migrate during the evaporation process, and how urea helps disrupt that process, leading to a more uniform distribution of the elements in the NCM cathode particles [8].

This study aims to build off those studies by aiming to figure out the behaviour of the inside flow. To do this, numerous experiments will be conducted on the FASP system, and the temperature of three different points will be measured inside the burner chamber. These experiments will be compared to a CFD model, which will attempt to accurately simulate the fluid flow inside the system, and will also measure the temperature, at those same three points. The experiments done on the FASP will include varying the carrier gas flow rate, the air flow rate into the premix chamber, and the methane flow rate into the premix chamber. Any discrepancy in trends between the data of the experiment and that of the simulation will have to be thoroughly analysed for an explanation for the deviation to be found. This will help generate a better understanding of the system, and hopefully help to generate a good foundation for a more complex simulation model to follow, in the future. On top of that, some powder-producing experiments will be conducted to help identify which factors cause changes in the size and the morphology of the particles produced in the powder, hopefully informing for future directions of investigation.

2. Background

2.1 Particle Morphology

In the drying process of a droplet, the morphology of the resulting particle is highly dependent on the temperature of the heating. The droplets start by undergoing an initial shrinkage, which happens at a constant rate. This occurs until a shell is formed on the surface. If the conditions are correct, and a dry shell is present, this then forms a solid particle. However, if the conditions are incorrect, the particle morphology can be badly affected. At lower temperature, the moisture is brought out of the droplet at a slower rate leading to an initial lower rate of shrinkage and moisture loss. However, this increases the risk of deformation, shrinkage, and collapse of the particle. At higher temperatures, the droplets can inflate, blister or break due to the rate of moisture loss and initial shrinkage being too high for a well-shaped particle to form. The particle can be formed at a gas temperature lower than the boiling point of the liquid, or moisture, being evaporated [9].

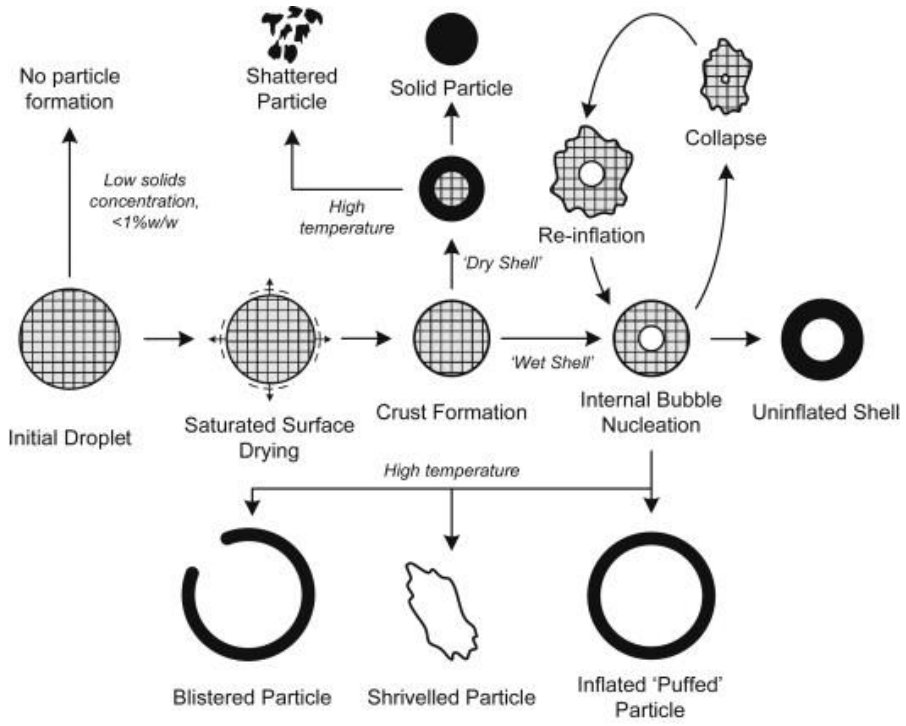

Figure 1: The following image shows the various morphologies of particles that come from the heating of droplets [9].

The 'wet shell' deformation occurs when droplets have incomplete drying because they are being heated by gas that is too cold and brings out the moisture too slowly. The 'high temperature' label, on Figure 1, indicates the morphology of various particles where the moisture extraction occurs too quickly.

2.2 Reynolds Number

The Reynolds number is a dimensionless quantity that is used to predict flow patterns, in fluid dynamics. It helps to determine if a flow is to follow the laminar regime or the turbulent one. The Reynolds number is a ratio of the inertial forces to the viscous forces in a flow.

$$Re = \frac{\rho v L}{\mu} \quad (1)$$

Equation 1 shows the formula for Reynolds number, where ρ represents the density of the fluid, v represents the speed of the flow, L represents the characteristic length, and μ represents the dynamic viscosity of the fluid. The typical cut off point for laminar flow is when the Reynolds number reaches 2300. At that point, the smooth flow starts to transition into a turbulent flow. When the Reynolds number reaches 4000, the flow is turbulent.

2.3 Entrance Length

The entrance length in fluid dynamics refers to the distance it takes a for a flow to become fully developed after entering a pipe. Fully developed means that the velocity profile is stable and will

remain constant throughout the rest of the pipe until there is a force that compels the velocity profile to change.

$$L_e = 0.06 * Re * D \quad (2)$$

Equation 2 shows the entrance length equation, in a laminar regime, where Re is the Reynolds number and D is the diameter of the pipe. It is an important length to calculate in research, as it helps to inform where in the simulation the pipe flow should be checked in order to confirm whether a flow is following the laminar regime or not. It helps the analysing of CFD simulations.

3. Experimental Set-Up

3.1. Set Up for the Solution

The first step is to clean the beaker and the measuring cylinder with deionized water. Deionized water is water that has had all the positive and negative ions removed. It is therefore very pure. It is good for cleaning due to its lack of ions, which is why deionized water can help eliminate contaminants on the surface of the apparatus in the lab [10]. Deionized water is also used for experiments as it is purer than tap water in the lab. Once the beaker and the measuring cylinder are cleaned, the next step is to measure out a set volume of deionized water in the measuring cylinder and then add that deionized water to the beaker.

A piece of weighing paper is placed on the mass scale. The mass scale is then set to zero. Hydrated cobalt nitrate, hydrated manganese nitrate, hydrated nickel nitrate, nickel hydroxide, and hydrated lithium nitrate are then measured on the scale so that the ratio of lithium to nickel to cobalt to manganese equals 10 parts to 8 parts to 1 part to 1 part. The molar ratio of hydrated nickel nitrate and nickel hydroxide is set at one to one, so half of the nickel comes from each compound. The amount of each compound added depends on the concentration of the suspension that was being used for each experiment, the concentrations used throughout the experimentation were 0.1M, 0.25M, and 0.5M.

All of the compounds are then added to the beaker with the deionized water. A magnetic stirrer is used and the beaker is placed onto a magnetic platform and left there for about 20 minutes while the stirrer is producing an even mix of the elements throughout the suspension. A transparent plastic cover is placed on top of the beaker to ensure no suspension is spilled during the stirring process.

3.2. Set Up for the Horizontal Experiments

Once the suspension is fully mixed, it is placed into a syringe. This syringe is attached to an ultrasonic atomizer, of frequency 130kHz, that automatically keeps the flow rate of our system at 0.6ml/min. The suspension goes through the atomizer which turns the suspension into a fine mist. An ultrasonic atomizer is used as it generates uniform droplets, as opposed to the 2-fluid-atomizer [11]. Another advantage of the ultrasonic atomizer is that it makes the setup of the system simpler; there is no need for an air compression source.

The system itself is made up of many sections, including 3 preheating sections, a reducer, an inner tube, a premix chamber, a burner chamber, and a collector. The preheating sections are all temperature-adjustable and are set to various temperatures, with the first preheating section having the lowest temperature, the second preheating section having the middle temperature, and the final preheating section having the highest temperature. The first preheating section is 24 inches long and has a diameter of 2 inches. The second preheating section is 18 inches long and has a diameter of 2 inches. The final preheating section is 18 inches long and has a diameter of 2 inches. All 3 preheating sections are lined up horizontally in the ‘horizontal' experiments.

The reducer has the job of decreasing the diameter of the system. It decreases the diameter from 2 inches to 1 inch. The reducer was designed with an inward curve to reduce the amount of particle buildup within this section. The reducer is 1.9 inches in length.

Following on from the reducer is the inner tube, which is 2 inches long and has a diameter of 1 inch. The inner tube carries the particles, and the flow, from the preheating sections into the burner chamber. The inner tube is surrounded by the premix chamber, which has a mixture of air and methane. The mixture of gases in the premix chamber, enters the burner chamber through a heat-resistant ceramic honeycomb. The heat-resistant ceramic honeycomb is used to ensure the gas enters the burner chamber uniformly. The methane in the powder experiments was set at 1.1l/min, whilst the

air mass flow rate was set at 22l/min for the gas mixture entering the premix chamber. The ceramic honeycomb was made so that the methane-air mixture entered the burner chamber with an inner diameter of about 1 inch and an outer diameter of about 3 inches.

The end of the ceramic honeycomb and the inner tube mark the start of the next section which is the burner chamber. This is where the methane-air mixture and stream from the inner tube join together. The flame within the system happens right at the end of the ceramic honeycomb at 800°C. This flame only directly affects the methane-air mixture, and not the stream from the inner tube. The methane is initially ignited by an electric pulse at the start of the experiment, which is subsequently turned off after the flame is detected by a rapid temperature rise in the system. This rapid temperature rise is detected by 3 separate thermocouples, which all have their tips on the centre of the burner chamber, and are 0.5 inches, 3 inches, and 7 inches away from the start of the burner chamber respectively. The methane keeps this flame alight for the entire experiment until the methane tap is turned off. The burner chamber has a 3-inch diameter and is 12 inches in length. The final section of flow is 12 inches, with a diameter of 2 inches.

Lastly, a downward 90° bend leads the particles and the flow into the collector. The collector has filter paper and metal gauze to ensure all the particles are collected and it forms a fine black powder, which is the lithium-ion cathode material. The collector is wrapped in an insulating blanket to ensure no major heat loss to the environment.

