A High-Order, Flow-Alignment-Based Compartmental Modelling Method

by

Alexandru Andrei Vasile

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Chemical Engineering

Waterloo, Ontario, Canada, 2025

© Alexandru Andrei Vasile 2025

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Bruno Blais Associate Professor Dept. of Chemical Engineering Polytechnique Montréal
Supervisor(s):	Nasser Mohieddin Abukhdeir Associate Professor Dept. of Chemical Engineering University of Waterloo
	Hector Budman Professor Dept. of Chemical Engineering University of Waterloo
Internal Member:	Marc G. Aucoin Professor Dept. of Chemical Engineering University of Waterloo
	Mario Ioannidis Professor Dept. of Chemical Engineering University of Waterloo
Internal-External Member:	Fue-Sang Lien Professor Dept. of Mechanical and Mechatronics Engineering University of Waterloo

Author's Declaration

This thesis consists of material which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

For Section 4.1 and Chapter 5, the author developed and implemented the documented method as well as obtained the numerical results. Marc G. Aucoin, Hector Budman, and Nasser Mohieddin Abukhdeir were involved in the conceptualization of the work. All authors read and approved the final manuscript.

For Chapter 6, Ittisak Promma set up and ran the CFD simulations in `OpenFOAM` that were used as input for the work. Matthew Peres Tino implemented the initial version of the reaction parsing used to obtain the oxygen dissolution, the author heavily modified and expanded upon it. The author did the rest of the derivations, simulations, and analysis on their own.

For Chapter 6 and Chapter 7, Yuvraj Aseri implemented the initial version of the code for converting the `OpenFOAM` results and meshes into the `OpenCCM` format, the author modified and expanded upon it. The author did the rest of the derivations, simulations, and analysis on their own.

Abstract

Industrially-relevant chemical engineering processes, such as stirred tank bioreactors in the pharmaceutical sector, inherently operate across multiple scales and involve complex, multiphysics, and multiphase interactions. Modelling of these systems is essential for their design, optimization, control, and operational troubleshooting; these processes are often too intricate for experimental approaches alone, with trial runs proving prohibitively costly or key metrics being difficult or impossible to measure. Traditionally, modelling such systems has relied on simplified design equations or idealized models, such as the continuously stirred tank reactor (CSTR). However, these approaches lack the explanatory power required to capture real-system outcomes, such as concentration gradient formation. With advancements in computational capabilities, computational fluid dynamics (CFD) simulations have become standard for investigating specific questions within these systems. Nonetheless, certain critical applications, such as extended simulations of microorganism growth or real-time predictive control, remain impractical due to their high computational demands.

Reduced Order Models (ROMs) offer a middle ground between the simplistic CSTR models and the computationally intensive CFD simulations. ROMs trade off some of the generality and accuracy of CFD simulations in exchange for a substantial reduction in computational cost, often by several orders of magnitude. This work focuses exclusively on a specific type of ROM: Compartmental Models (CMs). They are underpinned by the assumption of one-way coupling between the hydrodynamics and mass transport of reactive species. CMs are constructed through a two-step process. First, the domain is divided into non-overlapping compartments using a set of criteria; next, each compartment is represented by one or more simplified models. This network of models decouples mass transport from hydrodynamics and reduces the number of degrees of freedom on which the conservation of mass of the reactive species needs to be solved. This reduction is particularly important for bioreactors, where hundreds of coupled nonlinear reactions are common.

Current compartmental modelling methods exhibit several limitations, such as a disconnect between the criteria used for compartment identification and their subsequent modelling, an assumption that each compartment is well-mixed, a reliance on manual compartmentalization or manual intervention, and a non-prescriptive framework that is challenging to adapt to new geometries. This work introduces a novel compartmental modelling method based on flow alignment. The velocity field is analyzed and split into compartments within which the flow is unidirectional. Each compartment is then modelled as a series of 1D Plug Flow Reactors (PFRs). Benchmarking this method against the

state-of-the-art method demonstrates that it yields more accurate results while achieving computational speeds that are orders of magnitude faster than traditional CFD simulations.

Further, many current CM approaches simplify three-dimensional geometries by either modelling two-dimensional cross-sections and relying on rotational symmetry or by using a uniform grids of compartments. The developed method is extended to fully three-dimensional two-phase stirred tank systems without using these assumptions. It successfully compartmentalizes the distinct recirculatory regions generated by the impellers, eliminating the manual *ad hoc* intervention required by past methods. Mixing time and concentration predictions at probe locations are validated against CFD simulations, other CMs, and experimental data. The proposed *general* method performs as well or better than past CMs which were tailor made for the stirred tank geometry. Further, the model's capability to handle complex spatially varying reactions is demonstrated by simulating oxygen dissolution into the liquid phase, accurately capturing spatial gradients in dissolved oxygen concentration.

Lastly, a significant limitation in previous compartmental modelling work is the reliance on a single velocity snapshot or a time-averaged steady-state velocity field. For instance, in the case of vortex shedding from a cylinder in the laminar flow regime, neither time-averaged velocity-based CM nor an ensemble of CMs based on discrete velocity snapshots accurately captures the impact of the inherently non-stationary flow topology. The non-stationary nature of such flow fields is addressed by employing projection mappings to cycle through a series of compartmental models, allowing dynamically updating their shape, number, location, and connections. This approach successfully captures the oscillation period of the flow and demonstrates promise in representing non-stationary flow behaviours accurately.

In summary, this work advances the field of compartmental modelling by unlocking their the application to complex, industrially-relevant systems by developing a generalized, alignment-based method. This method extends the capability of CMs to handle both time-varying and fully three-dimensional multiphase flows without requiring manual intervention. The approach is validated through benchmarking against CFD simulations, other CM approaches, and experimental data, demonstrating improvements in computational efficiency and accuracy.

Acknowledgements

First and foremost, I would like express my deepest gratitude to my fiancé and my family, without whose unwavering encouragement and support none of this would have been possible.

I extend my sincere thanks to my supervisors, Professor Nasser Mohieddin Abukhdeir and Professor Hector Budman, for their guidance, mentorship, and trust. I am particularly grateful for their flexibility and understanding throughout this journey, which enabled me to complete this work despite the challenges of these past years.

A big thank you to Ittisak Proma, Arshia Fazeli, James Lowman, and Matthew Tino for all of the technical discussions we have had over the years. Thank you to past and present members of the COMPHYS group for all of our discussions, shared meals, and camaraderie throughout the years.

I would like to express my appreciation to the Department of Chemical Engineering at the University of Waterloo for providing exceptional opportunities and support throughout both my undergraduate and graduate studies. I am particularly indebted to Judy Caron, whose patience, dedication, and expertise have been invaluable.

I would like to thank my committee members, Professor Bruno Blais, Professor Fue-Sang Lien, Professor Marc G. Aucoin, and Professor Mario Ioannidis for their feedback on this work.

Lastly I would like to thank Sanofi Pasteur Limited, the Natural Sciences and Engineering Research Council of Canada (NSERC), and the University of Waterloo for their financial support and the Digital Research Alliance of Canada for the use of their computational resources.

Dedication

I dedicate this thesis to my fiancé and my family; you mean the world to me.

Table of Contents

Examining Committee ii

Author's Declaration iii

Statement of Contributions iv

Abstract v

Acknowledgements vii

Dedication viii

List of Figures xiii

List of Tables xxiii

1 Introduction 1

1.1 Research Motivation 1

1.2 Objectives 5

1.3 Structure of Thesis 6

2 Background 7

2.1 Ideal Versus Non-ideal Stirred Tank Behaviour 7

2.2 Single-Phase Computational Fluid Dynamics . . . 10
2.3 Multi-phase Computational Fluid Dynamics . . . 11
2.3.1 Interface-Capturing Models . . . 12
2.3.2 Two-Fluid Model . . . 14
2.3.3 The Euler-Lagrange Model . . . 16
2.4 Turbulence Modelling . . . 18
2.4.1 Single-Phase Simulations . . . 18
2.4.2 Multi-Phase Simulations . . . 21
2.5 Impeller Motion . . . 24
2.6 Reduced Order Modelling . . . 27
2.7 Compartmental Models . . . 28
2.8 Numerical Methods . . . 31
2.8.1 Coefficient of Mixing . . . 31
2.8.2 Finite Element Method . . . 32
2.8.3 Finite Volume Method . . . 34
2.8.4 Finite Difference Method . . . 35
2.8.5 Differential Algebraic Equations . . . 36

3 Literature Review **40**
3.1 Compartmentalization . . . 40
3.2 Modelling of Compartments . . . 46
3.3 Dynamic Compartments . . . 48
3.4 Other Topology-Based Methods . . . 50
3.5 Handling of Zero-flux Planes . . . 53

4 Compartmental Modelling Methods **57**
4.1 Method for Single-Phase 2D Systems . . . 57
4.1.1 Pre-processing . . . 58

4.1.2 Compartmentalization . . . 59
4.1.3 Compartment Modelling . . . 64
4.1.4 Validation . . . 69
4.1.5 Kullback-Leibler Divergence . . . 70
4.2 Extension to 3D Multi-Phase Closed Systems . . . 71
4.2.1 Convection-Guided Compartment Merging . . . 71
4.2.2 Switch to Facet Fluxes . . . 72
4.2.3 Extension to Multi-phase Flows . . . 73
4.2.4 Handling of Closed Systems . . . 74
4.2.5 Element-to-DOF and DOF-to-Element Mappings . . . 75
4.3 Multi-CM Methods for Transient Flows . . . 77
4.3.1 Flow Snapshot Selection Method . . . 78
4.3.2 Snapshot Ensemble CM Method . . . 80
4.3.3 Dynamic CM Projection Method . . . 81

5 Application to Representative 2D Domain **84**
5.1 CFD Simulation . . . 84
5.2 Compartmentalization . . . 87
5.3 Compartment Modelling . . . 89
5.4 Conclusions . . . 96

6 Application to a 3D Two-Phase Stirred Tank Bioreactor **97**
6.1 Stirred Tank Bioreactor . . . 97
6.2 Compartmentalization . . . 101
6.3 Mass Transfer Simulation . . . 103
6.4 Interphase Mass Transport Simulation . . . 107
6.5 Flow Oscillations . . . 111
6.6 Conclusions . . . 113

7 Accounting for Time-Varying Flows **114**
7.1 Chosen Geometry . . . 115
7.2 CFD Simulation . . . 116
7.3 Snapshot Subsampling . . . 117
7.4 Compartmental Model Simulations . . . 119
7.4.1 Time-Averaged Compartmental Modelling Approach . . . 119
7.4.2 Ensemble-Based Compartmental Modelling Approach . . . 121
7.4.3 Dynamic-Projection-Based Compartmental Modelling Approach . . . 122
7.5 Conclusions . . . 125

8 Conclusions and Recommendations **127**
8.1 Conclusions . . . 127
8.2 Recommendations for Future Work . . . 129

References **133**

APPENDICES **143**

A Implementation Details of Method in OpenCCM **144**
A.1 Summary . . . 144
A.2 Statement of Need . . . 145
A.3 Features . . . 146
A.4 User Interface . . . 146
A.5 Reaction Config File . . . 147
A.6 Examples of Usage . . . 148
A.6.1 RTD Curves . . . 150
A.6.2 Reactions . . . 150

List of Figures

1.1 Illustration of the fundamental approach of mesh-based compartmental modelling. A CFD mesh is systematically divided into five distinct compartments, with the division occurring at the element level. Each compartment is then represented as a continuously stirred tank reactor (CSTR), and these compartments are interconnected to form a network that approximates the original flow domain. From Ref. [20] 4

2.1 Representative geometries of the two extremes in bioreactor sizes showing that they feature the same basic elements: narrow aspect ratio, immersed impeller, and a sparger. Also notable is the inclusion of baffles on the larger reactor to aid with mixing, a feature not needed on the smaller reactor. **(a)** is a 15 mL Ambr microbioreactor from [2]. **(b)** is a 30 m^3 industrial bioreactor from [3]. 8

2.2 Illustration showing three different flow regimes of interest for multi-phase simulations. Grey is used for the gas phase while blue is used for the liquid phase. **(a)** Liquid phase continuous and gas phase dispersed. **(b)** Liquid phase continuous and gas phase dispersed in one region of the domain and continuous in another. **(c)** Both phases continuous, each confided to different regions of the domain and separated by a large contiguous interface. 11

2.3 Overview of VOF methods. **(a)** The exact interface overlayed by the numerical grid and the values of the indicato rfunction. **(b)** Simple Line Interface Calculation (SLIC) reconstruction of the interface. **(c)** Piecewise Linear Interface Calculation (PLIC) reconstruction of the interface. $\boldsymbol{m} = -\nabla C$ is the outward facing unit normal. Showing how the surface is reconstructed from the field. From [5]. 13

2.4 Illustrations representing the most popular interphase momentum transfer terms for which closures are used to approximate $\boldsymbol{\mathcal{M}}_i$. Created by Arshia Fazeli, used with permission. 15

2.5 Illustration of an Euler-Lagrange simulation using DEM particles to simulate the slurry-slurry, slurry-sand, and sand-sand interactions. While computationally intensive, such approaches are required when flow between particles, as well as their interactions, are crucial to achieving the goal of the simulation. From [31]. 17

2.6 Overview of the different approaches to single-phase turbulence modelling, showing the spectrum of approaches available ordered by computational intensity and accuracy. From [33]. 19

2.7 Overview of a stirred tank reactor mesh that has been prepared for either an MRF or an SM simulation. **(a)** Overall view of the tank and its mesh. **(b)** Close-up view of the mesh used for the rotating impeller domain. **(c)** View of the bottom of the tank and core of the structured mesh. From [12]. 25

2.8 Snapshots of a transient SM-based simulation of an inert tracer injected into a multi-Rushton impeller tank. The figure focuses on the inter-impeller region, with a Rushton impeller being located at the top and bottom of the frame. The macroinstability resolved by the SM method is seen in the middle of each frame as it oscillates and precesses. With an MRF-based simulation, that plane would be horizontal and fixed. From [40]. 26

2.9 Example of a POD-based ROM for the hydrodynamics of a flow-past a cylinder scenario. **Left:** Labelled solution manifold showing the time-evolution of the coefficient to the 3 POD modes, with certain key locations labelled. **Middle-left:** Three different flow snapshots labelled by their position on the manifold. **Middle-right:** The three most important modes used to simulate the system. **Right:** The trajectory of the system in the state-space of the three modes for both the original CFD simulation and the ROM simulation based on the three most important modes. From [41]. 27

2.10 Creation of the compartmental model from the CFD mesh. **(a)** The initial CFD mesh from which the compartments are to be made. **(b)** The result of the first step, each mesh element is coloured based on the compartment that it belongs to. **(c)** The result of the final step, each compartment is modelled as a single CSTR and a network is created based on the compartment connectivity. 29

2.11 A wastewater treatment pond manually split into 3 compartments, referred to as zones in the figure, based on expert analysis of the flow field.. Each compartment is then modelled as a network of CSTRs, with connections between CSTRs, both within and across compartments, also based on expert knowledge of the flow. From [51]. 31

2.12 Comparison of a CG (top) and DG (bottom) approach capturing the same known solution, each using linear interpolants. The discontinuities in the DG version allow it to better approximate the nonlinear solution, but they will require special treatment for evaluating the value at the element bound. From [60]. 33

2.13 Diagram of the pendulum system in Cartesian coordinates. x and y are position, u and v are velocity components in x and y, g is acceleration from gravity, θ is the angle between the rod and the vertical, and λ is the tension per unit length of the rod. From [62]. 37

3.1 Illustrative example of a structured compartmentalization in a cylindrical tank. **(A)** Radial-axial plane. **(B)** Radian-tangential plane. **(C)** Each internal compartment is connected to 6 others. From [54]. 41

3.2 Illustration of the manual custom-made compartmentalization of the stirred tank used by Nauha et al. from different angles. **Left**: vertical cross-section of the compartments showing the overall structure. **Top-right:** horizontal cross-section at the plane of the impellers, showing special irregular compartments for each impeller blade. **Bottom-right:** horizontal cross-section in all other parts of the tank. The compartments are large toroids that span the full 360° of the reactor. From [48]. 43

3.3 Overview of the automatic compartmentalization of a stirred tank reactor using the method of Tajsoleiman et al. **(a)** The geometry of the domain showing 3 bottom mounted Rushton turbines and 4 baffles. **(b)** Zones based on changes in radial velocity. **(c)** Zones based on changes in axial velocity. **(d)** Compartments resulting from intersecting the zones in (b) and (c). From [24] 45

3.4 Manual compartmentalization of the flow profile in a burner. **Left:** Three main regions are identified based on a visual inspection of the flow profile. **Right:** Each region is split into multiple compartments, and each compartment is modelled by one or more CSTRs/PFRs, with volumes and flowrates fitting to simulated data. From [50]. 47

3.5 Illustration of how changing reactor working volume can be incorporated into a CM method. CFD results are available at 40, 60, and 90 m^3 and used to create CM. For volumes in-between, the CM is created by interpolating between the two nearest CFD-based CM. Current volume is tracked, and the compartmentalization is changed only once the volume has changed by a set amount from its starting value. From [77]. 49

3.6 Inert tracers a and Tr are used to identify compartments and choose how many CSTRs-in-series to model each compartment with. Tr is injected with a unique value from each domain inlet, allowing for regions of influence to be identified. Streamlines of constant value Tr are used to delineate compartments. a is used to quantify the age of the fluid, and split each streamline into pieces of roughly uniform age. From [79]. 51

3.7 Cross-section of the CFD and compartmentalization results for flow in a pipe with a sudden expansion. **(a)** Flow topology visualized through velocity streamlines coloured by velocity magnitude. **(b)** Resulting compartments: orange - inlet, green - outlet, purple - splay, red - recirculation, blue - stagnation zone. From [21]. 52

3.8 Streamlines visualizing flow topology of the same geometry under 4 different conditions. Critical points used for compartmentalization are labelled as 1–19. From [80]. 53

3.9 Visualization of the flow topology, using a line integral convolution, of the region between two Rushton impellers from a stirred tank MRF-based CFD simulation. The dashed horizontal line in the middle of the figure indicates the horizontal plane that forms at which the velocity, and thus convective transport, is in the horizontal direction. The yellow and orange regions represent two different CSTR-based compartments which span either side of this plane. The use of the CSTR model for each compartment causes "short-cutting", i.e. causes mass to move artificially quickly across the plane. From [12]. 54

3.10 Structured compartmentalization of the recirculatory region around the top impeller. The distribution of compartments, their connections, and flowrates are obtained from design equations and expert knowledge of the system. From [3]. 55

4.1 Projection of the higher-order FEM result onto a zero-order interpolant for a one-dimensional mesh. The original higher-order FEM interpolant is shown in blue while the final zero-order interpolant is shown in orange. **Left:** Projection on the original mesh. **Right:** Projection on the refined mesh. . 58

4.2 Rejecting a neighbour based on the angle between the direction vector of the compartment's seed element and that of the neighbour element. The blue and orange arrows represent the direction vectors of the seed and neighbouring elements, respectively. 60

4.3 Rejecting a neighbouring element, B, because it is upstream of the compartment. The blue arrow is the velocity vector of an element a in the compartment, and the green arrow is the outward facing unit normal to the purple facet between a and b. 60

4.4 Example showing if the flow from a to b is adequate. Blue arrows on the left represent the velocity vector for each mesh element. On the right, the blue arrows represent the flowrate through each facet of the elements, labelled as Q1–Q7. 61

4.5 Identifying the bounding facets for compartment A. The facets of all elements in A are collected. Internal facets will appear twice, once for each element in the compartment it shares. Bounding facets are identified as those appearing only once. 62

4.6 Splitting the bounding surface of compartments A through E based on the compartment on the other side. **Left:** Compartments identified. **Middle:** Bounding facets for each compartment identified in green. **Right:** Bounding facets split grouped, shown using colour, based on which two compartments they are between. 62

4.7 **Left:** Compartment network shown as labelled regions in physical space. **Right:** The same compartment network represented as an undirected graph. 63

4.8 Illustrative example of the merging of compartments. After compartment B is merged into A, compartment C is left with a single connection, with A. This is unphysical as mass cannot be balanced around C, thus compartment C is then also merged into A. 63

4.9 Illustrative example of splitting the surface between compartments A and B based on the sign of the flux. The sub-surface between points L and M has net flow from $A \to B$, the portion from M to N has no net flow and was removed, and the portion from N to P has a net flow from $B \to A$. . . 64

4.10 Illustrative example of the computation of spatial ordering of compartment connections. **1.** Compartment with each sub-surface represented by a different colour. Arrows indicate the direction of flow through the surface. d_{avg} is the average direction vector of the compartment. **2.** Spatial distribution of flow through each surfaces. **3.** Flow-weighted representative position of each surface. **4.** Position vector of each surface connected to average direction vector of the compartment. **5.** Position vectors projected onto the average direction vector. **6.** Dot-product results from 5. normalized from 0.0, most upstream, to 1.0, most downstream. Upward facing arrows represent inlets into the compartment, and downward facing arrows represent outlets. . . 65

4.11 Illustration of the resulting PFR-in-series from a compartment with 5 connections, 3 of which are spatially close to each other. **Left:** Without merging, this configuration produces 4 PFRs, 2 of which are significantly smaller than the rest. **Right:** With merging, this configuration produces 2 PFRs, halving the number of degrees of freedom for this compartment. 68

4.12 The construction of a PFR sequence used to model a compartment with 4 connections, 2 inlets and 2 outlets, with different spatial locations. The result is a sequence of 3 PFRs. 68

4.13 Simplified example showing the different merging result between (a) the original method and (b) the new method. Green represents the compartment chosen for merging and red represents the rejected one. **(a)** The original method chooses compartment B despite the insignificant amount of flow between A & B. **(b)** The new method chooses compartment C, despite the higher angle difference, because of the much higher flow between them. . . 71

4.14 Illustration of two adjacent mesh element with different velocity vectors in each. If the element's velocity vector is used for calculating the flow through the red facet, then different values are obtained if calculated from element A versus element B. 72

4.15 Calculation of volumetric flowrate for different flow regimes through the red facet on the right side of each domain. u_o represents the local velocity of the orange phase, and u_b the local velocity of the blue phase. **(a)** For a single phase system the calculation is straight forward. **(b)** For a stratified two-phase system the calculation requires an adjustment to the surface area used in the calculation since the orange phase does not **(c)** For a dispersed two-phase system the area must again be adjusted to account for the presence or absence of individual bubbles/droplets. 73

4.16 Input Compartment networks of a closed system to the optimization in Step 3 in Section 4.1.3. **(a)** Original closed system. **(b)** With virtual connections added it is possible for both compartments to have no net outflow. 75

4.17 Labelling the mesh elements of a compartment made of 3 PFRs with 3 DoFs per PFR. The mapping, (c), allows for CM simulation results to be projected onto the mesh, or spatially varying CFD results on the mesh to be projected onto the CM. 76

4.18 Snapshots of the flow profile showing the characteristic periodic steady-state of the system. Colouring based on velocity magnitude with an LIC filter overlayed. 77

4.19 Illustration of how a period signal is sampled and its frequencies used to dictate snapshots to use for compartmentalization. **(a)** Compound sine wave showing data for multiple periods. **(b)** Power spectrum of the data. **(c)** The sampled profile based on all identified frequencies. **(d)** The sampled profile based on only the most important frequency. 79

4.20 Illustration of a hypothetical tracer concentration based on the oscillatory flow described previously, showing how the ensemble creates an envelope that encompasses the oscillations. 80

4.21 Illustration of how a sampled periodic signal is used to created a period set of CMs. **(a)** Periodic signal with snapshots labelled and coloured according to their corresponding CM. **(b)** The resulting compartments and their projection connections to form a closed loop. 81

4.22 Illustration of the projection matrix's role in moving the state between two different CM. **(a)** Mesh elements labelled by DoF k from CM i. **(b)** The same elements labelled by DoF j from CM $i+1$. **(c)** The same elements labelled by source DoF in CM i and destination DoF in CM $i+1$. 82

5.1 Schematic of the simulation domain with boundary conditions with the inlet, outlet, and no-slip boundary conditions labelled. The domain has a height of 1, the inlet length of 2, the outlet length is 10 and the radius of the semi-circle is 2 where all units are non-dimensional. 85

5.2 Line integral convolution (LIC) visualization of the computed velocity field coloured by velocity magnitude. **(a)** The full domain. **(b)** The recirculatory region. 86

5.3 Effect of mesh refinement levels on CFD simulation results. **(a)** Inlet tracer concentration as a function of time. **(b)** CFD RTD. 86

5.4 Visualizations of the compartmentalization of **(a)** the full reference domain, **(b)** the recirculatory region, and **(c)** the directed spatial graph of centroids for each compartment, nodes, and connections between them, edges. Yellow represents mesh elements not compartmentalized, while the remaining colours, red to blue, signify different compartments. Note that the spatial position of the nodes have been modified slightly to decrease overlap and increase readability. 88

5.5 Spatial graphs showing the different compartment networks with nodes representing each CSTR, or PFR, and edges representing the connections between them. **(a)** The PFR-based CM. **(b)** The CSTR-based CM. Note that the spatial position of the nodes have been modified slightly to decrease overlap and increase readability. 89

5.6 RTD functions computed using PFR/CSTR-based CM and directly using CFD for the range of mesh refinements. 90

5.7 KL divergence between RTDs computed using the CSTR/PFR-based CMs and directly through CFD for the range of mesh refinements. 92

5.8 Effect of minimum compartment size on the RTD predicted from the resulting CM. **(a)** The PFR-based CM proposed in this work. **(b)** Previous state-of-the-art CSTR-based CM. 94

6.1 Cross-sectional view, with dimensions labelled, of the stirred tank used for the simulations. From [3]. 98

6.2 Overview of a stirred tank reactor mesh that has been prepared for either an MRF or an SM simulation. **(a)** Overall view of the tank and its mesh. **(b)** Close-up view of the meshing used for the rotating impeller domain. **(c)** View of the bottom of the tank and core of the structured mesh. From [12]. 100

6.3 Cross-section of the tank showing time-averaged fields: **(a)** Gas phase volume fraction. **(b)** Liquid velocity with LIC. **(c)** Turbulent energy dissipation rate. **(d)** Pressure. 101

6.4 Cross-section of the tank demonstrating the the developed compartmentalization method correctly captures the horizontal no-flux planes. **(a)** Velocity profile with LIC. **(b)** Compartmentalization of the entire liquid phase of the reactor, colouring by compartment ID. **(c)** Side by side of velocity and compartments. 102

6.5 Inert tracer concentration along the x-z plane as a function of time. Concentration normalized by the steady-state well-mixed concentration. . . . 103

6.6 Comparison of normalized tracer concentration of the three sensor ports. **(c)** bottom. Values normalized by the steady-state well-mixed concentration. Coloured dots represent experimental data from Vrabél et al. [3]. Coloured dashed lines are CM results from Vrabél et al. [3]. Black line represent average concentration from angle-sweep results of the CM simulation from this work, shaded grey regions represent bounding envelope for the results. Grey dashed line represented steady-state concentration. 105

6.7 Trajectory of CoM as a function of time. Black line is the trajectory of the CM from this work. Grey horizontal line is the target CoM of 0.05. Thick blue line is the range of values from literature [3, 90, 22, 12]. 106

6.8 Dissolved oxygen concentration along the x-z plane as a function of time. 108

6.9 Steady-state CM simulation results. **(a)** Oxygen saturation concentration. **(b)** Dissolved oxygen concentration. **(c)** Dissolved oxygen concentration as a percentage of oxygen saturation concentration. 109