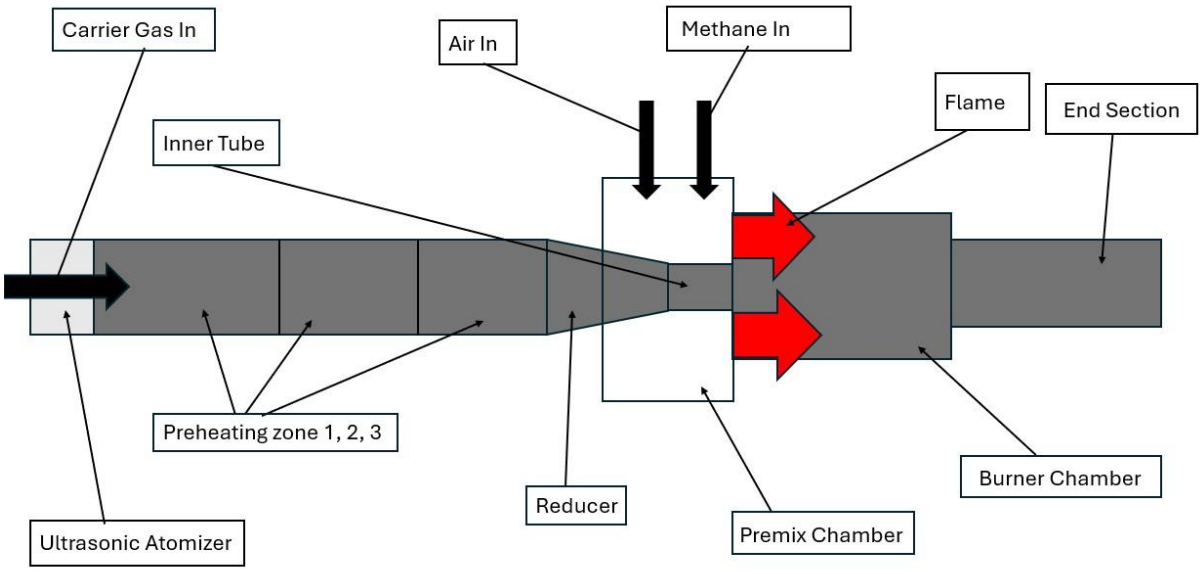

Figure 2: This is the experimental set up for the horizontal setup. The flow goes from left to right. The methane and the air enter the premix chamber and enter the burner chamber through the ceramic honeycomb and then the flame. The carrier gas enters at the start of the system, this is where any suspension would enter the system. After exiting the end section, the flow will go through a downwards bend and into the collector.

3.3. Set Up for Vertical Experiments

The setup for the vertical experiments is nearly identical to the horizontal experiments. The only difference is that the first two preheating sections are lined up vertically, as opposed to horizontally. The flow is directed downwards through the atomizer and then through the first two preheating sections. Then a 90° bend directs the flow back to the horizontal plane, where the flow enters the third

preheating section. The rest of the setup is unchanged. This setup was used for all the powder producing experiments.

3.4. Set Up for Simulations

The simulations were conducted within the student version of the ANSYS Fluent application and on a laptop, and therefore the power of the simulations was limited due to the processing power of the laptop and the licensing limit for the student version of ANSYS. For fluid physics, the license limitation, as set by ANSYS, is 1 million cells and nodes [12]. Due to the processing power of the laptop used, none of the simulations conducted were anywhere near this limit.

The first step in setting up the simulation involves creating the structure for the fluid flow field. SolidWorks was used to recreate the structure that the fluid goes through in the system. Given that ANSYS Fluent was used for the final simulations, the structure made in SolidWorks was solid. The SolidWorks files were then imported into ANSYS Discovery where split lines were placed where different sections started and finished so that accurate boundary conditions could be used throughout the system. Once these sections had been created, the file was imported into ANSYS Mesh, where a mesh was made. Meshing is the process where shapes become elements that can be simulated by the software. If a mesh is inaccurate, the quality of the simulation can be seriously affected, so a key part of the setup was creating a good mesh quality. This was done by testing several element size values and measuring if the flow in the simulation matched that of the theory.

To determine the correct theory against which the simulation should be measured, the Reynolds number of the flow must be evaluated. To evaluate the Reynolds number, Equation 1 is utilized. Given that many different boundary conditions were tested, there was various Reynolds numbers within these tests. The highest fluid flow in this study was calculated at 0.2466m/s, which is when the carrier gas is set for 30l/min. The Reynolds number at this flow rate and at 20°C, which is the initial temperature of the fluid as it exits the atomizer and enters the first preheating section, was calculated at 831. Since Re<2300, the flow is judged to be laminar at the start of the preheating sections. The highest temperatures in the preheating sections reach 275°C, where the Reynolds number is calculated at 285; this is assuming the flow speed stays the same throughout, which it should do as governed by the assumption the system is closed and therefore the mass flow rate remains constant throughout the sections. The Reynolds number for the inner tube was calculated at 567. Given these Reynolds' numbers, it can be safely assumed the preheating sections and the inner tube can be modelled in the laminar flow regime. However, the whole system is not all laminar. A second stream is introduced at the start of the burner chamber. The two streams begin to mix throughout the burner chamber. This inherently causes turbulence in the flow because of the perpendicular pressures between the two streams. Therefore, the flow from the burner chamber onwards is no longer in the laminar regime, instead it is turbulent. Therefore, the simulation model needs to be separated into two: the laminar flow regime and a separate simulation for the turbulent regime. It is not possible to make a simulation that changes from laminar to turbulent flow in ANSYS, hence it is necessary to break the problem into two halves.

For the first half of the simulation, to check if the flow accurately follows the laminar theory, one must check the flow profile after the flow has fully developed, so the entrance length must be obtained. This length can be calculated by using Equation 2. The entrance length for the preheating sections would be 0.844m if the gases remained at room temperature, but given the Reynolds number decreases with temperature, as shown in the discussion above, so does the entrance length. Given the fluids get heated quickly by the preheating zones, a better estimate of entrance length can be made at 150°C, roughly halfway between the room temperature and our highest temperature of 275°C. The entrance length, calculated at 150°C gives us an entrance length of 0.45m, which is 18 inches.

$$v_{laminar} = 2 * v_{flow} * \left(1 - \left(\frac{r}{R}\right)^2\right) \quad (3)$$

Equation 3 shows the theoretical profile of a fully developed laminar flow. To test if a mesh was of good enough quality to be used for a full simulation, the mesh would be put in a simulation and the results would be plotted against the theoretical curve. It was determined that a mesh size with an element size of 5mm provided a good basis for the simulation. The laminar section of the simulation alone required approximately 200,000 elements and nodes, so it was not possible to get a much smaller element size within the licensing limit. It also provided simulation data that was comparable to the theoretical values within the laminar section; the simulation got 91% of the expected value of velocity in the fully developed section.

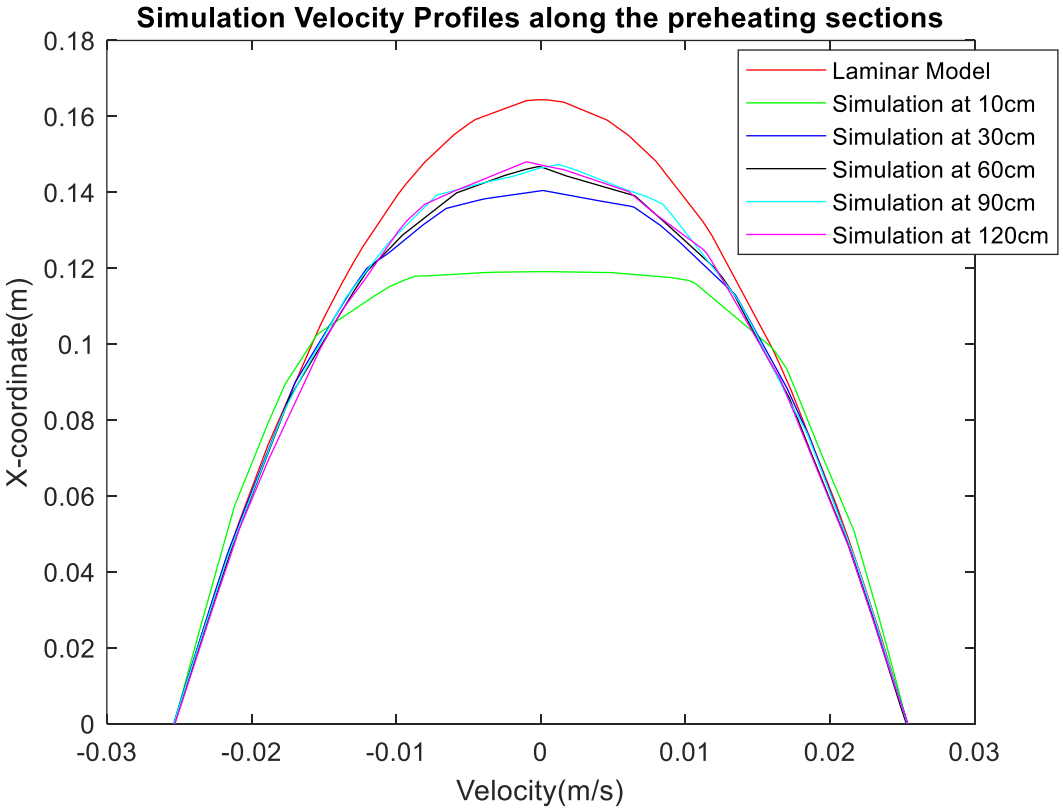

Figure 3: The velocity profile of air at various points along the preheating zones, when input into the first preheating section. The velocity of air is set at 0.0822m/s, equating to a flow rate of 10l/min. The preheating zones are at 150°C, 250°C and 275°C. The results show undeveloped flow at 10cm and 30cm in. The flow is developed at 60cm, 90cm, and 120cm, and for the rest of the preheating zone. The entrance length is therefore between 30cm and 60cm, the predicted entrance length is in this range of 45cm.