6.10 Plot of dissolved oxygen concentration as a function of height from the bottom of the reactor along the vertical line at $x = -0.3$ and $y = 0$. Red lines represent the location of impellers, blue lines represent the location of the horizontal planes halfway between impellers. The formation of the characteristic Rushton impeller recirculatory regions are visible as as regions with small gradients separated by abrupt changes in concentration. 110

6.11 Glucose concentrations at three different sample ports shows the presence of large-scale oscillations in the measure values. The black squares, "6.48 m", represent the sensor at the top of the reactor, close to the glucose feed port. From [92]. 112

7.1 Labelled 2D vortex shedding domain with boundary names, domain dimensions, and probe location labelled. 115

7.2 Snapshots of the time-varying flow and the time-averaged flow CFD simulations. **Left:** Velocity magnitude with LICs. **Right:** Tracer concentration. . . . 116

7.3 CFD time-varying simulation measurements of oscillations in fields. **(a)** Velocity magnitude over whole simulation. **(b)** Tracer concentration over multiple oscillatory periods. Time in (b) is measured from the injection of the tracer at $t = 10$ s. . . . 117

7.4 Power spectrum on the subset of velocity magnitude measurements. . . . 118

7.5 Sampled velocity magnitude with chosen snapshots labelled. Cyclic colourmap used to emphasize that the first and last compartments along this period are connected. . . . 119

7.6 Comparison of CFD and time-averaged flow field-based CM tracer simulation results at the probe location. . . . 120

7.7 Comparison of CFD and ensemble CM tracer results at the probe location. 121

7.8 The snapshots used for the projection compartmentalization. **(a)** Concentration for full length of coupled tracer simulation. **(b)** Closed loop of CM with Δt labelled. . . . 123

7.9 Comparison of CFD and dynamic projection CM tracer simulation results at the probe location. . . . 124

A.1 Visualization of hydrodynamics from CFD simulation with line integral convolutions indicating local flow direction and color corresponding to velocity magnitude. . . . 149

A.2 Visualization of flow-informed compartmentalization with coloring corresponding to compartment number. . . . 149

A.3 Undirected graph of the compartment network resulting from both (i) flow-information compartmentalization and (ii) the use of spatially-varying compartment approximations (PFRs). . . . 149

A.4 Residence time distribution curves for CFD and CM simulations. . . . 150

List of Tables

5.1 Summary of the mean residence time, variance, and skewness of the RTD from the CFD simulation across the range of tested mesh densities. 87

5.2 Mean residence time from the CFD, PFR-CM, and CSTR-CM simulations at the different tested mesh densities. 91

5.3 Variance of RTD from the CFD, PFR-CM, and CSTR-CM simulations at the different tested mesh densities. 91

5.4 Skew of RTD from the CFD, PFR-CM, and CSTR-CM simulations at the different tested mesh densities. 91

5.5 Computational costs, in cpu-s, of the CFD and CM RTD simulations over the range of tested mesh densities. The CM times are split into Offline, one-time cost for creating the CM, and Online, running the simulation, costs. Neither the CFD nor Offline CM times include the cost of solving for the steady-state velocity profile. 93

5.6 Effect of minimum compartment size on the RTD statistics of the PFR-based and CSTR-based CM. The CFD statistics included are those obtained from the same mesh used for creating both types of CMs. 95

6.1 Physical properties used for the two-phase `OpenFOAM` simulation. 99

6.2 Physical properties used for the oxygen dissolution simulation. 108

7.1 Comparison of the three proposed approaches for CM simulations of systems with time-varying flows. 126

Chapter 1

Introduction

1.1 Research Motivation

Stirred tank reactors are fundamental to chemical engineering processes, particularly in systems involving multiple phases and coupled physical phenomena. These reactors serve as the primary configuration across various industrial applications, with particular significance in pharmaceutical manufacturing [1]. They are characterized by complex air-liquid interfaces with interphase mass transport mechanisms, such as oxygen dissolution, while simultaneously accommodating multiple physical processes including species reactions and heat transfer. The scale of these systems spans several orders of magnitude, from laboratory units of 0.015 L to industrial vessels exceeding 30,000 L [2, 3]. This wide range of volumes presents unique challenges, as increasing reactor dimensions typically lead to more pronounced gradient formation and greater complexity in maintaining optimal operating conditions.

The substantial variation in reactor scales presents significant challenges for laboratory experimentation, particularly in developing robust scaling relationships that reliably predict industrial-scale behaviour. Traditional experimental approaches face several key limitations. Many critical process parameters cannot be measured directly or continuously, and those that can be measured are often restricted by practical constraints of sensor technology and placement. For instance, while dissolved oxygen or pH might be monitored at specific points, comprehensive spatial distribution data is not available. The inherent limitations of physical sensor, both in terms of spatial resolution and the potential for process interference, have driven the increasing adoption of computational methods. Consequently, computer simulations, particularly computational fluid dynamics (CFD) simulations, have

become an indispensable tool in modern engineering practice, offering insights into otherwise inaccessible process dynamics and enabling more comprehensive process optimization strategies.

Despite significant advances in high-performance computing infrastructure and numerical methods, the simulation of multiphase multiphysics processes remains computationally intensive, particularly for industrially relevant scenarios. Stirred-tank bioreactors present unique simulation challenges, primarily due to three fundamental complexities:

1. The size of industrial-scale bioreactors generate extensive systems of nonlinear equations after spatially discretizing the governing equations. This complexity is compounded by the necessity to couple intricate biological reaction phenomena with hydrodynamics at each spatial discretization point.

2. The inherently unsteady nature of these processes [4], combined with convection-dominated flows and moving boundaries, necessitates transient simulations with fine temporal resolution to maintain numerical stability [5].

3. A significant temporal scale disparity exists between microorganism metabolism and multiphase hydrodynamic phenomena. The metabolic reactions occurs over timescales of days but the hydrodynamics impose timesteps of fractions of a second, significantly increasing simulation costs.

Historically, these challenges have rendered comprehensive multiphysics CFD simulations of industrial-scale bioreactors largely infeasible. Continual advances in GPU hardware have decreased the time and monetary cost of these class of simulations [6], though it still remains too slow and too expensive for many use cases. Applications requiring simulation capabilities, such as design optimization and model predictive control (MPC), have traditionally relied on idealized Continuous Stirred Tank Reactor (CSTR) models due to the intractability of full-scale simulations. However, this simplification inherently eliminates the spatial variations crucial to process understanding.

In pharmaceutical manufacturing applications, these three-phase systems (continuous liquid phase with dispersed gas and microorganisms) have predominantly been analyzed under this assumption of perfect mixing. While the assumption holds reasonably well for bench- and laboratory-scale operations, its validity deteriorates significantly with increasing reactor size. Multiple studies have demonstrated the invalidity of this well-mixed approximation for industrial-scale reactors [7, 8, 9, 10].

Even decoupling the hydrodynamics and metabolism through a fixed velocity field, the combination of complex reaction kinetics and large number spatial degrees of freedom,

regularly exceeding 10^6, remains computationally prohibitive for most practical applications [4, 11, 12]. The work of Promma et al. [12] illustrates this challenge. The *E. coli* metabolism model they use, which, despite considering only four pathways and species (glucose, acetate, dissolved oxygen, and biomass), requires approximately 0.65 seconds to compute metabolic fluxes for 10^6 mesh elements. While this computation time might appear modest in isolation, it often becomes significantly larger than that of the hydrodynamics. This is especially true of CFD approaches which rely on a large number of fast-to-compute explicit time-steps, e.g. Lattice-Boltzmann simulations [13].

Given prohibitively high computational costs of full CFD simulations and the inadequacy of ideal CSTR approximations [10], there is a clear need for model order reduction techniques that can maintain acceptable accuracy while substantially reducing computational complexity. Reduced order models (ROMs) bridge this gap, offering a compromise between the oversimplified CSTR and computationally intensive CFD simulation. By sacrificing some of the generality and precision of CFD simulations, ROM methods can reduce computation time by several orders of magnitude. While the field of model order reduction encompasses numerous approaches [14, 15], compartmental models (CMs) emerge as a particularly promising class, leveraging *a priori* knowledge of process hydrodynamics [16]. The knowledge is used to significantly reduce the number of spatial degrees of freedom in order to make feasible the simulation of the complex reactions occurring in the system. This is achieved by taking advantage of the disparate time-scales of the system, the hydrodynamics occur much faster than the mass transport of the different species, and the pseudo-one-way coupling, mass transport of species and reactions can have very little impact on the hydrodynamics, to fully decoupled the hydrodynamics from the reactions and make use of a steady-state or time-averaged hydrodynamics. They are particularly well suited for bioreactor applications because of the complexity and large computational cost per degree of freedom of the metabolic growth equations. Recent comprehensive reviews have demonstrated the broad applicability of CMs across various chemical engineering unit operations, including crystallizers, bubble columns, stirred tank reactors, gas turbines, and wastewater treatment systems [17, 16]. Similar methods have found success in building energy management applications [18], highlighting the versatility of this approach.

The compartmental modelling approach consists of two fundamental steps: 1) compartment identification and 2) modelling of the individual compartments. This approach systematically partitions the spatial domain based on the distribution of user-selected fields, e.g. velocity or concentration, into regions with similar characteristics. As illustrated in Figure 1.1, the two-dimensional domain is decomposed into five distinct compartments. The governing equations within each compartment are then simplified, and the system is approximated as a network of interconnected compartments. While this reduction in

spatial resolution introduces some loss of accuracy, it yields a model with substantially fewer states that can be solved much more efficiently. The computational advantages are significant; for instance, Nadal-Rey [19] reported that their CFD simulation required 30-60 days using 32-48 CPU cores, while an equivalent compartmental model achieved results in just 2 seconds on a single core.

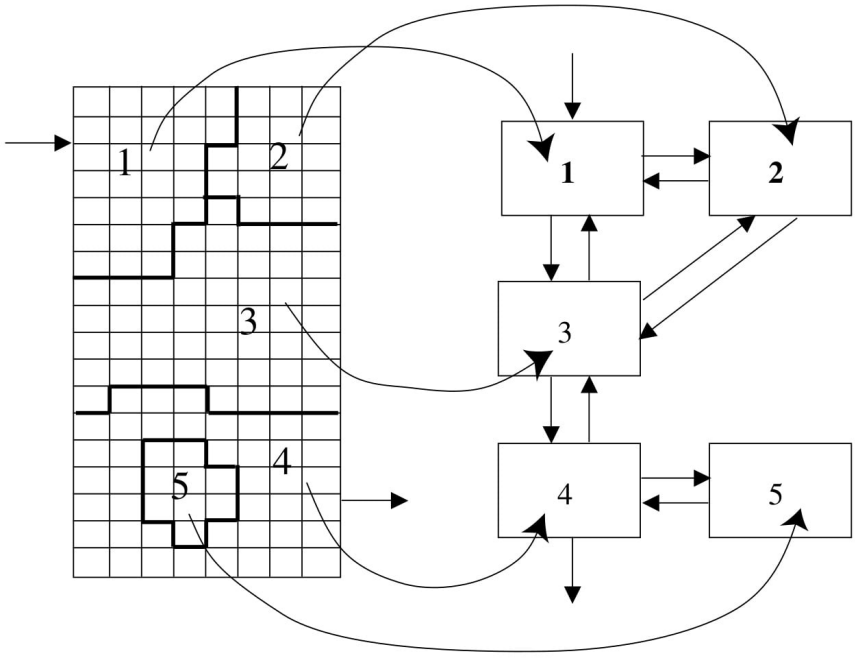


Figure 1.1: Illustration of the fundamental approach of mesh-based compartmental modelling. A CFD mesh is systematically divided into five distinct compartments, with the division occurring at the element level. Each compartment is then represented as a continuously stirred tank reactor (CSTR), and these compartments are interconnected to form a network that approximates the original flow domain. From Ref. [20]

The computational efficiency of CMs enables applications that were previously impractical:

1. Real-time state prediction to support model predictive control implementation for complex spatially-varying reaction systems.

2. Reduced reliance on CFD simulations during process design and troubleshooting. The efficiency of CMs allows for extensive parameter space exploration within the

same computational budget as a single CFD simulation, enabling rapid identification of promising conditions for subsequent detailed analysis.

3. Development of computationally lightweight CFD-derived digital twins. The quasi-one-way coupling would enable equipment manufacturers to develop pre-computed libraries of CFD-based hydrodynamic solutions across operational ranges. These libraries would integrate with CM-based digital twins to provide real-time predictions of spatial mass transport and reaction dynamics.

However, current compartmental modelling approaches in the literature exhibit several significant limitations, including: the use a uniform grid of compartments, the need for manual intervention, manual construction based on expert knowledge of the domain, not prescriptive enough and requires users to pick the fields to compartmentalize by, and inconsistencies between compartmentalization criteria and compartment modelling approaches. These limitations are examined in detail in Chapter 3.

The topology-based approach developed by Donnelly [21] offers a promising foundation for addressing these challenges through its use of fundamental flow characteristics for compartment identification and modelling. This work builds upon Donnelly's idea of a topology-based approach to develop a more comprehensive method that addresses the identified limitations.

1.2 Objectives

The overall objective of this work is to develop a unidirectional flow compartmental modelling method based on analysis of CFD hydrodynamic results. To that end, the specific objectives of this work are:

1. Develop a new method for identifying compartments based on local flow topology.
2. Develop a method for modelling compartments using higher-order models in a way that is congruent with the identification of the compartments.
3. Analyze the performance of the resulting CM by comparing it to reference CFD simulations and to CSTR-based compartmental model.
4. Validate CM predictions on an industrially-relevant 3D two-phase system.

5. Demonstrate the benefit of the proposed CM by simulating spatially-varying inter-phase mass transfer in a two-phase system.