The graph above shows that the fluid acts in the way expected, albeit with some value error, probably from the suboptimal mesh, but the entrance length is as predicted. The velocity profile stays the same from 45cm onwards until the fluid is close to the reducer, where the diameter shrinks rapidly. As the fluid approaches the reducer, the decreased pressure in the reducer causes a pressure gradient force on the fluid near the end of the preheating zones. This force causes an increase in the velocity of the air and causes the fluid to deviate from the laminar profile. The problem with starting after this effect has occurred in the inner tube is that the flow in the inner tube never fully develops. The entrance length in the inner tube is calculated at over 6 inches and the inner tube is only 2 inches long. Therefore, the flow profile in the simulation does not follow the laminar profile theory at the very end of the final

preheating section, nor in the reducer or the inner tube. Given the simulation may not be 100% accurate due to various imprecisions in the mesh or any potential other inaccuracies in the simulation, it seems more sensible to take our starting point where the flow is still following the laminar theory. Therefore, the turbulent simulation starts near the end of the third preheating section. The turbulent section cannot start too far into the laminar section as that would create an inaccurate air flow profile before the inner tube and could therefore affect the accuracy of the simulation in the turbulent section. However, it must start within the laminar section, as it provides a predictable base from which to start the turbulent simulation from. Therefore, to create an accurate model for the turbulent section of this system, the place where this deviation occurs needs to be found. To check this phenomenon, the velocity profiles of the preheating section were checked near the start of the reducer.

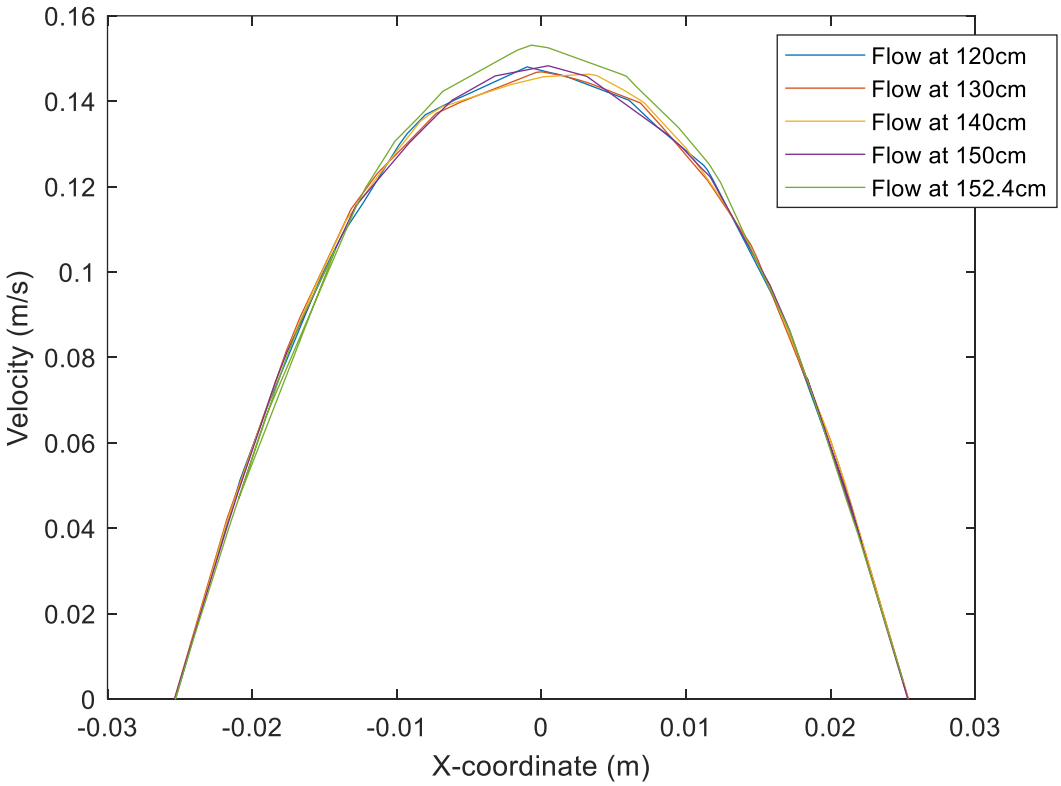

Figure 4: The velocity profile of the fluid towards the end of the preheating sections. The graph shows us that the pressure gradient force does not affect the flow in the first 150cm of the preheating zone, the effect starts somewhere within the final 2.4 cm of the preheating zone.

These results show that the turbulent section of the simulation should start with 2.4cm of the preheating zone left. The velocity was modelled to the power of 2 at 150cm in, and the temperature was measured to the power of 4. These models were input as the inlet values for the temperature and velocity for the turbulent section of the simulation.

$$v_{inlet} = a - b\left(\sqrt{x^2 + y^2}\right)^2 \tag{4}$$

$$T_{inlet} = a + b\left(\sqrt{x^2 + y^2}\right) + c\left(\sqrt{x^2 + y^2}\right)^2 + d\left(\sqrt{x^2 + y^2}\right)^3 + e\left(\sqrt{x^2 + y^2}\right)^4 \tag{5}$$

Equations 4 and 5 show the functions that model the velocity and the temperature 150cm into the laminar section simulation. The letters a to e represent various coefficients that are calculated separately for every simulation. The x and y represent the distance from the centre of the pipe in the x and y directions. They are combined to create a radial distance from the centre of the pipe, which means this equation can translate a two-dimensional model onto the three-dimensional simulation. The velocity was modelled to the power of 2 as that is the same power used in the laminar velocity flow equation, Equation 3 in this paper. This model makes sense as the velocity profile of the flow should follow a negative squared relationship. When the coefficients were tested against the simulation results, in MATLAB, the fit was found to be very appropriate, for the horizontal setup especially. For the vertical setup, there was a change in the velocity profile; the flow became asymmetrical in the z-plane. The flow in the bottom half of the pipe was quicker than that of the upper half; this effect was exacerbated more at greater velocities. The 90° bend after the first two preheating zones seems to have created this effect. This can be explained by Bernoulli's principle which states that for an incompressible and frictionless fluid, the total mechanical energy of a streamline is constant, which means that if the velocity increases, then the pressure must decrease. As these particles enter the bend, they tend to want to maintain angular velocity. This results in quicker flow on the outer radius of the bend, which leads to the bottom half of the pipe in the third preheating zone. This creates a lower pressure in this region. This lower pressure and higher velocity profile remains until the end of the third preheating zone. The effect almost gets cancelled out at the lowest flow rate simulated, at 10l/min, but remains prominent in the higher flow rates. To model this, a factor 4 model was attempted, with the same structure as Equation 5. This was to account for Bernoulli's principle affecting the simulation flow, and it created a seemingly better fit in two dimensions. However, when the model was transferred to the third dimension, the accuracy broke down and the factor 4 model suddenly became less accurate than the factor 2 model; the factor 2 model had the same structure as Equation 4. Therefore, the factor 2 model was maintained for the vertical setup, even if it was not ideal. A factor 4 model was used to model the temperature because the outer diameters of the flow are slower and therefore have more time to get heated by the preheating zones. This allowed the gas to reach a greater temperature as more energy could be used in heating the outer gases than the inner gases, especially since the flow is considered laminar in this section. Therefore, a model with a double peak, with a small trough was required, and a factor 4 model represented the simplest possible solution. The model also matched the data from the simulation very closely, which helped confirm the correctness of the model.

These models were calculated in MATLAB and input as the functions for the inlet velocity and inlet temperature, the inlet for the turbulent simulation being 150cm into the preheating sections. The simulation then travels for 2.4cm before it reaches the reducer, then the inner tube, and then the burner chamber. The burner chamber combines the two flows and has three temperature-measuring thermocouples. The flow from the premixing section comes into the burner chamber at different temperatures, depending on the ratio of the flow between the air and the methane. Therefore, an estimate for the temperature of the flame must also be made. To do this, the concept of adiabatic flame temperature was used. Adiabatic flame temperature is the maximum temperature that can be achieved with the given reactants. To calculate the flame temperature, the stoichiometry of the combustion reaction is needed, the energy released is needed, the specific heat of the products is needed, and the initial temperatures and conditions are needed.

$$CH_4 + 2O_2 \rightarrow CO_2 + 2H_2O + Energy \quad (6)$$

Equation 6 represents the chemical reaction of the combustion in this system. The system uses air, which is 21% oxygen, and methane gas. In every experiment conducted the oxygen was in excess and all the methane had combusted. To calculate the energy released, the moles of reacted methane were multiplied by its Lower Heating Value (LHV), which is 802kJ/mol. The LHV was used because the water vapor never condenses in the system due to the temperatures within the system.

$$Q = n_{CH_4} * LHV \quad (7)$$

Equation 7 represents the heat energy produced by the combustion, where n_{CH_4} represents the moles of methane that were combusted, and LHV represents the Lower Heating Value. To calculate the specific heat capacity of the gas mixture coming from the flame, the number of moles for each gas is needed. Therefore, the excess moles of oxygen were calculated and so too were the moles of the products of carbon dioxide and water vapor. The moles of unreacted nitrogen were calculated as 78% of the original moles of air. The specific heat of each gas was set at 1050K to give a first estimate of the specific heat capacity of the mixture. This was calculated by multiplying the number of moles for each gas by the specific heat capacity of each gas at 1050K. The products of these calculations were then added to each other, and then this sum was divided by the total number of moles in the mixture. This calculation provided flame temperatures that ranged from 1276K-1707K. These temperatures far exceed the initial 1050K used, so the process was redone at 1500K and then once more at 1400K to give a more accurate calculation of the specific heat capacities. The calculations done, at 1400K, produced the specific heat capacity values of the gas mixtures used in the simulations. All moles were calculated from 1 minute of flow.

$$C_{p,gas\ mixture} = \frac{C_{p,N_2}^{1400K} * n_{N_2} + C_{p,excess\ O_2}^{1400K} * n_{excess\ O_2} + C_{p,CO_2}^{1400K} * n_{CO_2} + C_{p,H_2O}^{1400K} * n_{H_2O}}{n_{N_2} + n_{excess\ O_2} + n_{CO_2} + n_{H_2O}} \quad (8)$$

Equation 8 shows the equation used to calculate the specific heat capacity for the mixture of gases coming from the flame, in the burner chamber. C_p represents various heat capacities, and n represents various molar amounts of different gases. Next, the mass of the gas per minute is calculated. This mass, in conjunction with the Q and C_p, from Equations 7 and 8, is used to calculate the adiabatic flame temperature.

$$T_{flame} = \frac{Q}{C_{p,gas\ mixture} * m_{gas\ mixture}} \quad (9)$$

Equation 9 shows the equation used to calculate the adiabatic flame temperature. Q represents the heat energy, calculated in Equation 7; $C_{p,gas\ mixture}$ represents the specific heat capacity of the gas mixture, calculated in Equation 8; $m_{gas\ mixture}$ represents the mass of the gas in 1 minuet of flow. The flame temperature calculated, from Equation 9, was then used as the inlet temperature for the burner chamber, from an inner diameter of 1 inch to an outer diameter of 3 inches. The velocity for this inlet was derived from the given mass flow rate of the gas mixture divided by the area of the burner chamber inlet. This alongside the calculated model input into the final 2.4cm of the preheating sections, using Equations 4 and 5, formed the boundary conditions for the turbulent simulation.