6. Extend the method to handle time-varying velocity fields.

7. Compare computation times of the proposed method in different scenarios against past compartmental modelling methods and against CFD.

8. Implement and release the methods in an open-source package.

1.3 Structure of Thesis

This thesis is comprised of seven other chapters.

Chapter 2 establishes the theoretical foundation necessary for this work. The chapter presents an overview of stirred tank bioreactors and challenges associated with scaling them up, it examines single- and multi-phase CFD simulation methods, it introduces reduced-order modelling, and it provides an overview of compartmental modelling.

Chapter 3 presents a comprehensive review of state-of-the-art compartmental modelling methods. The review focuses on approaches for compartmentalization, strategies for compartment modelling, methods for modelling dynamic compartments, topology-based methods for compartmental modelling, and treatment of zero-convective-flux planes.

Chapter 4 develops the compartmental modelling method. It begins with a method for single-phase systems and then extends it to function for complex three-dimensional multi-phase systems. The chapter concludes with approaches for analyzing time-varying flows and presents two methods for developing CMs from them.

The subsequent three chapters validate the methods proposed in Chapter 4. Chapter 5 validates the method using a representative two-dimensional geometry, comparing the flow-informed approach against existing CSTR-based models and CFD simulations. Chapter 6 extends the validation to an industrially relevant three-dimensional two-phase stirred tank bioreactor, demonstrating successful handling of zero-flux planes, accurate prediction of mixing times and concentration profiles, and effective simulation of spatially varying inter-phase mass transfer. Chapter 7 addresses the limitations of compartmental modelling for inherently transient flows and evaluates two methods for modelling such systems.

Lastly, Chapter 8 summarizes the conclusions of this work and makes recommendations for future research.

Chapter 2

Background

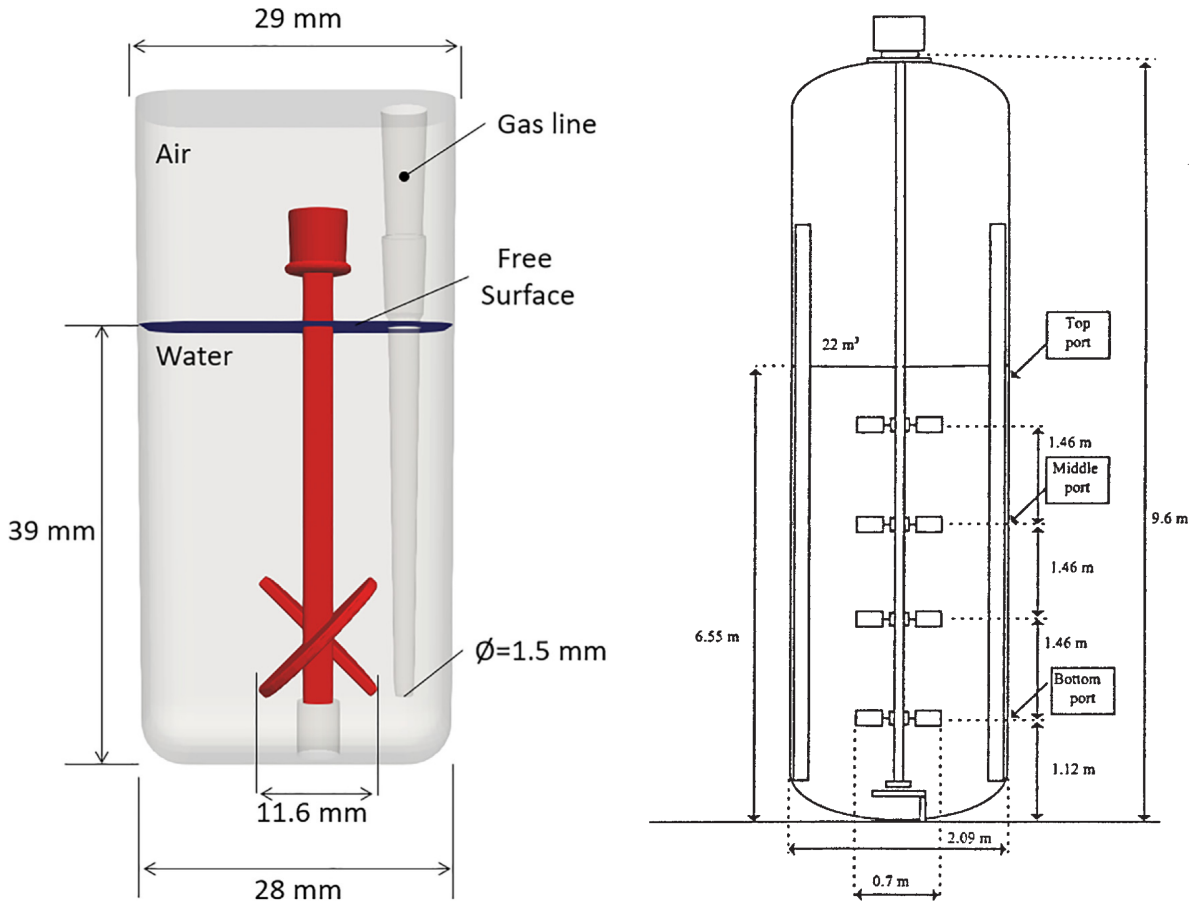
Large-scale stirred tank bioreactors have significant operational challenges that impact process efficiency and product quality. This chapter examines these challenges through the lens of computational fluid dynamics (CFD), with particular emphasis on the multiphase systems characteristic of industrial bioreactors. The discussion encompasses both the fundamental physical phenomena that must be modelled and some of the computational approaches available for their simulation. Next it provides an overview of reduced order modelling approaches, specifically compartmental models, and their need in balancing computational tractability with physical accuracy. The chapter concludes with a description of the different numerical methods used to generate the results of the thesis, focusing on the finite difference discretization scheme used to solve the system of partial differential equations derived from the compartmental model, including the treatment of the resulting differential algebraic equations.

2.1 Ideal Versus Non-ideal Stirred Tank Behaviour

Stirred tank reactors represent the predominant reactor configuration in various industries, particularly in pharmaceutical manufacturing. Their widespread adoption stems from both historical precedent and practical advantages, including straightforward design principles and lower operational requirements.

While some reactors are operated under anaerobic conditions, many applications require aerobic environments maintained through bottom-mounted air spargers. These air bubbles, combined with the presence of a deformable liquid-air free surface between the top of the

liquid and the headspace, create an inherently multi-phase system. Further complicating these systems is the wide range of scales over which they are build; their size spans several orders of magnitude in volume, ranging from bench-scale units of 15 mL to industrial-scale vessels exceeding 30 m^3. Figure 2.1 illustrates the geometry of bioreactors at either size extreme.



(a) (b)

Figure 2.1: Representative geometries of the two extremes in bioreactor sizes showing that they feature the same basic elements: narrow aspect ratio, immersed impeller, and a sparger. Also notable is the inclusion of baffles on the larger reactor to aid with mixing, a feature not needed on the smaller reactor. **(a)** is a 15 mL Ambr microbioreactor from [2]. **(b)** is a 30 m^3 industrial bioreactor from [3].

The significant range in operating volumes presents fundamental challenges for labo-

ratory experimentation, particularly in developing representative scaled models that can reliably predict or diagnose performance issues in industrial-scale reactors. For example, temperature is typically regulated using a heating/cooling jacket mounted on the vessel surface. As reactor size increases, the surface area for heat transfer grows slower than the heat-generating volume. This scaling relationship increases system response time and often necessitates modifications to operating procedures to ensure that control can be maintained.

Mass transfer presents similar scale-dependent challenges. As reactor dimensions increase, the formation of gradients become more pronounced due to the interaction between convective transport and reaction kinetics. In aerobic reactors, oxygen transfer occurs through two mechanisms: sparged air bubbles and the free surface interface. While mass transfer through the free surface can account for up to 30% of total oxygen transfer in microreactors and benchtop systems [2], this contribution becomes negligible at larger scales due to the decreasing ratio of free surface area to total volume. Consequently, dissolved oxygen concentrations become dependent on the sparger system. This shift creates oxygen concentration gradients, as oxygen addition occurs at a single point, due to practical design constraints, while consumption occurs throughout the reactor volume.

Similar gradients develop in the concentrations of dissolved species such as glucose. Maintaining homogeneity at larger scales requires proportionally higher convective transport rates for a given reaction rate. However, increased convection demands greater energy input and generates elevated shear rates near impeller regions, potentially causing cell damage and death. Thus it cannot be increased enough to prevent gradients. Additionally, substrate consumption rates can vary by three orders of magnitude during a single cultivation [22, 23]. The common practice of positioning feed ports above the liquid level, while beneficial for contamination control and cleaning procedures, further contributes to mixing challenges.

Global gradients are not the only concern; local concentration extremes present additional challenges. The introduction of pH adjustment solution can create zones of lethal pH levels, while concentrated nutrient feeds may induce metabolic overflow, resulting in unwanted byproduct formation. Moreover, as organisms circulate through these zones of varying concentrations, they may experience selective pressures leading to unintended mutations, particularly during extended cultivation periods [24, 10].

The prediction and mitigation of these and other undesired behaviour requires more than experimental investigation alone, as both economic factors and practical limitations constrain their feasible scope. Computational approaches provide essential predictive capabilities and enable the evaluation of alternative operational strategies without the need

for physical process modifications.

Despite widespread recognition of gradient formation, many historical modelling approaches, particularly in process control applications, have relied on idealized assumptions of perfect mixing. While more sophisticated models exist, their implementation presents significant practical challenges. CFD simulations effectively characterize hydrodynamic behaviour and could, theoretically, be coupled with cellular metabolism models to provide comprehensive operational predictions. However, integrating even basic metabolic models with CFD simulations becomes computationally intractable due to the inherent complexity of cellular metabolism and the coupled nature of hydrodynamic and mass transfer phenomena. This computational limitation necessitates alternative approaches for coupling CFD hydrodynamic solutions with metabolic growth models, such as compartmental modelling, which will be examined in subsequent sections.

The alternative approaches explored and developed in this work will all be based on CFD results. To perform these CFD simulations, appropriate mathematical models are required. Since only the hydrodynamics of the system are solved, the discussion about mathematical models will be limited to those required for simulating the conservation of mass and momentum for the liquid and gas phases. The discussion will start with single-phase systems and build up to multi-phase ones.

2.2 Single-Phase Computational Fluid Dynamics

To simulate the movement of liquid in the single-phase simulations performed in Chapter 5 and Chapter 7 the Navier-Stokes equations will be used. These equations simulate the movement of the liquid by modelling the conservation of mass and momentum of the phase. The conservation of mass and momentum for an incompressible Newtonian fluid, in the absence of reaction, is expressed as:

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0 \tag{2.1}$$

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{\nabla} \cdot (\boldsymbol{u}\boldsymbol{u}) = -\boldsymbol{\nabla} P + \nu \Delta \boldsymbol{u} + \boldsymbol{g} \tag{2.2}$$

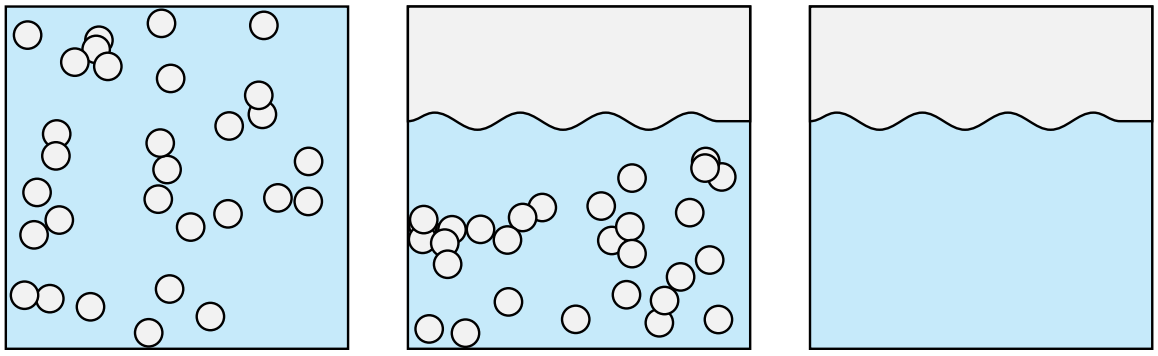
where $\boldsymbol{u}$ is the velocity, P the modified pressure, ν the kinematic viscosity, and $\boldsymbol{g}$ the force of gravity.

This formulation is applicable to incompressible Newtonian single-phase systems under all flow conditions. However, the computational requirements increase significantly with turbulent flows, making direct simulation impractical. Various approaches to manage these

computational demands through turbulence modelling are examined in Section 2.4. Before examining turbulence models, it is appropriate to consider the extension of these governing equations to multi-phase systems, allowing for a unified discussion of turbulence modelling across both single- and multi-phase scenarios.

2.3 Multi-phase Computational Fluid Dynamics

The end-goal of this work requires CFD simulations of multi-phase systems. Such systems can be categorized into three main groups, as illustrated in Figure 2.2, each requiring a specialized CFD model. This section will focus on an overview of models for dispersed and separated flows.



(a) Dispersed (b) Transitional (c) Separated

Figure 2.2: Illustration showing three different flow regimes of interest for multi-phase simulations. Grey is used for the gas phase while blue is used for the liquid phase. **(a)** Liquid phase continuous and gas phase dispersed. **(b)** Liquid phase continuous and gas phase dispersed in one region of the domain and continuous in another. **(c)** Both phases continuous, each confided to different regions of the domain and separated by a large contiguous interface.

Dispersed flows are characterized by a continuous phase containing distributed domains of secondary phases, such as small gas bubbles dispersed in a liquid. These systems can, in principle, be modelled using: an interface-capturing model, which tracks the movement of each phase and the interface between them; a Two-Fluid Model (TFM), which treats phases as interpenetrating continua; or an Euler-Lagrange model, which tracks discrete particles or bubbles within the continuous phase. In practice, tracking the interface of

each individual bubble becomes prohibitively expensive and only the TFM and the Euler-Lagrange model are used.