Table 1: This table shows the flame temperatures and the corresponding specific heat values at each combination of methane and air flow rate tested.

Methane Flow Rate (Liters/minuet)	Air Flow Rate (Liters/minuet)	C_p (kJ/kg-K)	Temperature (K)
1.0	15.0	1.338	1634
1.0	17.5	1.320	1465
1.0	20.0	1.307	1333
1.0	22.5	1.296	1229
1.1	22.5	1.304	1313
1.2	22.5	1.313	1395
1.3	22.5	1.321	1476
1.4	22.5	1.330	1556

In ANSYS Fluent, the energy equation was used due to the system's high-temperature changes, heat transfer, and combustion. The laminar model was used in the laminar section, and the k-epsilon model was used in the turbulent section of the model. The k-epsilon model was used as it provides a good balance between computational efficiency and accuracy. There are two equations in the k-epsilon model: the turbulent kinetic energy and the rate of dissipation of turbulent kinetic energy. The turbulent kinetic energy, k, models the energy in turbulent eddy currents and it measures the energy per unit mass associated with eddies in turbulent flow. The more chaotic the low, the greater the contribution of k in this model. The rate of dissipation, epsilon, describes how quickly the turbulent energy is dissipated into thermal energy because of the fluid's viscosity. It measures the decay of turbulence over time and distance. The k-epsilon model is used with enhanced wall treatment as it provides a smooth transition between the viscous sublayer and the fully turbulent region. Therefore, with wall treatment the boundary layer can be accurately modelled. A steady simulation was used, as transient simulations caused floating point errors. The computer used simply could not handle the extra complexity of a transient simulation. Each simulation was run until the mass flow rate and the facet average temperature both stabilized at a value; both were measured at the outlet. This ensured that both the temperature and the velocity of the flow were accurately simulated, whilst also ensuring the simulations had a finite ending. For these simulation measurements, a coupled pressure-velocity coupling scheme was chosen as this scheme provided the quickest convergence for both the mass flow rate and the facet average temperature. The complex geometry and two separate flows, ultimately, the reason why the coupled scheme provided quicker and more accurate convergence. The last settings input into the simulations were the temperatures and velocities. The velocities input into the inlet and the burner chamber inlet were calculated by the various flow rates divided by their respective inlet areas. The temperatures input for the 3 preheating zone temperatures were the temperatures used in the experiments, whilst the temperature from the flame inlet flow was the determined from the adiabatic flame temperature calculations. It was determined that pressure would not be applied as an outlet boundary condition as it only impacted the simulation by less than 0.5%, and it increased the complexity of the calculation for the computer.

4. Results

4.1 Results for changing carrier gas flow rate

The first experiment conducted was to vary the carrier gas flow rate, which enters the system at the start of the preheating zones. This test was conducted on both the horizontal and vertical setups, with the preheating zones set at 150°C, 250°C, and 275°C respectively, and with 1l/min of methane and 15l/min of air entering at the inlet located at the start of the burner chamber.

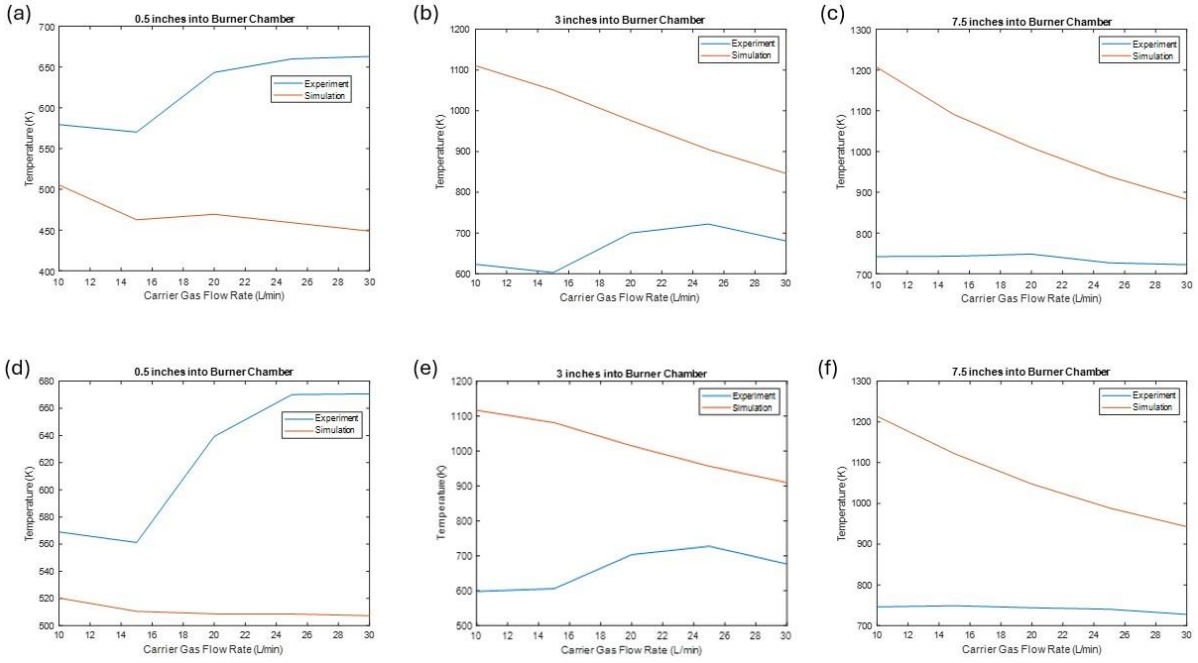

Figure 5: The following figure shows the temperature of the thermocouples at 0.5 inches, 3 inches, and 7.5 inches from the start of the burner chamber, all located in the centre of the pipe. (a)-(c) shows the results with the horizontal setup, whereas (d)-(f) shows the results with the vertical setup. The simulations show a decrease in temperature when the carrier gas flow rate is increased, this is not seen in the experiments.

The simulations show that increasing the carrier gas flow rate will decrease the temperature. When the carrier gas goes through the preheating sections at a greater velocity, the air is traveling in the preheating sections for less time. This means that the preheating sections put less heat energy into the carrier gas, giving it a decreased temperature, especially in the centre of the pipe where the flow is at its greatest velocity. This means that the flow coming through the inner tube into the burner chamber is at a lower temperature. Given that the first thermocouple is only 0.5 inches into the burner chamber, the flow coming from the inner tube barely interacts with the flow coming from the flame by the first thermocouple. This creates the effect seen on the two graphs at 0.5 inches. When the gases get to the next thermocouple, there will have been a substantial mixing of the two streams, hence why the temperature values for the thermocouple are at a much higher temperature, for the simulation, at around 1110K for both vertical and horizontal simulations, whereas the temperatures at the first thermocouple are from 450K-520K. The temperature loss from increasing the carrier gas flow rate also significantly increases, and the gradient becomes much steeper. The gradient for the horizontal simulation went from -2.8 K min/litre to -13.3 K min/litre. This is a substantial gradient decrease. There is a sharp temperature gradient between the two streams in the burner chamber, meaning that the

heat flux will be large. Therefore, if the cooler stream has an increased velocity, it will have significantly less time to be heated, leading to a much lower temperature. The effect of this is a more pronounced version of that seen in the first thermocouple. The third thermocouple is 7.5 inches, so much further downstream, and the same analysis applies as in between the first and second thermocouples. The gradient increases again to -16.2 K min/litre, for the horizontal section, and the temperatures are greater, at around 1210K for both the vertical and horizontal setups.

Given this analysis, it is surprising to see that the experimental results have not yielded a similar trend. The temperature significantly increases at the first thermocouple, as carrier gas increases, with an average gradient of +4.2 K min/litre, the opposite of the simulated trend. The average gradient is +4.2 K min/litre. The second thermocouple has a slightly shallower gradient of +2.9 K min/litre, whilst the third thermocouple has a flat gradient. In reality, the flow is not perfectly laminar, the flow from the end of the inner tube does not follow the laminar profile, as in Equation 4. Disturbances to the flow in the reducer cause small turbulent eddies in the inner tube, and consequently to the flow coming into the middle of the burner chamber. At higher velocities of carrier gas this effect should increase, which would increase the rate of mixing of the two streams. Given the stream from the flame has such great heat compared to the heat of the stream coming from the inner tube, any increase in mixing of the two streams will cause a noticeable increase in temperature. This explains why the increase of carrier gas affects the first thermocouple especially. Beyond that, the increased carrier gas flow rate helps decrease the heat loss to the environment, as the gases have less time to lose heat when they are moving at a greater velocity before they reach the thermocouples. Given the temperature gradient exceeds 1000K in some experiments, and the lack of insulation on the burner chamber, heat loss can be quite significant. Additionally, even though the streams are well mixed by 3 inches, they are not fully mixed, so the increased eddies from a higher carrier gas flow rate could also help increase the mixing at this point too, helping to increase the temperature at this point of the burner chamber. These two factors explain the increased flow rate of the carrier gas causing an increased temperature in the second thermocouple. The third thermocouple is much further downstream, meaning the two streams should be fully mixed, meaning the increased eddies from the inner tube section should have little to no impact on the third thermocouple, which is placed 7.5 inches into the burner chamber. Therefore, the measurement here is dominated by the heat lost to the environment and the proportion of hotter gas and colder gas entering the system, as the streams should have fully mixed. Increasing the carrier gas flow rate will decrease heat lost to the environment. However, given the carrier gas is much cooler than the stream from the flame, increasing its proportion will decrease the theoretical temperature in an adiabatic system. These two effects seem to have cancelled each other out in this experiment, leading to a flat gradient seen for the relationship between carrier gas flow rate and temperature at the third thermocouple.