Separated flows, featuring large-scale phase boundaries, require interface-capturing techniques to accurately model the evolution of the interfaces. A typical example is the gas-liquid interface at the surface of a stirred tank vessel. The more complex transitional flows exhibit both dispersed elements as well as large-scale phase separation, as seen in aerated stirred tank vessels where bubbles rise and burst through the free surface, and required a hybridized approach.

2.3.1 Interface-Capturing Models

Interface-capturing methods explicitly resolve phase boundaries within the computational domain. While theoretically capable of tracking all interfaces, including those of individual bubbles or droplets, the computational requirements would be prohibitive for practical applications. As a result, these methods are typically used for modelling large-scale interfaces, such as gas-liquid free surfaces or large discrete phases, while smaller-scale phenomena are handled through hybridization with complementary methods.

Several interface-capturing frameworks exist, including Volume of Fluid [5], phase-field [5], and level-set methods [25]. While these methods were not directly implemented in the current research, they warrant discussion as they impose significant computational costs on simulating two-phase stirred tank bioreactor systems. Each method offers distinct advantages and limitations in capturing complex interface dynamics, particularly relevant to the free-surface behaviour and gas-liquid interactions characteristic of stirred tank bioreactors.

The first, is the Volume of Fluid (VoF) method. It uses a scalar indicator function, α, to designate phase presence within each mesh element:

$$\alpha = 0 \qquad \text{phase 1} \tag{2.3}$$

$$\alpha = 1 \qquad \text{phase 2} \tag{2.4}$$

$$0 < \alpha < 1 \qquad \text{interface} \tag{2.5}$$

This approach tracks phase distribution throughout the domain, with interface reconstruction occurring in elements where α takes intermediate values, as illustrated in Figure 2.3. While interface tracking is implicit, the reconstructed interface is sharp. Phase evolution, and consequently interface movement, is governed by the advection equation:

$$\frac{\partial \alpha}{\partial t} + \boldsymbol{\nabla} \cdot (\alpha \boldsymbol{u}) = 0 \tag{2.6}$$

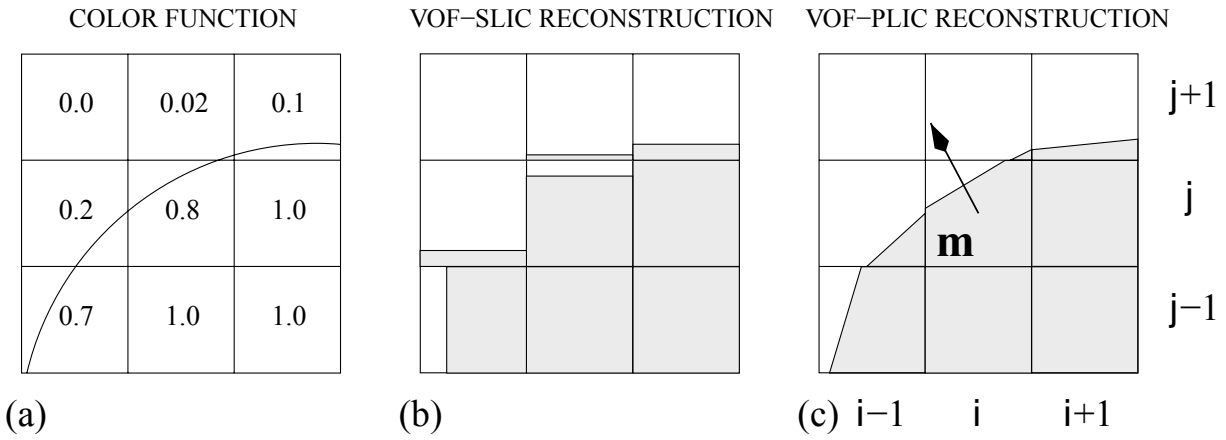


Figure 2.3: Overview of VOF methods. **(a)** The exact interface overlayed by the numerical grid and the values of the indicato rfunction. **(b)** Simple Line Interface Calculation (SLIC) reconstruction of the interface. **(c)** Piecewise Linear Interface Calculation (PLIC) reconstruction of the interface. $\boldsymbol{m} = -\nabla C$ is the outward facing unit normal. Showing how the surface is reconstructed from the field. From [5].

The primary advantages of this method lie in its conservation properties and ability to maintain sharp interfaces.

The phase-field method shares similarities with the VoF method, utilizing a scalar field indicator, ϕ, to designate phase regions:

$$\phi = -1 \qquad \text{phase 1} \tag{2.7}$$

$$\phi = 1 \qquad \text{phase 2} \tag{2.8}$$

$$-1 < \phi < 1 \qquad \text{interface region} \tag{2.9}$$

$$\tag{2.10}$$

The interface is represented as a diffuse region of finite thickness, characterized by $\phi = 0$. Interface dynamics are governed by the Cahn-Hilliard equation:

$$\frac{\partial \phi}{\partial t} + \boldsymbol{u} \cdot \boldsymbol{\nabla} \phi = M \frac{\partial F(\phi)}{\partial \phi} \tag{2.11}$$

where M represents the mobility coefficient and F denotes the free energy functional. A key advantage of this approach is its explicit treatment of surface tension.

Finally, the level set method employs a signed distance function, ϕ, to characterize the interface location and distinguish between phases. This function represents the shortest

distance from any point to the interface, with its sign indicating the phase:

$$\phi < 0 \qquad \text{phase 1} \tag{2.12}$$
$$\phi > 0 \qquad \text{phase 2} \tag{2.13}$$
$$\phi = 0 \qquad \text{interface} \tag{2.14}$$

A primary advantage of this approach is its precise interface definition, represented by the zero-level set ($\phi = 0$). This mathematical formulation enables accurate calculation of interface geometric properties, such as normal vectors and curvature, making it particularly suitable for problems where interface topology undergoes significant changes. However, a fundamental limitation of the method is its non-conservative nature, which can lead to mass conservation errors during extended simulations. This characteristic often necessitates additional computational procedures to maintain mass conservation.

2.3.2 Two-Fluid Model

To address the high computational cost associated with explicitly tracking each bubble/particle using an interface-capturing method, the TFM was derived from interface-capturing models using volume and time averaging. Instead of modelling the two phases separated by a distinct boundary, it models them as interpenetrating continua. Multiple different derivations of the TFM exist, including the most popular one by Ishii [26] and a promising alternative by Brennen [27]. Although Brennen's formulation demonstrates stronger physical justification and preliminary studies indicate superior performance [28], the results presented in Chapter 6 are produced using `OpenFOAM` which uses Ishii's derivation. The governing equations of the two-fluid model, comprised of conservation of mass, conservation of momentum, and the phase fraction constraint, are presented in Equation (2.15).

$$\begin{aligned}
\frac{\partial \alpha_i \rho_i}{\partial t} + \boldsymbol{\nabla} \cdot (\alpha_i \rho_i \boldsymbol{u}_i) &= \mathcal{I}_i \\
\frac{\partial \alpha_i \rho_i \boldsymbol{u}_i}{\partial t} + \boldsymbol{\nabla} \cdot (\alpha_i \rho_i \boldsymbol{u}_i \boldsymbol{u}_i) &= -\boldsymbol{\nabla} \cdot (\alpha_i p_i) + \boldsymbol{\nabla} \cdot (\alpha_i \boldsymbol{\tau}_i) + \boldsymbol{\nabla} \cdot (\alpha_i \boldsymbol{\tau}_i^t) \\
&\quad + \alpha_i \rho_i \boldsymbol{g} + \boldsymbol{\mathcal{M}}_i \\
\sum_i \alpha_i &= 1
\end{aligned} \tag{2.15}$$

where α_i is the volumetric phase fraction of the i^{th} phase, $\boldsymbol{u}_i$ the phasic velocity, ρ_i the phasic density, $\mathcal{I}_i$ the total interphase mass transport term, p_i is the phasic pressure, $\boldsymbol{\tau}_i$

is the stress tensor, $\boldsymbol{\tau}_i^t$ is the Reynolds stress tensor, $\boldsymbol{g}$ is the acceleration due to gravity, and $\boldsymbol{\mathcal{M}}_i$ is the total interphase mass transfer term. The various time/volume averaging indicators have been omitted for readability. The turbulence term $\boldsymbol{\nabla} \cdot (\alpha_i \boldsymbol{\tau}_i^t)$ is omitted from the simulation when operated in the laminar regime.

Because the model does not explicitly capture individual droplets/bubbles of the dispersed phase, $\boldsymbol{\mathcal{M}}_i$ is of paramount importance; it is responsible for the entirety of the interactions, such as drag, which are otherwise captured explicitly by more computationally expensive models. Figure 2.4 illustrates several of the most common terms. It has been shown that the choice of terms to include, and even the closure used for a given term, has a significant impact on the stability and well-posedness of the resulting model [28].

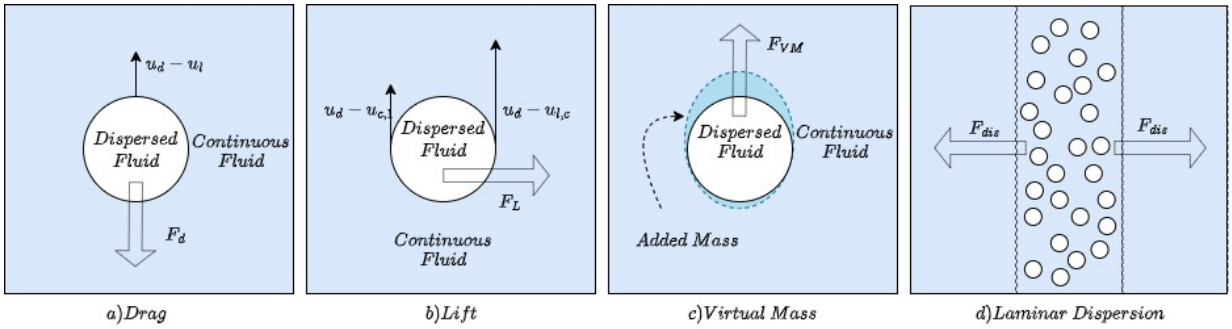


Figure 2.4: Illustrations representing the most popular interphase momentum transfer terms for which closures are used to approximate $\boldsymbol{\mathcal{M}}_i$. Created by Arshia Fazeli, used with permission.

The advantages of the TFM include substantially reduced computational requirements compared to Euler-Lagrange models. Further, while not achieving the accuracy or stability of interface-capturing models, Euler-Euler models can be used to approximate free-surface behaviour, a capability notably absent in Euler-Lagrange models.

However, it also has several significant drawbacks. The first is that it relies heavily on empirical closures, which can limit its generalizability. The method also does not explicitly track individual particles, making it difficult to capture particle history effects, coalescence phenomena, breakup, and particle-particle interactions in general. Additionally, fundamental physical constraints, such as maximum packing fractions and the requirement to maintain phase fractions between 0 and 1, are not inherently satisfied by the model formulation. Ensuring these constraints are satisfied adds additional computational costs and require specialized solvers [29].

2.3.3 The Euler-Lagrange Model

The Euler-Lagrange approach represent an alternative framework to modelling dispersed two-phase flows. In these models, the continuous phase is governed by the Navier-Stokes equations, while the dispersed phase is described through Newton's Second Law of Motion and, where applicable, mass conservation equations accounting for inter-phase transport phenomena. Multiple variants exist, each tailored to specific operational conditions and scenarios.

The selection of an appropriate model variant is largely determined by the Stokes number, the ratio of particle relaxation time to the convective timescale. At very low Stokes numbers, particles effectively act as flow tracers, following fluid streamlines. Conversely, at high Stokes numbers, $>> 1$, particle trajectories become increasingly independent of fluid motion due to the increasing dominance of their inertia.

For scenarios involving very low Stokes numbers, the simplest modelling approach treats particles as perfect flow followers; this significantly reduces the computational cost by eliminating the need to solve Newton's Second Law. While this simplification has limited general applicability, it can be particularly suitable for modelling microorganisms, which often exhibit neutral buoyancy and high surface area-to-volume ratios, resulting in Stokes numbers much less than one [30].

For particles with higher Stokes numbers, Newton's Second Law must be explicitly solved. That said, if the dispersed phase occupies a sufficiently small volume fraction, particle-particle interactions can be reasonably neglected. However, in both of the previous approaches, the interphase momentum transfer terms, the forces acting on each particle, are accounted for using closures, similar to Euler-Euler models.

As particle-particle interactions increase in importance, more costly simulations are required where particle are modelled using the Discrete Element Method. Figure 2.5 shows an example of a complex Euler-Lagrange simulation of a slurry flow through sand. Both the slurry and sand particles are explicitly modelled, the flow between them explicitly captured, and particle-particle and particle-wall collisions modelled.