The temperature is also greater than the simulated temperature at all carrier gas flow rates at the first thermocouple, pointing to the flow from the inner tube being higher than simulated, or that there is more heat flux between the streams at the start of the burner chamber than being simulated. A higher temperature inner tube flow could be the result of lower velocities in the preheating sections of the pipe than simulated. A greater heat flux could be possible if the flow coming out from the flame is not straight, meaning the air particles could start at the correct velocity to ensure a consistent mass flow rate, but could be coming out at many different angles. The simulation assumes the flow coming off the flame starts its path exclusively in one dimension, but this does not seem realistic given that particles coming off a flame may fly in all directions. This would result in some particles mixing with the flow from the inner tube much quicker than simulated. The second explanation seems more plausible given the temperature difference between the simulation and the experiment is over 100K. Additionally, the fact that in all simulations at the 3-inch and 7.5-inch thermocouple, the simulated temperature is much higher than the experimental temperature points to the fact that the adiabatic flame temperature calculated is greater than the actual flame temperature, or that there are significant heat losses to the environment in that period, or both. This indicates that the simulation should yield greater temperatures than the experimental results, or at least even results, at the 0.5-inch thermocouple as well. However, the opposite is true. The simulated temperature is greater than the

experimental temperature at all points of the burner chamber, except at the beginning of the chamber, which seems to indicate that the mixing of the gases is quicker than the simulations indicate, especially at the start.

The differences between the vertical and horizontal results seemed most pronounced for the simulation at the first thermocouple. This was due to the asymmetrical factor 4 model, like that seen in Equation 5, not mapping into the radial three-dimensional space well enough. This led to the use of a less accurate 2-factor model having to be used in simulation, as seen in equation 4, which could not consider the asymmetry of the flow, due to Bernoulli's principle. The modelled flow from the preheating sections only has an impact on the first thermocouple, hence why the only noticeable difference in these tests were those values. All other experimental and simulated results show great similarity between the two setups and the temperatures measured.

4.2 Results for changing the air flow rate

The second set of experiments conducted was to vary the air flow rate into the premix chamber, which enters the system at the start of the burner chamber, from 15 l/min to 22.5 l/min, in equal increments of 2.5 l/min. This test was conducted on both the horizontal and vertical setups, with the preheating zones set at 150°C, 250°C, and 275°C respectively, with a carrier gas flow rate of 20 l/min, and a methane flow rate of 1 l/min.

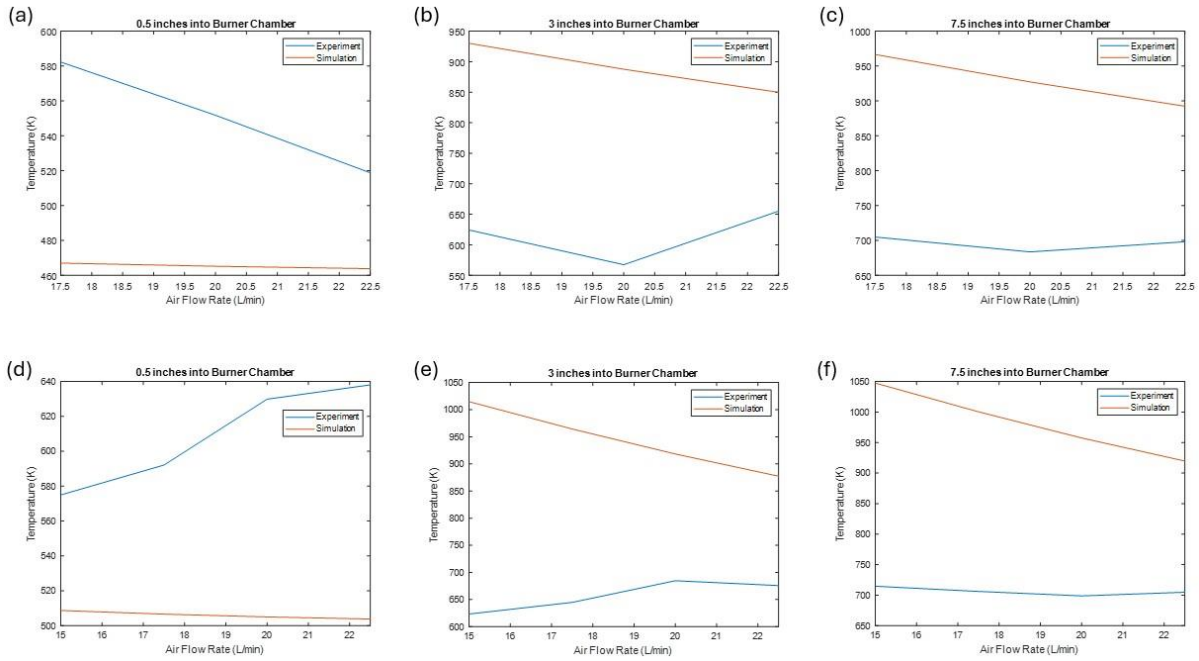

Figure 6: This figure shows the thermocouple temperatures at the stated distance into the burner chamber. (a)-(c) shows the results from tests done with the horizontal setup, whilst (d)-(f) shows the results from tests done with the vertical setup. The simulation showed that increasing the air flow rate will decrease the temperature in the burner chamber. The experimental results showed varied results.

As seen in Table 1, as the airflow rate increases, the temperature of the flame decreases. This is due to the increased mass of the reacted gas mixture being heated by the same amount of combusted methane, meaning the same amount of heat energy is heating a greater amount of air. This decreases the amount of heat the flame can add to the system, decreasing temperatures. This effect should be more pronounced in the second and third thermocouples when the two streams have had more time to

mix. As expected, this bears out in the simulations, both in the vertical and the horizontal setups. The gradient in the first thermocouple is -0.7 K min/litre, in the horizontal simulation and -0.8 K min/litre in the vertical simulation. This gradient increases for the second thermocouple to -16 K min/litre for the horizontal and -18 K min/litre for the vertical, a substantial increase. Both gradients slightly decrease for the third thermocouple, but the difference is very minor. The lower the flame temperature gets due to the increased air flow, the lower the simulated temperature gets.

The first thermocouple temperature of the horizontal setup experiment decreased with an increased air flow rate, at a gradient of -13 K min/litre. This is further evidence of greater mixing between the two streams towards the start of the burner chamber, as seen in the experiments conducted with changing the carrier gas flow rate. The decreased temperature of the flame had a significant impact on the temperature of the first thermocouple. However, in the vertical setup, the first thermocouple increased in temperature as the airflow rate decreased. This seems an odd result given the flame temperature is decreased. This is probably down to experimental error. After all, the temperature shift throughout the experiment is only 50K, which at these temperature levels is not an enormous difference. The relaxation time between changing the variables in this test was only 10 minutes, and within 10 minutes the system does not fully reach equilibrium. The temperatures in the experiment were greater across all air flow rates at the first thermocouple, once again pointing to greater mixing between the streams at the early stages of the burner chamber than being simulated.

The gradient in the experimental second and third thermocouples suggests that increasing the airflow temperature does not have a massive impact on the temperatures in the burner chamber. This is not in line with the simulations, but there are many reasons why this occurs. First of all, the flame is not adiabatic, no flame can be fully adiabatic. This would slightly decrease the range of flame temperatures in the experiment compared to the simulations, as the temperatures should scale down together with the same ratio. One advantage that an increased air flow rate brings is the increased overall flow rate of gas within the system. Given the stream is moving with greater velocity, there is less potential for heat loss. After all, the methane is still theoretically generating the same amount of heat energy, it is just being dispersed over a greater volume. The second and third thermocouples seem to be affected by two variables, the flame temperature and the amount of heat loss due to the environment. The two seem to have effectively equalized each other out.

The differences between the horizontal and vertical results are quite small, most simulation and experimental results were very similar for these two different setups.

4.3 Results for changing methane mass flow rate

The third set of experiments conducted was to vary the methane flow rate into the premix chamber, which enters the system at the start of the burner chamber, from 1 l/min to 1.4 l/min, in equal increments of 0.1 l/min. This test was conducted on both the horizontal and vertical setups, with the preheating zones set at 150°C, 250°C, and 275°C respectively, with a carrier gas flow rate of 20 l/min, and an air flow rate of 22.5 l/min.

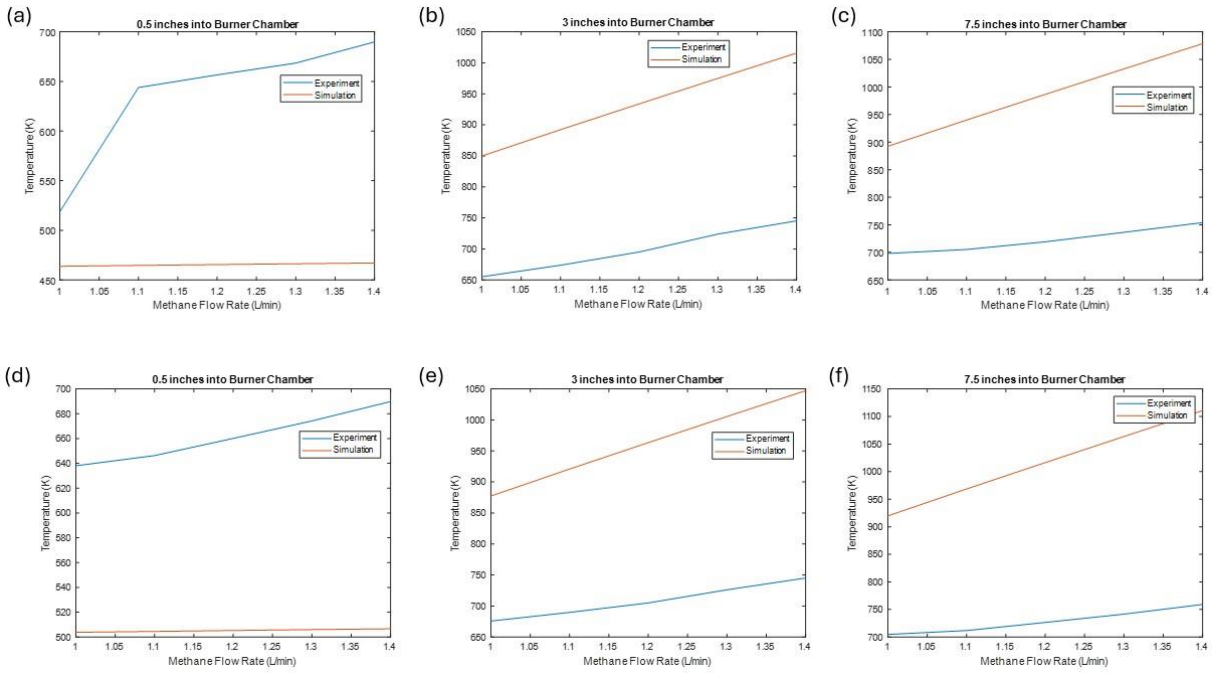

Figure 7: This figure shows the thermocouple temperatures at the stated distance into the burner chamber. (a)-(c) shows the results from tests done with the horizontal setup, whilst (d)-(f) shows the results from tests done with the vertical setup. The simulation showed that increasing the methane flow rate will increase the temperature in the burner chamber. The experimental results showed agreed with this positive trend.