DEM simulations can be categorized into two distinct approaches based on their resolution requirements and accuracy objectives. The first approach, "unresolved simulations", uses closure models for inter-phase momentum transfer, similar to those used in Euler-Euler frameworks. These simulations can be implemented with either one-way or two-way coupling between phases. While this method does not resolve detailed flow features, such as particle boundary layers, it permits the use of meshes elements significantly larger than the particles, substantially reducing computational costs. This approach can feasibly handle

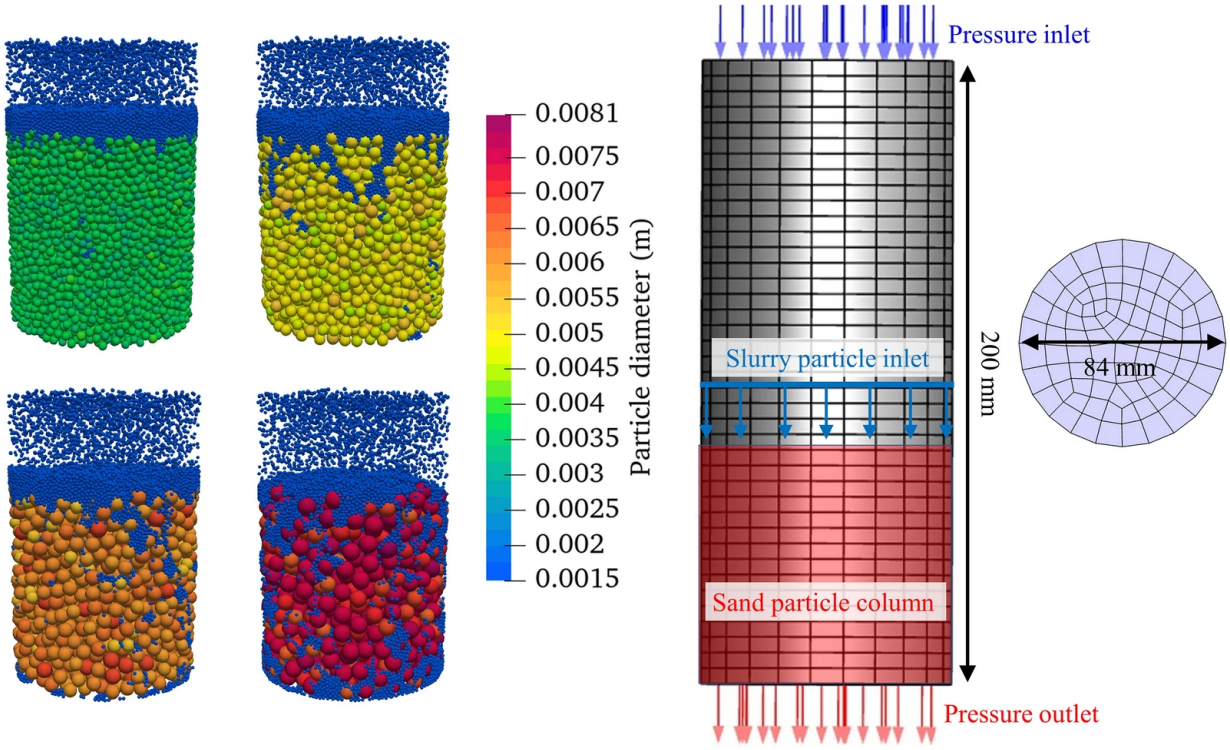


Figure 2.5: Illustration of an Euler-Lagrange simulation using DEM particles to simulate the slurry-slurry, slurry-sand, and sand-sand interactions. While computationally intensive, such approaches are required when flow between particles, as well as their interactions, are crucial to achieving the goal of the simulation. From [31].

tens of millions of particles using high end desktop hardware, i.e. multiple GPUs, making it particularly suitable for systems where particle dimensions are much smaller than the domain characteristic length [32].

The second, high-resolution "resolved simulations", approach resolves the boundary layer around particles and explicitly calculates inter-phase momentum transfer through surface integration over individual particles. While this method provides higher accuracy, it comes at a substantial computational cost. The need for mesh elements significantly smaller than particle dimensions reduces both the feasible domain size and particle count by several orders of magnitude compared to "unresolved simulations".

Advantages of Euler-Lagrange approach include straightforward tracking of individual particle histories and heterogeneities; biological processes such as growth, cell division, and dissolution can be explicitly modelled at the particle level. When necessary, particle-

particle interactions can be directly incorporated into the simulation framework. However, significant drawbacks exist. The computational cost is substantially higher than Euler-Euler models, particularly when modelling large particle populations. Time-step size is severely constrained by the need to properly resolve particle-particle and particle-wall interactions to prevent particles passing through each other or through walls. Finally, the method cannot independently represent free surface behaviour, necessitating hybridization with interface-capturing models.

2.4 Turbulence Modelling

Stirred tank systems typically operate in the turbulent regime, adding complexity to both single- and multi-phase simulations. While the previously discussed models are theoretically applicable to turbulent flows, their direct application is computationally intensive. Direct Numerical Simulation (DNS) of single-phase systems is restricted primarily to academic applications because of this. The challenge stems from the need to resolve turbulent eddies across multiple spatial scales, requiring extremely fine mesh resolution and small time steps.

Practical approaches to reduce computational demands are based on the insight that smaller eddies carry proportionally less energy while requiring a disproportionately more computational resources. These methods replace explicit eddy modelling with closure terms that approximate their collective effects on the flow field. This discussion will first address single-phase systems, as they form the foundation for more complex multi-phase turbulence models.

2.4.1 Single-Phase Simulations

As shown in Figure 2.6, there are two classes of turbulence models: Large-Eddy Simulation (LES) and Reynolds Averaged Navier-Stokes (RANS)-based models. LES is developed by filtering out eddies below a certain size; the small eddies carry very little energy but dominate the computational cost of the simulation. While more accurate than RANS, it will not be discussed further since it is too computationally expensive to run given the software and hardware available for this research.

RANS-based models are developed from the NS equations by decomposing the flow into a time-averaged mean, $\langle \boldsymbol{u} \rangle$ and a fluctuating component, $\langle \boldsymbol{u} \rangle'$ [34]:

$$\boldsymbol{u} = \langle \boldsymbol{u} \rangle + \boldsymbol{u}' \tag{2.16}$$

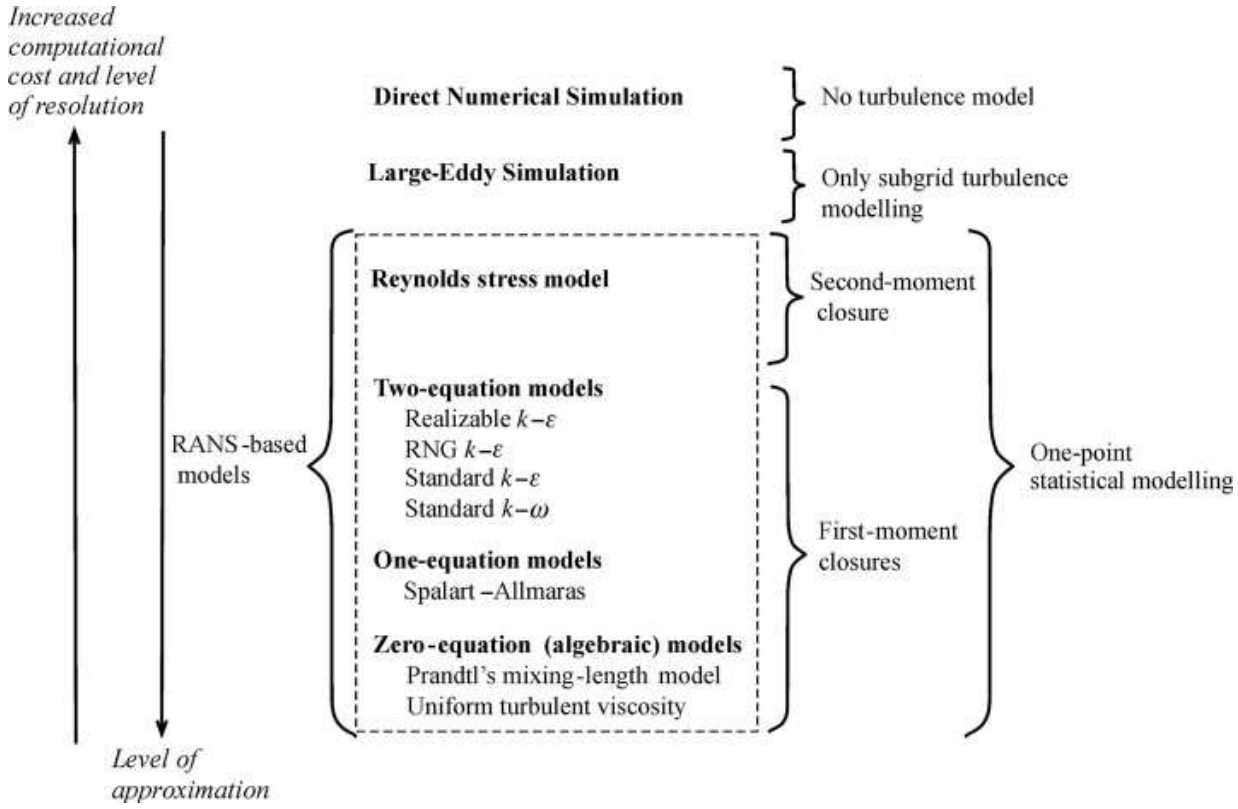


Figure 2.6: Overview of the different approaches to single-phase turbulence modelling, showing the spectrum of approaches available ordered by computational intensity and accuracy. From [33].

The resulting equation is then time-averaged and most, but not all, of the fluctuating components are average out to zero. The resulting time-averaged conservation equations, for an incompressible Newtonian liquid, are:

$$\boldsymbol{\nabla} \cdot \langle \boldsymbol{u} \rangle = 0 \tag{2.17}$$

$$\frac{\partial \langle \boldsymbol{u} \rangle}{\partial t} + \boldsymbol{\nabla} \cdot (\langle \boldsymbol{u} \rangle \langle \boldsymbol{u} \rangle) = -\boldsymbol{\nabla} \langle P \rangle + \nu \Delta \langle \boldsymbol{u} \rangle + \boldsymbol{g} - \boldsymbol{\nabla} \cdot \boldsymbol{\tau} \tag{2.18}$$

where $\boldsymbol{\tau} = -\rho \langle \boldsymbol{u}' \boldsymbol{u}' \rangle$ is the Reynolds-stress term, the negative of the Reynolds-stress tensor divided by density which accounts for the effects of turbulence on the mean stress. Since $\boldsymbol{u}'$ is not modelled explicitly, $\boldsymbol{\tau}$ requires a closure, the choice of which has profound impacts on the performance and accuracy of the simulation.

There are two types of models that are generally used: Boussinesq approximation-based

models (two parameters) and so-called Reynolds stress models (six parameters). Only the Boussinesq-based k-ε model will be covered as it is the most common model for free and wall turbulence and is the one used in this work. It models $\boldsymbol{\tau}$ as proportional to the gradient of the mean velocity [33]:

$$\frac{\boldsymbol{\tau}}{\rho} = -\left\langle \boldsymbol{u}'\boldsymbol{u}' \right\rangle = \nu_T \left(\boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle + \boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle^T \right) - \frac{2}{3} k \boldsymbol{\mathcal{I}} \tag{2.19}$$

where k is the turbulent kinetic energy:

$$k = \frac{1}{2} \left\langle \boldsymbol{u}'\boldsymbol{u}' \right\rangle \tag{2.20}$$

and ν_T is the turbulent viscosity which also requires a closure. For the k-ϵ model that closure is:

$$\nu_T = C_\mu \frac{k^2}{\varepsilon} \tag{2.21}$$

where C_μ is a model parameter, and ε represents the rate at which turbulent kinetic energy is dissipated as thermal energy:

$$\varepsilon = \nu \left\langle \boldsymbol{\nabla}\boldsymbol{u}' \cdot \boldsymbol{\nabla}\boldsymbol{u}'^T \right\rangle \tag{2.22}$$

To use the model, two transport equations are required, one for k and one for ε [33]:

$$\begin{aligned} \underbrace{\frac{\partial k}{\partial t}}_{\text{Accumulation}} + \underbrace{\left\langle \boldsymbol{u} \right\rangle \cdot \boldsymbol{\nabla} k}_{\text{Convection}} = & \underbrace{\nu_T \left[\left(\boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle + \boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle^T \right) : \boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle \right]}_{\text{Generation}} \\ & + \underbrace{\boldsymbol{\nabla} \cdot \left[\left(\nu + \frac{\nu_T}{\sigma_k} \right) \boldsymbol{\nabla} k \right]}_{\text{Diffusion}} \\ & - \underbrace{\varepsilon}_{\text{Dissipation as thermal energy}} \end{aligned} \tag{2.23}$$

$$\begin{aligned} \underbrace{\frac{\partial \varepsilon}{\partial t}}_{\text{Accumulation}} + \underbrace{\left\langle \boldsymbol{u} \right\rangle \cdot \boldsymbol{\nabla} \varepsilon}_{\text{Convection}} = & \underbrace{C_1 \frac{\varepsilon}{k} \nu_T \left[\left(\boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle + \boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle^T \right) : \boldsymbol{\nabla} \left\langle \boldsymbol{u} \right\rangle \right]}_{\text{Generation}} \\ & + \underbrace{\boldsymbol{\nabla} \cdot \left[\left(\nu + \frac{\nu_T}{\sigma_\varepsilon} \right) \boldsymbol{\nabla} \varepsilon \right]}_{\text{Diffusion}} \\ & - \underbrace{C_2 \frac{\varepsilon^2}{k}}_{\text{Dissipation}} \end{aligned} \tag{2.24}$$

where C_1, C_2, σ_ε, and σ_k are model parameters.

The drawback of this model is that, on its own, it incorrectly predicts $\boldsymbol{\tau}$ near the walls because the Boussinesq approximation assumes that $\boldsymbol{\tau}$ is isotropic, which is not true in the vicinity of immersed objects.

2.4.2 Multi-Phase Simulations

The Two-Fluid Model (TFM) equations in Equation (2.15) incorporate both volume averaging and Reynolds time averaging. The turbulent stress term, $\boldsymbol{\tau}_i^t$, requires closure modelling similar to single-phase systems. This work employs the mixture k-ϵ model implemented in `OpenFOAM`, following the formulation by Behzadi et al. [35], with effective gas density and bubble-induced turbulence modifications from Lahey et al. [36]. The closure for $\boldsymbol{\tau}_i^t$ is based on the Boissinesq hypothesis and is expressed as:

$$\frac{\boldsymbol{\tau}_i^t}{\rho_i} = -\nu_i^t \left(\boldsymbol{\nabla} \langle \boldsymbol{u}_i \rangle + \boldsymbol{\nabla} \langle \boldsymbol{u}_i \rangle^T - \frac{2}{3} \boldsymbol{\mathcal{I}} \boldsymbol{\nabla} \cdot \langle \boldsymbol{u}_i \rangle \right) + \frac{2}{3} k_i \boldsymbol{\mathcal{I}} \tag{2.25}$$

The turbulence diffusion coefficient follows a similar closure to the single-phase model, where $\nu_i^t = C_\mu \frac{k_i^2}{\varepsilon_i}$, necessitating conservation equations for both k_i and ε_i.

Initial multiphase turbulence models, such as that developed by Kataokat et al. [37], adapted single-phase methods by attributing turbulence exclusively to the continuous phase. However, this approach remains valid only for systems with minimal dispersed phase concentrations. A key distinction in multiphase systems is the presence of interphase transport terms for kinetic energy, e.g. bubble-induced turbulence. Behzadi et al. [35] advanced the framework established by Gosman et al. [38], where the latter introduced dispersed phase turbulence through a response function C_t, relating the Reynolds stress terms of the two phases:

$$\mathrm{R}_d = C_t^2 \mathrm{R}_c \tag{2.26}$$

Which, after applying the Boussinesq approximation, relates the turbulent fluctuations of the two phases as :

$$\boldsymbol{u}_d' = C_t \boldsymbol{u}_c' \tag{2.27}$$

Behzadi's extension implements a more sophisticated response function, tracking mixture's turbulent kinetic energy k_m and dissipation rate ε_m. This enhancement enables accurate modelling at higher dispersed phase concentrations, including phase inversion scenarios, while maintaining consistency with single-phase k-ε behaviour as $\alpha_d \to 0$.

The `OpenFOAM` implementation of Behzadi et al.'s model incorporates two significant modifications from Lahey [36]. The first introduces an effective gas phase density, Equation (2.40), to account for entrained liquid mass in bubble motion. The second implements an interphase transport term (Equation (2.30)) to represent bubble-induced turbulence effects. With these modifications, the transport equations for mixture turbulent kinetic energy (k_m) and dissipation rate (ε_m) are expressed as:

$$\underbrace{\frac{\partial \rho_m k_m}{\partial t}}_{\text{Accumulation}} + \underbrace{\boldsymbol{\nabla} \cdot \left(\rho_m \left\langle \boldsymbol{u}_m \right\rangle k_m\right)}_{\text{Convection}} = \underbrace{G_k^m}_{\textit{Generation}} + \underbrace{\boldsymbol{\nabla} \cdot \left(\frac{\mu_m^t}{\sigma_k} \boldsymbol{\nabla} k_m\right)}_{\text{Diffusion}} - \underbrace{\rho_m \varepsilon_m}_{\text{Dissipation}} + \underbrace{S_k^m}_{\text{Interphase Transport}} \tag{2.28}$$

and

$$\underbrace{\frac{\partial \rho_m \varepsilon_m}{\partial t}}_{\text{Accumulation}} + \underbrace{\boldsymbol{\nabla} \cdot \left(\rho_m \left\langle \boldsymbol{u}_m \right\rangle \varepsilon_m\right)}_{\text{Convection}} = \underbrace{\left(C_1 \frac{\varepsilon_m}{k_m}\right) G_k^m}_{\text{Generation}} + \underbrace{\boldsymbol{\nabla} \cdot \left(\frac{\mu_m^t}{\sigma_\varepsilon} \boldsymbol{\nabla} \varepsilon_m\right)}_{\text{Diffusion}} - \underbrace{\left(C_2 \frac{\varepsilon_m}{k_m}\right) \rho_m \varepsilon_m}_{\text{Dissipation}} + \underbrace{\left(C_3 \frac{\varepsilon_m}{k_m}\right) S_k^m}_{\text{Interphase transport}} \tag{2.29}$$

where C_1, C_2, C_3, σ_k, σ_ε are model parameters.

The interphase transport term for turbulent kinetic energy, modelling bubble-induced turbulence, is [36]:

$$S_k^m = \max(0, \alpha_d - \alpha_p) \frac{C_p \alpha_c \alpha_d \rho_c}{d} \left(|\boldsymbol{u}_r|^3 + \left(\frac{C_d \nu_m}{d}\right)^{4/3} |\boldsymbol{u}_r|^{5/3} \right) \tag{2.30}$$

where α_p is a model parameter for specifying the dispersed phase volume fraction below which bubble-induced turbulence is neglected, and d the bubble diameter.

The above equations are obtained from the individual phase equations through averag-

ing using following definitions:

$$k_m = \left(\alpha_c \frac{\rho_c}{\rho_m} + \alpha_d \frac{\rho_{d,\text{eff}}}{\rho_m} C_t^2\right) k_c \tag{2.31}$$

$$\varepsilon_m = \left(\alpha_c \frac{\rho_c}{\rho_m} + \alpha_d \frac{\rho_{d,\text{eff}}}{\rho_m} C_t^2\right) \varepsilon_c \tag{2.32}$$

$$\langle \boldsymbol{u}_m \rangle = \frac{\alpha_c \rho_c \langle \boldsymbol{u}_c \rangle + \alpha_d \rho_{d,\text{eff}} \langle \boldsymbol{u}_d \rangle C_t^2}{\alpha_c \rho_c + \alpha_d \rho_{d,\text{eff}} C_t^2} \tag{2.33}$$

$$\mu_m^t = \frac{\left(\alpha_c \mu_c^t + \alpha_d \mu_c^t C_t^2\right) \rho_m}{\alpha_c \rho_c + \alpha_d \rho_{d,\text{eff}} C_t^2} \tag{2.34}$$

$$\mu_d^t = c_t^2 \left(\frac{\nu_c}{\nu_d} \frac{\rho_{d,\text{eff}}}{\rho_c}\right) \mu_c^t \tag{2.35}$$

$$G_k^m = \alpha_c G_k^c + \alpha_d G_k^d \tag{2.36}$$

$$k_d = C_t^2 k_c \tag{2.37}$$

$$\varepsilon_d = C_t^2 \varepsilon_c \tag{2.38}$$

$$\rho_m = \alpha_c \rho_c + \alpha_d \rho_{d,\text{eff}} \tag{2.39}$$

$$\rho_{d,\text{eff}} = \rho_d + C_{\text{vm}} \rho_c \tag{2.40}$$

where $\rho_{d,\text{eff}}$ is the effective density of the dispersed phase, ρ_d the phasic density of the dispersed phase, C_{vm} is the virtual mass coefficient, and ρ_c is the density of the continuous phase.

Finally, the equation for C_t used is:

$$C_t(\alpha_d) = 1 + (C_{t0} - 1) e^{-f(\alpha_d)} \tag{2.41}$$

where $f(\alpha_d)$ is given by:

$$f(\alpha_d) = 180\alpha_d - 4.71 \times 10^3 \alpha_d^2 + 4.26 \times 10^4 \alpha_d^3 \tag{2.42}$$

and C_{t0} is the value of C_t at $\alpha_d \approx 0$, obtained from the following set of equations:

$$C_t = \frac{3+\beta}{1+\beta+2\rho_d/\rho_c} \tag{2.43}$$

$$\beta = \frac{2A_d l_e^2}{\rho_c \nu_c \mathrm{Re}_t} \tag{2.44}$$

$$l_e = C_\mu \frac{k_c^{1.5}}{\varepsilon_c} \tag{2.45}$$

$$A_d = \frac{3\alpha_d \rho_c C_d |\boldsymbol{u}_r|}{4d} \tag{2.46}$$

$$\mathrm{Re}_t = \frac{|\boldsymbol{u}''| l_e}{\nu_c} \tag{2.47}$$

$$|\boldsymbol{u}''| = \sqrt{2k_c/3} \tag{2.48}$$

where A_d is the drag coefficient on a bubble, l_e the eddy length scale, Re_t the turbulence Reynolds number, C_μ is a model parameter

Equation (2.28) and Equation (2.29) are solved and Equation (2.31) and Equation (2.32) are used to related the mixture values to those of the continuous phase, which are them related to those of the dispersed phase using Equation (2.37) and Equation (2.38).

2.5 Impeller Motion

Solving the governing equations for single- and multi-phase flows requires spatial discretization of the domain into a connected network of polytopes. Using a conformal meshing approach, geometric features including walls, baffles, and impellers must be explicitly represented by mesh facets aligned with their surfaces. This explicit geometric representation presents a particular challenge when modelling systems with moving components, such as stirred tank reactors, where the mesh must accommodate dynamic boundaries. The most physical, and naive, approach would be to directly deform or reconstruct the mesh at each time step. The associate cost of the mesh manipulation and the subsequent projection operation would make even small simulations computational infeasible. Consequently, several alternative methods have been developed to model impeller motion more efficiently.

The two discussed are the Multiple Reference Frame (MRF) method and the Sliding Mesh (SM) method. Both methods still use conformal meshes, but instead of a single mesh for the whole domain, the domain is split into two meshes: one for the rotating domain, e.g.

around the impeller shaft and the impellers in a tank; and one for the stationary domain, e.g. the rest of the tank. This is shown from different angles in Figure 2.7; specifically notice the cylindrical zones around the impellers in Figure 2.7(a).

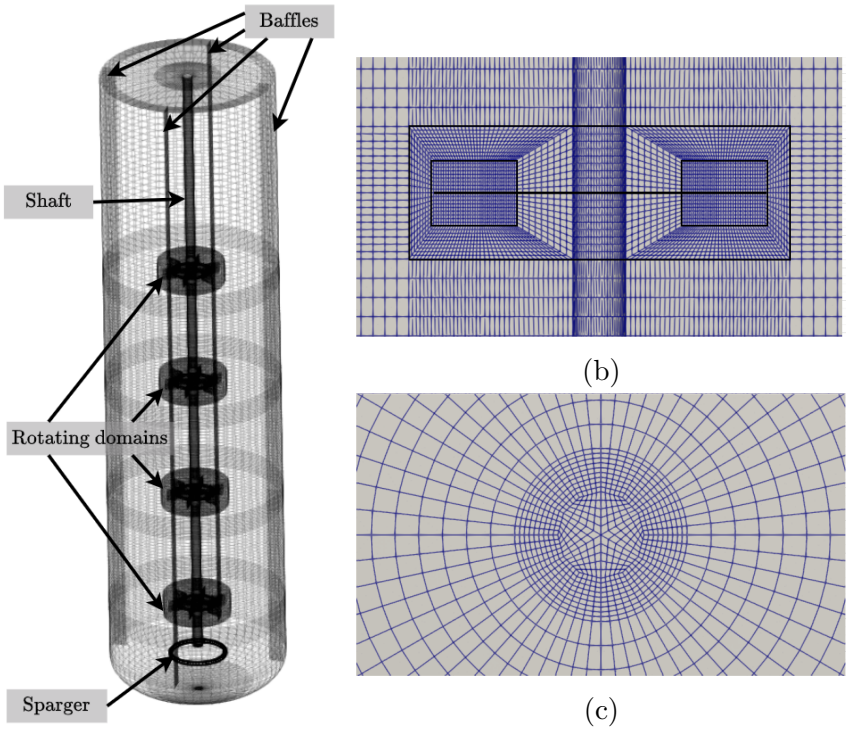


Figure 2.7: Overview of a stirred tank reactor mesh that has been prepared for either an MRF or an SM simulation. **(a)** Overall view of the tank and its mesh. **(b)** Close-up view of the mesh used for the rotating impeller domain. **(c)** View of the bottom of the tank and core of the structured mesh. From [12].

While both methods can use the same mesh, their approach to performing the simulations is fundamentally different. The MRF method, as the name implies, uses multiple frames of references to solve the governing equations without needing to move any of the mesh. Both the stationary domain and rotating domain are solved from a reference frame in which the domain appears stationary. This means that the rotating domain is solved in a non-inertial frame of reference, and additional terms need to be added to the governing equations [39]. The boundary between the two domains also requires boundary conditions for transporting mass, momentum, etc. between the two frames. Because it uses a fixed

mesh, the simulation is significantly faster but this comes at the cost of not capturing the true time-varying nature of the flow.

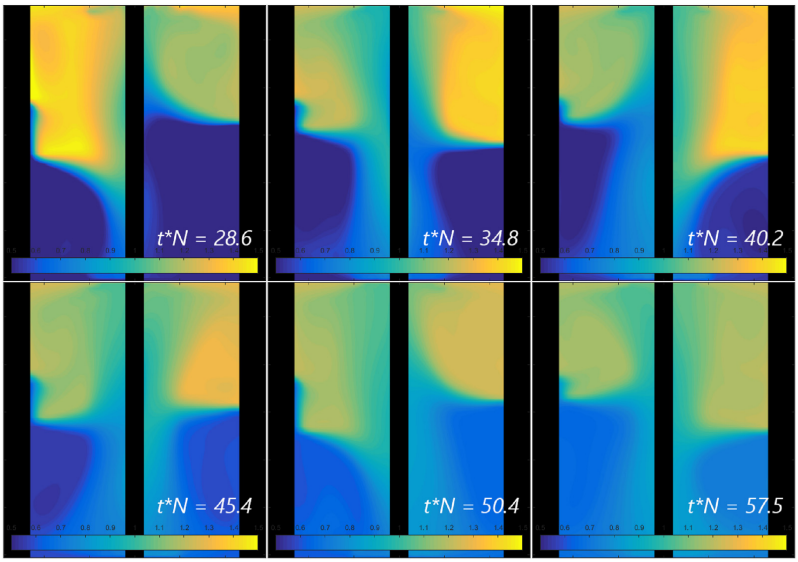


Figure 2.8: Snapshots of a transient SM-based simulation of an inert tracer injected into a multi-Rushton impeller tank. The figure focuses on the inter-impeller region, with a Rushton impeller being located at the top and bottom of the frame. The macroinstability resolved by the SM method is seen in the middle of each frame as it oscillates and precesses. With an MRF-based simulation, that plane would be horizontal and fixed. From [40].