As seen in Table 1, the expectation is that an increased flow rate of methane would increase the flame temperature. This would consequently increase the temperature of the entire system. The oxygen is always at a surplus in all data points trialled in this experiment. Because of this, additional methane means more combusted methane, in this experiment. Therefore, more heat energy can be generated to heat the same amount of air. This leads to a greater flame temperature.

This should result in a positive gradient, meaning the temperature should increase as the methane flow rate increases. The temperature should also rise the further down the pipe the flow is measured. This is due to the increased opportunity for the fluids to mix, resulting in higher temperatures measured at the centre of the pipe. The simulations agree with the expectation, that the temperature gets greater with a greater methane flow rate and with further progression into the pipe. The gradient is almost flat at the first thermocouple. The gradients of the second and third thermocouples are similarly positive at about 450 K min/litre. The third thermocouple was operating at a higher temperature than the second thermocouple, as expected. This was the case for both the horizontal and vertical setups.

The experimental results showed a very similar trend to the simulation results. The experimental results followed a consistent positive gradient in both setups, at all three thermocouples. In addition, the temperature was increasing the further downstream it was measured, as expected.

However, one key difference is that the temperature of the first thermocouple compared to the methane flow rate is greater than that in the simulations, this is also true for the gradient of the temperature. This points to the fact that the flame in the burner chamber has a greater impact on the flow coming out of the inner tube than the simulations suggest. This seems to be due to the flow coming out of the flame not being perfectly straight, therefore causing more mixing between the two flows at an early stage in the pipe flow than in the simulations. Another notable difference between the experimental results and the simulations is the significantly lower temperatures at the second and

third thermocouples. This is due to heat loss from the pipe to the environment and the flame not being perfectly adiabatic.

Once again, there is little discernible difference between the horizontal and vertical setups in both the experimental and simulation data, for this experiment. From all the tests done, both setups give very similar temperature profiles in the burner chamber. There are other aspects of the system that are affected by the bend, like the velocity profile, especially in the preheating zones, or the particle deposition rate on the pipe walls, but this cannot be measured with the apparatus available.

4.4 Changing concentration of NCM811 - powder production

The next experiment conducted was powder-producing. As discussed in the setup, a suspension is prepared and placed into a syringe connected to an ultrasonic atomizer. The variable changed in this experiment was the concentration of the suspension placed in the syringe. The flow rate was constant at 10 l/min, with the air flow rate at a constant 22 l/min and the methane flow rate at a constant 1.2 l/min. The three preheating zones were set at 100°C, 130°C and 275°C. Finally, the flow rate for the suspension was set at 0.6ml/min. The concentrations used were 0.1M, 0.25M, 0.5M.

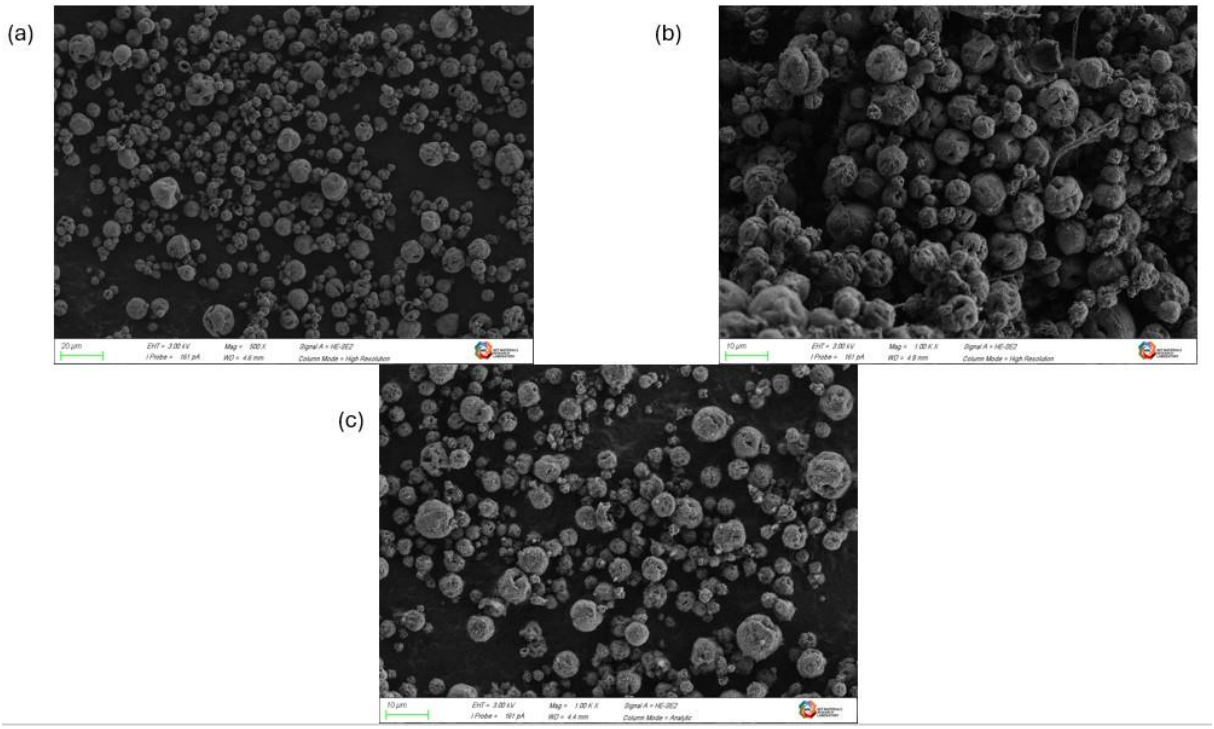

Figure 8: This is the imaging for the varying concentrations experiment. (a) shows the imaging from the 0.5M sample, with a scale bar of 20µm. (b) shows the imaging from the 0.25M sample, with a scale bar of 10µm. (c) shows the imaging from the 0.1M sample, with a scale bar of 10µm.

Upon inspection of the photo, all samples showed variability in particle size. This range was from 2µm to 13µm in the sample. The reason for such low variation is because of the ultrasonic atomizer, which provides a marked improvement versus 2-fluid atomizers, which create particles from the scale of nanometers to 100s of micrometers. There are other shared similarities between the samples which include having an imperfect spherical shape, with wrinkles and small cracklings at the surface. Every sample has a fluffy surface.

The main difference between the samples is the sizes of the particles. A clear distinction can be made between the sizes of the samples of 0.5M and 0.1M with just an eyeball inspection. However, given the diameter-to-volume ratio is x^3, it makes visual inspection difficult to distinguish the 0.25M and 0.5M samples, without statistical analysis. Therefore, statistical analysis was done using ImageJ analysis, and the size of the particles was measured. The 0.5M sample had a particle size of 10.2µm±2.2µm, the 0.25M sample had a particle size of 8.1µm±2.0µm, and the 0.1M sample had a particle size of 5.9µm±1.7µm. The 0.5M sample has a diameter 1.7 times the size of the 0.1M diameter. When this value is cubed, to estimate the volumetric ratio of the two samples, it equals 5.17, meaning the volume of the 0.5M sample is 5.17 times larger than that of the 0.1M sample. This is almost identical to the concentration ratio of these two samples, which is 5:1. The same analysis performed between the 0.25M sample and the 0.5M sample reveals the 0.5M sample has a size 2.00 times bigger than the 0.25M sample. This volumetric ratio is the same as the concentration ratio to 3 significant figures. These results show that an increase in concentration directly correlates to an increase in volume, meaning if the concentration was to be doubled, so too would the average volume.

4.5 Changing preheating temperature - powder production

The next experiment conducted was powder-producing. The variable that changed in this experiment was the temperatures of the preheating zones. The flow rate was constant at 5 l/min, with the airflow rate at a constant 22 l/min and the methane flow rate at a constant 1.2 l/min. The concentration was set at 0.25M. Finally, the flow rate for the suspension was set at 0.6ml/min. The preheating temperature combinations were 60°C, 80°C, 275°C; 80°C, 100°C, 275°C; and 100°C, 130°C, 275°C.

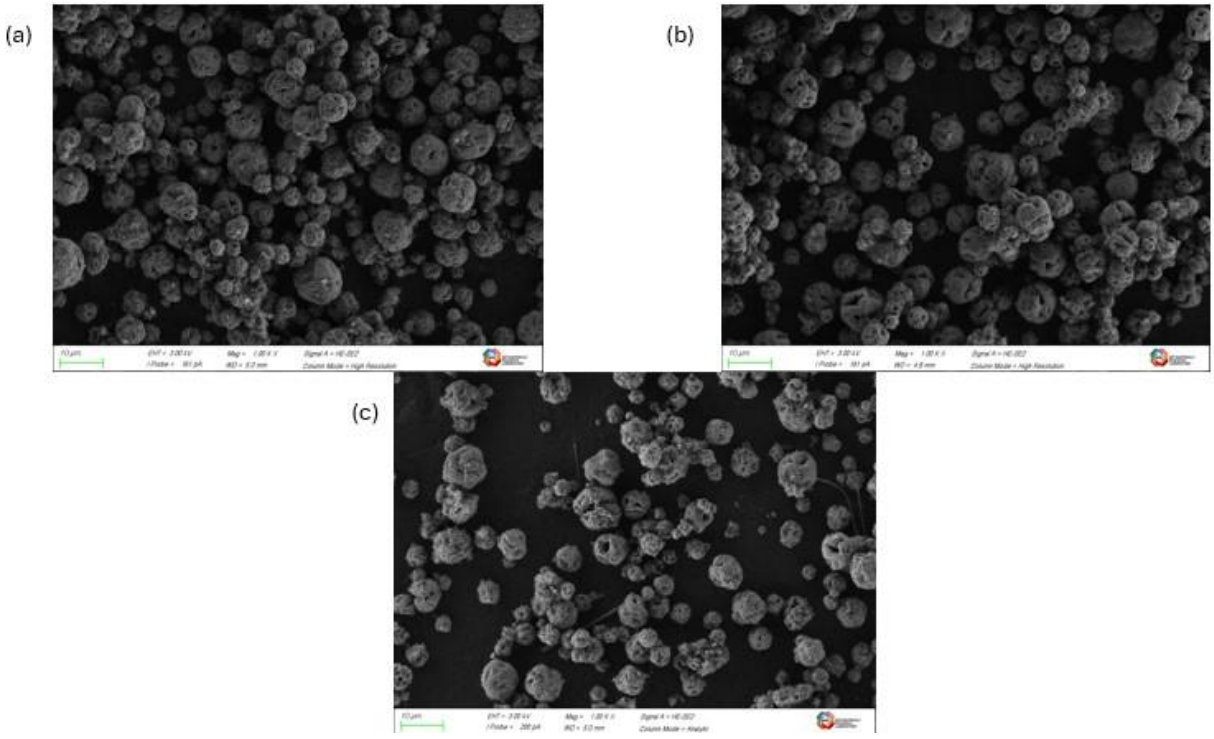

Figure 9: This is the imaging for the varying preheating zone temperatures experiment. (a) shows the imaging from the sample with temperature zones of 60°C, 80°C, and 275°C. (b) shows the imaging from the sample with temperature zones of 80°C, 100°C, and 275°C sample. (c) shows the imaging from the sample with temperature zones of 100°C, 130°C, and 275°C. All the scale bars in this figure are 10µm.

Given the concentration of the three samples is the same, it is no surprise to see that, upon inspection, the sizes of the three samples seem to be very similar. The particles range in size from 2µm to 12µm. As discussed in the background, the particle morphology is greatly impacted by the heating temperature. The three samples have folds and wrinkles, but the sample with the temperature zones of 60°C, 80°C, and 275°C has the least wrinkles and therefore it has the best-looking shape of the three samples. The other two samples seem to have similar crack sizes. The reason the sample with the lowest preheating zone temperature has the best shape is due to the slower evaporation rate, meaning moisture is more gently brought out from the surface to the environment. When the process is slowed, a flat and smoother surface is achieved. It appears that the heating rate was too quick once a temperature of 100°C or above was used, in the first two sections of the preheating zone. The evaporation of the moisture can happen below 100°C, and it seemed that keeping the temperature below the boiling point of water created a gentle enough release of the moisture to reduce the cracking and error on the surface. Another observation from this experiment is that the combination of 60°C and 80°C did not prevent the droplets from being full evaporated by the time the particles got to the third preheating zone. If that had been the case, the particles would have had significant cracking due to being heated at 275°C. This probably would have resulted in that sample having the worst morphology, instead of the best.

Another experiment was conducted with varying temperatures, apart from the flow rate was increased to 10 l/min. All the other constant variables remained the same, so the air flow temperature was 22 l/min, and the methane flow rate was 1.2 l/min. The concentration was set at 0.25M. Finally, the flow rate for the suspension was set at 0.6ml/min. Only two different temperature settings were trialled in this test. The two combinations tested were with preheating temperatures of 80°C, 100°C, 275°C and 100°C, 130°C, 275°C.

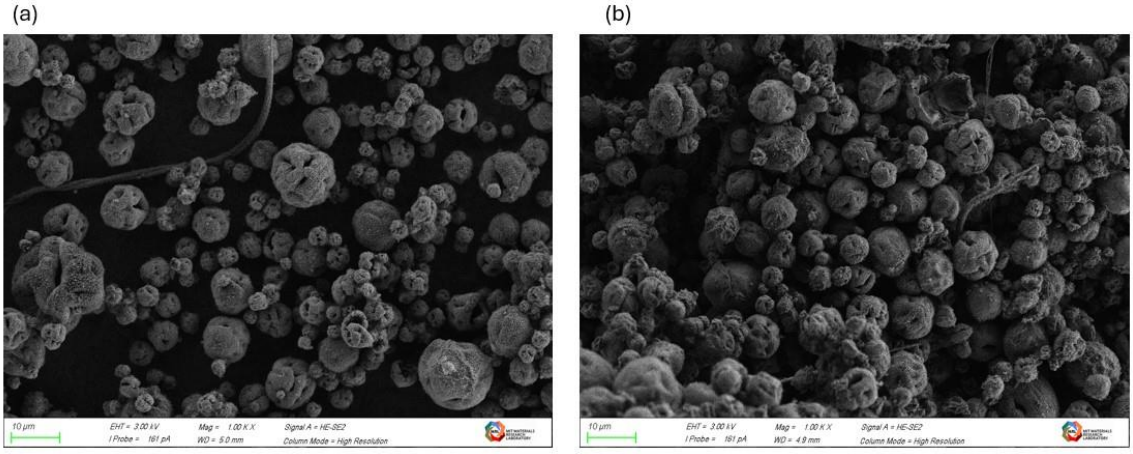

Figure 10: These images show us the particles produced from the experiments where the preheating temperature was varied. (a) shows the particles produced with preheating zones of 80°C, 100°C, and 275°C, whilst (b) shows the particles produced with preheated zones of 100°C, 130°C, and 275°C. The scale bars are set to 10µm for both (a) and (b).

The particles from the 80°C, 100°C, and 275°C trial have bigger variation than the other trial. There is a wider distribution of diameter. The particles in Figure 10(a) have a range from 2µm to 18µm, whereas the particles in Figure 10(b) have a range from 3µm to 12µm. The particles sizes, from the 80°C, 100°C, and 275°C trial, include many huge particles as well as many smaller particles. In contrast, the 100°C, 130°C, and 275°C trial have fewer big particles and fewer small particles, instead giving a much more tightly distributed set of diameters.

4.6 Changing carrier gas flow rate - powder production

The variable that changed in this experiment was the flow rate of the carrier gas, coming from the start of the preheating zones. The preheating zones were set at a constant temperature of 80°C, 100°C, and 275°C, with the airflow rate at a constant 22 l/min and the methane flow rate at a constant 1.2 l/min. The concentration was set at 0.25M. Finally, the flow rate for the suspension was set at 0.6ml/min. The two carrier gas flow rates used were 5 l/min to 10 l/min.

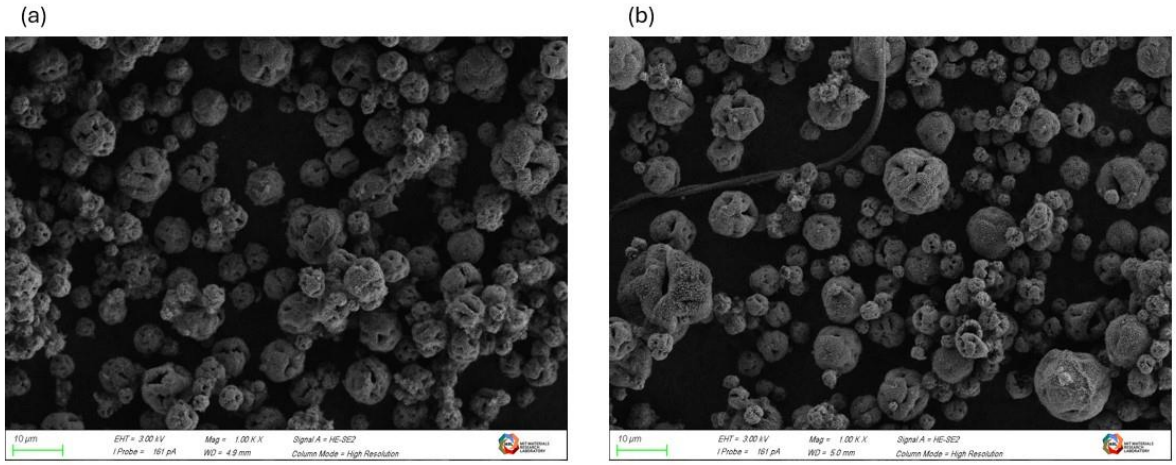

Figure 11: These images show us the particles produced from the experiments where the carrier gas flow rate was varied from 5l/min to 10l/min, with preheating temperatures at 80°C, 100°C, and 275°C for both tests. (a) shows the particles produced with a carrier gas flow rate of 5l/min, whilst (b) shows the particles produced with a carrier gas flow rate of 10l/min. The scale bars are set to 10µm for both (a) and (b).

The result shows very little difference between the two samples. The particles in both samples display a fluffy surface, with cracklings on the surface as well. Both sample sets also show big variances in particle diameter and also have similar sizes for average particle diameter. It seems that the preheating temperature and the concentration have a bigger impact on the particle size and morphology than the carrier gas flow rate does. The only way the carrier gas would affect the particle morphology is if the flow rate was too high, leading to the particle not having enough time to get rid of the moisture in the first two preheating zones. If this were the case, then the particle morphology would be negatively impacted and would have more cracks and a worse-looking surface, as the end of the heating would be done by the third preheating zone, at 275°C. That is not the case when working with the ranges of carrier gas flow rates that are in this experiment, which is from 5l/min to 10l/min.

An additional experiment was conducted with varying carrier gas flow rates, but the preheating zone temperatures were raised to 100°C, 130°C, and 275°C. All the other constant variables remained the same, so the air flow temperature was 22l/min, the methane flow rate was 1.2l/min. The concentration was set at 0.25M. Finally, the flow rate for the suspension was set at 0.6ml/min. Only two different temperature settings were trialled in this test. The two carrier gas flow rates tested were 5l/min and 10l/min.

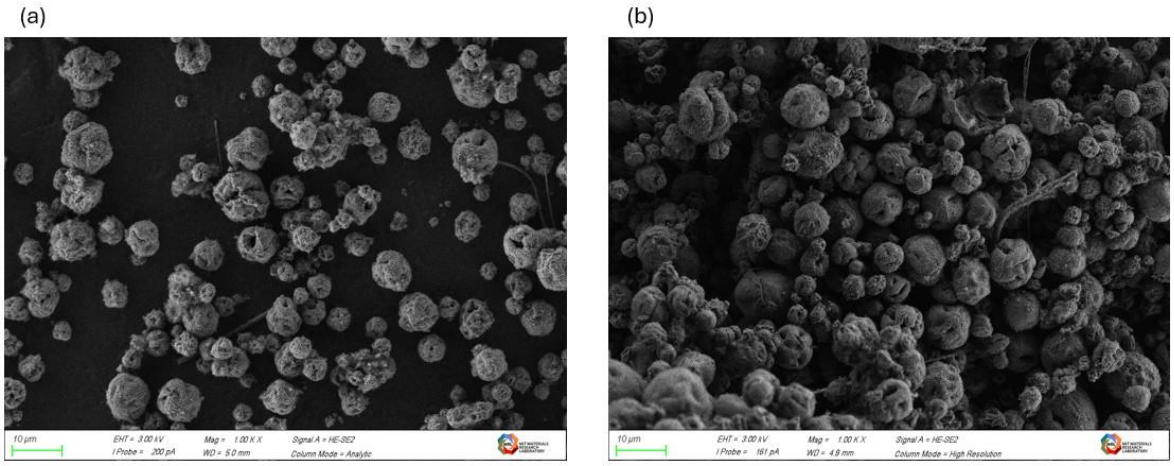

Figure 12: These images show us the particles produced from the experiments where the carrier gas flow rate was varied from 5l/min to 10l/min, with preheating temperatures at 100°C, 130°C, and 275°C for both tests. (a) shows the particles produced with a carrier gas flow rate of 5l/min, whilst (b) shows the particles produced with a carrier gas flow rate of 10l/min. The scale bars are set to 10µm for both (a) and (b).

The particles from both images look extremely similar, as was the case in the other carrier gas experiment. Once again, the carrier gas does not seem to have a big impact on the shape and the size of the particles. A much bigger effect is created when varying the concentration of the suspension or by varying the preheating temperatures, as seen in the other experiments. In the range of 5l/min to 10l/min, the carrier gas flow rate does not have a defining impact. As stated previously, if the carrier gas flow rate was to increase to a value that would prevent the total elimination of the moisture, that would have a significant impact on the surface quality of these particles. That is not the case, meaning a very little difference between the particles in these two images.

5. Conclusion

5.1 Discussion

There are many takeaways from the results of the experiments conducted in this study. For one, from the non-powder-producing experiments, the idea of turbulence from the inner tube flow ought to be considered when designing subsequent simulation models. The reducer's sudden reduction in diameter causes eddies to form. These swirls encourage the two fluids to mix in the burner chamber much more quickly than what happened in the simulations. These eddies are difficult to simulate due to their randomness, but the model was laminar in the inner tube anyway, which creates a divergence between reality and simulation. The simulations also had the velocity of the gas mixture, coming from the flame, in one dimension. This is also an incorrect assumption. There will be gas particles flying off in all directions coming off the flame, most of which will have a majority of their velocity in the forward direction but will also have some velocity in all three dimensions coming right off the flame. The model for the flame was made as a very hot inlet, rather than a flame, which causes less unpredictability in the direction of the gas particles coming off than a flame would. These two factors combined meant that the temperatures in the experimental thermocouple far exceeded the temperature of the simulated thermocouple.

Another key difference between the simulation and the experimental results was the factor of heat loss in the burner chamber. The burner chamber is not well insulated and operates at a very high temperature. This means a lot of heat loss will occur into the surrounding environment. The temperature gradient between the room temperature lab and the burner chamber can exceed 1000K at points, which is a huge difference. This factor of heat loss was not put into the simulation, causing massive temperature differences between the experimental results and the simulations at the second and third thermocouples, when the flow has travelled far enough for heat loss to become a major factor. Not much heat loss can occur in 0.5 inches, so the first thermocouple is largely unaffected by this phenomenon. This factor also meant that the higher flow rates resulted in greater temperatures than lower ones, as those flows experienced less heat loss to the environment than the slower flow rates. This meant that increasing the flow rate of the air into the premix chamber did not lower the temperatures measured at the second and third thermocouple, despite decreasing the flame temperature dramatically. The increased carrier gas flow rate also increased the temperature of the second thermocouple and did not decrease the temperature of the third thermocouple. That is even though increasing the carrier gas flow rate increases the proportion of colder gas in the system. The carrier gas has a temperature of around 500K at the end of the inner tube, the temperature of the air and methane mixture is 1634K, meaning the flow from the carrier gas is about 1100K cooler than the other stream. Despite this, the greater carrier gas flow rate increased the temperature at the second thermocouple and did not lower the temperature at the third thermocouple, where the effects of decreasing the heat loss, introducing a higher proportion of colder air, and increasing the turbulence coming out of the inner tube seemed to equalize.

Another factor that separated the simulation and the experimental results was the assumption the flame was adiabatic. Flames are never completely adiabatic, so this simulation assumption was always going to result in exaggerated flame temperatures. This combined with the heat loss factor to the environment meant that temperatures at the second or third thermocouples were always far lower in the experiments than in the simulations.

The horizontal and the vertical setups provided minimal differences in the experimental results, and very few differences in the simulations. The differences in the simulations should have been caused by an asymmetrical flow, caused by Bernoulli's principle, creating a different velocity and

temperature profile into the inner tube. Instead, the asymmetrical flow did not translate onto the radial model, as only symmetrical flows, with a maximum or minimum at the centre could be accurately modelled and input into the turbulent section accurately. This created a difference between two setups for a different reason, the shifted peak meant the values were slightly lower in the simulation than they should have been, but more importantly in the wrong place on the plane. This effect got exaggerated for both velocity and temperature at the greater carrier gas flow rates, in terms of peak magnitude and especially placement on the plane. This meant that the differences between the slower and quicker carrier gas rates was reduced in the vertical simulation compared to the horizontal simulation. The difference caused would have been small on the thermocouple readings, as demonstrated by the experimental results being very similar for both the horizontal and vertical setups. It would be interesting to test how other parts of the system are affected by changing the setup. For example, changing the setup would probably change particle deposition rates on the walls of the pipe flow and it most likely changes the velocity and temperature profile in the third preheating section, as seen in the simulation results.

For the powder producing experiments, it was not enormously surprising that increasing the concentration of the solution had a direct correlation on increasing the volume of the particles produced. An interesting and notable result was that the lower temperatures created a better-looking surface to the particle. It was notable that introducing temperatures of 100°C or above started to noticeably increase the level of deformation, specifically cracking, on the surface of the particles. It was also interesting that the moisture could be fully removed within the first two preheating zones for the test done with 60°C and 80°C as the temperature set for preheating zones 1 and 2. Otherwise, there would have been massive deformations on the surface due to extreme heat extracting the moisture at a much faster rate than wanted during the third preheating zone. The lack of impact of the carrier gas flow rate is, in hindsight, not surprising either. The moisture would have had to be not fully removed before the preheating zone for that to have a sizeable impact on the particle morphology. The reality is that the range the carrier gas flow rates were tested at was too low for impactful results for that experiment.

5.2 Future Directions

There are many variables that can be changed in both the experiment and the simulation as the system is extremely complex. First of all, the simulation model can be improved by creating a model that accurately transports the flow from the laminar section to the turbulent section if the flow is asymmetrical. Another step would be to calculate and include the heat flux from the burner chamber to the environment. This could be done to all sections of the system, but the largest heat loss likely lies in the burner chamber, so it would make sense to start there. Another change could be to increase the randomness of the flow coming off the flame. This would be a way to create a model for the flame inlet which would also help improve the accuracy of the simulation.

From an experimental standpoint, it would be interesting to test the impact of the carrier gas over a greater range, both for the powder-producing and non-powder-producing experiments. It would be interesting to see what carrier gas flow rate would be required to cause major cracking in the particles. It would also be interesting to see the effects that higher and lower flow rates than tested would have on the temperature of thermocouples. The same goes for the air flow rate, from the flame. Given the lack of surprising results produced by the varying methane flow rate experiment, it does not seem necessary to repeat these experiments with a wider range, as the results can be easily predicted. Specifically, for powder-producing experiments, it would be interesting to see if increasing the suspension flow rate into the atomizer has any impact on the quality of the particles produced. The larger the suspension flow rate, the greater impact the droplets have on the velocity and temperature

profiles within the system. A key piece would be finding the largest flow rate for this effect not to be too extreme so that good-looking particles are still produced. This would help increase the production rate of powder. In conclusion, many future steps can be made to fully optimise FASP considering the system's complexity. These further investigations could yield impactful results that could pave the way to faster and better production protocols that could ultimately lead to cheaper batteries for the average consumer.

References:

[1]: “Global Battery Market Size by Technology | Statista” [Online]. Available: https://www.statista.com/statistics/1339880/global-battery-market-size-by-technology/. [Accessed: 02-May-2024].

[2]: “Lithium-Ion Battery - Clean Energy Institute” [Online]. Available: https://www.cei.washington.edu/research/energy-storage/lithium-ion-battery/. [Accessed: 02-May-2024].

[3]: Khan, F. N. U., Rasul, M. G., Sayem, A. S. M., & Mandal, N. (2023). Maximizing energy density of lithium-ion batteries for electric vehicles: A critical review. *Energy Reports*, *9*, 11-21.

[4]: “What Is the NMC 811 Battery? What Are Its Features? | Battery Monday” [Online]. Available: https://www.grepow.com/blog/what-is-the-nmc-811-battery-what-are-its-features-battery-monday.html. [Accessed: 02-May-2024].

[5]: “Using Combustion to Make Better Batteries | MIT Department of Mechanical Engineering” [Online]. Available: https://meche.mit.edu/news-media/using-combustion-make-better-batteries. [Accessed: 02-May-2024].

[6]: Zhang, J., Muldoon, V. L., & Deng, S. (2022). Accelerated synthesis of Li (Ni0. 8Co0. 1Mn0. 1) O2 cathode materials using flame-assisted spray pyrolysis and additives. *Journal of Power Sources*, *528*, 231244.

[7]: Zhang, J., Muldoon, V. L., & Deng, S. (2023). Effects of the preheating temperature on flame-assisted spray pyrolysis of nickel-rich cathode materials. *Proceedings of the Combustion Institute*, *39*(1), 1165-1173.

[8]: Bhat, M., Luo, S., Zhang, J., Zhang, C., Zhou, B., & Deng, S. (2024). Multi-component precursor droplet evaporation in spray synthesis of cathode materials. *Chemical Engineering Journal*, *479*, 147417.

[9]: Handscomb, C. S., & Kraft, M. (2010). Simulating the structural evolution of droplets following shell formation. *Chemical Engineering Science*, *65*(2), 713-725.

[10]: “Chemicals and Solvents - Lab Pro Inc” [Online]. Available: https://labproinc.com/blogs/chemicals-and-solvents/the-importance-of-deionized-water-in-laboratory-applications. [Accessed: 02-May-2024].

[11]: Nii, S. (2016). Ultrasonic atomization. *Handbook of ultrasonics and sonochemistry*, 239-257.

[12]: “Download Ansys Student | Workbench-Based Simulation Tools” [Online]. Available: https://www.ansys.com/en-gb/academic/students/ansys-student. [Accessed: 02-May-2024].