The other, more accurate, method is the SM method. Unlike the MRF method, the SM method rotates the domains relative to one another, thus providing a much more accurate simulation. As the inner domain moves relative to the outer domain, special care must be taken with the mesh density at their interface so that quantities can be properly conserved and moved from one domain to the other. While mesh reconstruction is still avoided through interface interpolation, the SM method has significantly higher computational costs and may become numerical unstable [39]. The benefit of the method is that it captures the real time evolution of the flow as well as inherently non-stationary phenomena. For example, Haringa et al. show that using the MRF method instead of the SM method causes the simulation to impose artificial steady-states onto the system [40]. For the stirred tank reactor they simulated, the MRF method removes the oscillations of the horizontal plane between impellers. That plane, as shown in Figure 2.8, does not remain horizontal, but in-fact precesses around the impeller axis. This motion is captured by the SM method, but is entirely absent in an MRF-based simulation. Consequently,

the SM method is particularly valuable for applications where non-stationary effects are significant.

2.6 Reduced Order Modelling

The substantial computational cost of high-fidelity CFD simulations limit their applicability in many practical scenarios, particularly in real-time applications such as Model Predictive Control. Traditionally, process control applications have relied on idealized reactor models, such as Continuously Stirred Tank Reactors (CSTR) and Plug Flow Reactors (PFR), to circumvent these computational constraints. Reduced Order Models (ROMs) offer a middle-ground approach, maintaining essential flow physics while significantly reducing computational cost through dimensional reduction. Unlike heavily simplified idealized models, ROMs preserve the dominant flow features of the original system while decreasing the degrees of freedom, thereby achieving a balance between computational efficiency and physical accuracy.

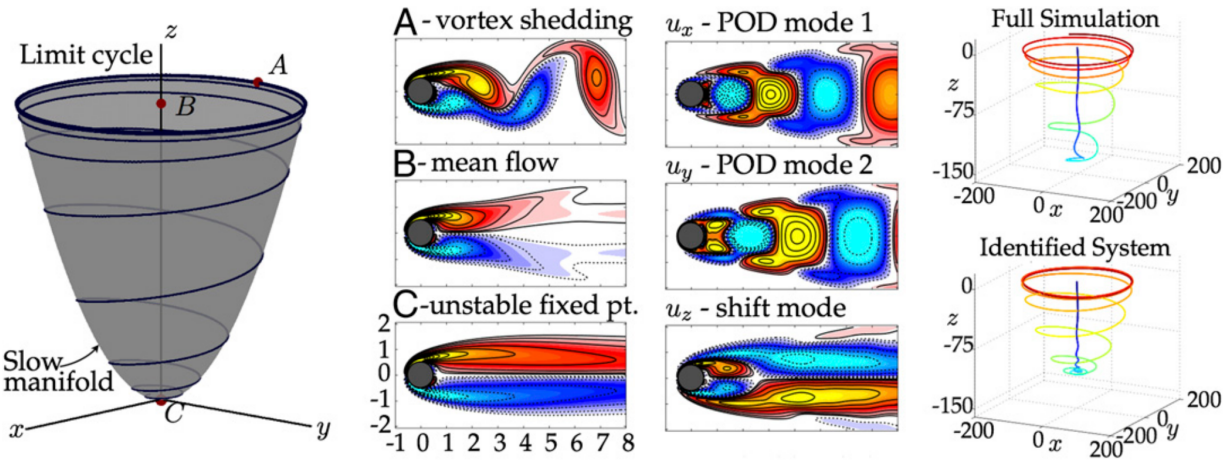


Figure 2.9: Example of a POD-based ROM for the hydrodynamics of a flow-past a cylinder scenario. **Left:** Labelled solution manifold showing the time-evolution of the coefficient to the 3 POD modes, with certain key locations labelled. **Middle-left:** Three different flow snapshots labelled by their position on the manifold. **Middle-right:** The three most important modes used to simulate the system. **Right:** The trajectory of the system in the state-space of the three modes for both the original CFD simulation and the ROM simulation based on the three most important modes. From [41].

There have been several different methods for creating ROMs, each designed for a different application. Proper Orthogonal Decomposition (POD) [42, 41] is a prevalent technique that identifies the most significant modes of a system, thereby reducing the dimensionality of the problem while retaining the most dominant features. Galerkin projection can then be used to project the original governing equations onto the reduced basis and model the dynamics of the system out in time. The work of Brunton et al. [41], illustrated in Figure 2.9, demonstrates a POD-based decomposition of cylinder wake flow. Their approach reduces the velocity field to three principal modes, which are then used to predict temporal flow evolution. Additionally, machine learning-based methods are increasingly being adopted to develop ROMs, leveraging data-driven strategies to capture complex nonlinear behaviours that traditional techniques may not adequately represent. For example Lee and Carlberg [43] use an autoencoder architecture in order to effectively produce a non-linear POD of the system.

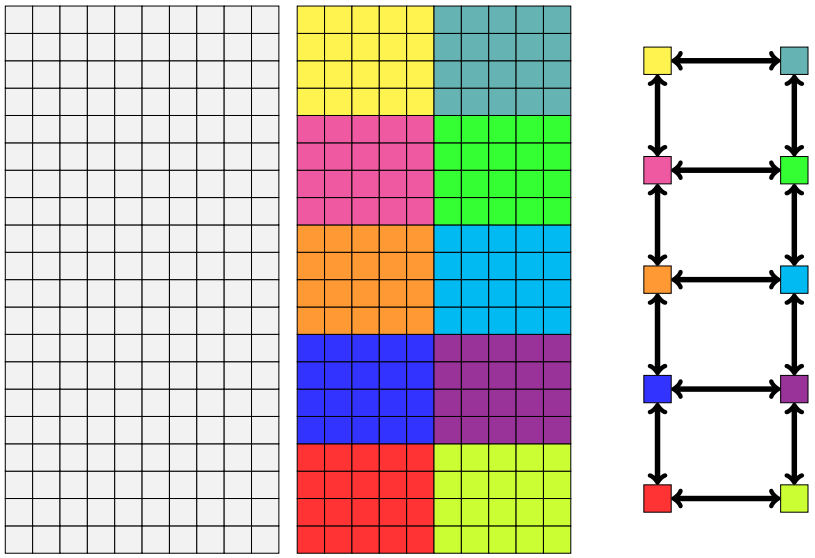
While these ROM approaches effectively reduce degrees of freedom within their defined basis set, this reduction is limited to the space of that reduced basis. In the cylinder wake example (Figure 2.9), although the system is described by three degrees of freedom corresponding to modes u_x, u_y, u_z, the reconstruction of the complete velocity field maintains the original system's degrees of freedom. Therefore, incorporating reaction kinetics presents challenges: either multiple simulations must be performed to generate a reaction basis set, or reactions must be solved in the full-dimensional space. As described earlier, the reaction system frequently involve hundreds of coupled non-linear reactions representing metabolic processes whose computational cost typically exceeds that of the hydrodynamics. It is for these very reasons that a ROM is needed and the simulation of the reaction on the original domain is unwanted, or infeasible. Compartmental models represent a class of ROMs particularly well suited for this applications.

2.7 Compartmental Models

While ROM techniques such as POD and Galerkin projection focus on reducing the dimensionality of detailed models, compartmental models (CMs) approach complexity by dividing the system into distinct regions, compartments. This segmentation facilitates the analysis of system dynamics by concentrating on the interactions and transfers between these compartments. To maintain clarity, it is important to explicitly lay out terminology: a compartment model refers to the reduced-order representation of an individual compartment, while a compartment**al** model refers to the complete network of interconnected compartment models. The rest of this section will provide an overview of how compart-

mental models are created, while a review of recent advances in compartmental modelling is covered in Chapter 3.

CM development follows a two-step process illustrated in Figure 2.10. The input to the compartmental modelling process is a combination of expert-knowledge and simulation/experimental results. Figure 2.10(a) shows the structured CFD mesh of a hypothetical domain. Several past works, especially older ones, lacked CFD results and relied on expert-knowledge and design equations for modelling the process, e.g. [3]. In these cases, both the shape of the compartments and the flowrate through their connections were manually selected based on expert-knowledge.



(a) (b) (c)

Figure 2.10: Creation of the compartmental model from the CFD mesh. **(a)** The initial CFD mesh from which the compartments are to be made. **(b)** The result of the first step, each mesh element is coloured based on the compartment that it belongs to. **(c)** The result of the final step, each compartment is modelled as a single CSTR and a network is created based on the compartment connectivity.

Compartmentalization, the first step in developing a CM, involves partitioning the domain into non-overlapping compartments. The result of this step is illustrated in Figure 2.10(b), where distinct colours denote different compartments. Historically, manual compartmentalization has dominated the literature. This approach has been employed

under various circumstances: when CFD data is unavailable [3], when substantial expert knowledge of the flow exists [44, 45, 9], when iterative model refinement is desired [46], when manual methods outperformed automated approaches [47], when CFD software limitations exist [48, 49], or when visual inspection proves more practical than programmatic implementation [50]. Automated compartmentalization approaches have emerged to accelerate the process, formalize expert knowledge, and reduce the dependence on user expertise [24, 21]. When leveraging CFD results, several methods have been explored: manual partitioning [51, 44, 52], uniform grid discretization [53, 54, 55], and more sophisticated approaches based on spatial property variations [24]. The different methods and their trade-offs is discussed further in Section 3.1.

Compartment modelling, the second step in CM development, involves applying ROMs to individual compartments, creating a network of coupled regions governed by systems of differential equations. Then, flowrates are calculated between neighbouring compartments and the results are used to created a directed graph, a network, of the compartments. Figure 2.10(c) shows a hypothetical compartment network from the compartments previously identified in (b).

The idealized CSTR serves as the fundamental building block for most compartmental modelling approaches. For each CSTR, the conservation of mass equation is:

$$\frac{dC_i}{dt} = \sum_{j}^{inlets} C_j \frac{Q_{j\to i}}{V_i} - C_i \sum_{k}^{outlets} \frac{Q_{i\to k}}{V_i} + R_i \tag{2.49}$$

where i is CSTR index, j upstream CSTR (or domain inlet) of i, k downstream CSTR to i, C well-mixed concentration, V reactor volume, $Q_{i\to j}$ volumetric flowrate from CSTR i to CSTR j, and R reaction rate per unit volume.

Initial research in this field employed multiple CSTRs within individual compartments, with their arrangement determined by either CFD results or established understanding of flow patterns. Notable examples include the work of Vrábel et al. [3] and Alvarado et al. [51], who implemented networks of interconnected CSTRs within each compartment. The configuration of these networks, including both the spatial distribution and number of CSTRs, was guided by expert knowledge, as illustrated in Figure 2.11. Subsequent CSTR-based approaches, exemplified by Rigopoulos et al. [20], adopted a simpler framework where each compartment is represented by a single CSTR. However, this choice of comes with several drawbacks, which are discussed in Section 3.2.

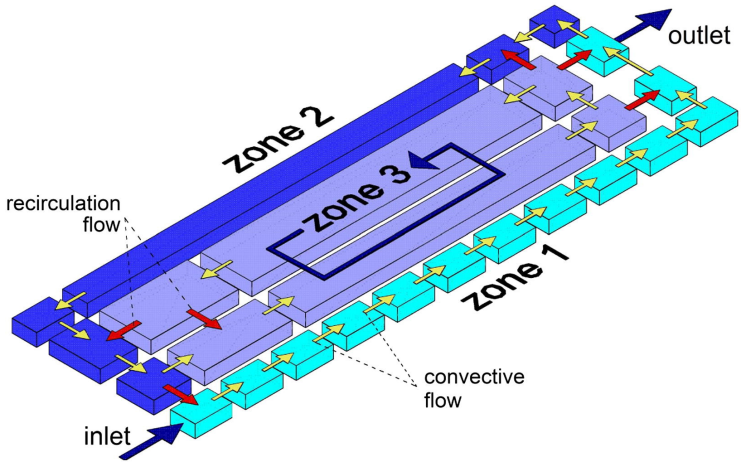


Figure 2.11: A wastewater treatment pond manually split into 3 compartments, referred to as zones in the figure, based on expert analysis of the flow field.. Each compartment is then modelled as a network of CSTRs, with connections between CSTRs, both within and across compartments, also based on expert knowledge of the flow. From [51].

2.8 Numerical Methods

This section reviews the numerical methods relevant to the development, simulation, and validation of the CM. It starts with a discussion of how mixing performance is quantified using CFD results. Then, a brief overview of the Finite Element and Finite Volume Methods, which form the basis of the reference CFD simulations, is provided. Next, the Finite Difference Method is then examined, focusing on the specific discretization scheme used to transform the PDEs governing the systems of PFRs into a system of coupled ODEs. The section concludes with an overview of Differential Algebraic Equations and their application to handling PFR boundary conditions, including a discussion of the chosen solution strategy.

2.8.1 Coefficient of Mixing

Because of impact that non-uniformity has on reactor performance, measures for quantifying it are needed. One such approach, which is based on a single measurement, is the analysis of the Residence Time Distribution (RTD) of the system. The time taken for a system to reach 95 or 99% of the steady-state value provides insight into how quickly and

thoroughly mixing occurs. Further insights can also be drawn from the shape of the RTD curve.

Various mixing metrics can be derived when multiple measurement points or CFD simulation data are available across a domain. The Coefficient of Mixing (CoM), introduced by Hartmann et al. [56], quantifies mixing performance through the volume-weighted standard deviation of concentration, as expressed in Equation (2.50). While CoM yields a single scalar value, it incorporates spatial concentration data from the entire domain, providing a more comprehensive assessment than single-point measurements used in RTD or concentration profiles. Additionally, CoM is applicable to closed systems, overcoming a key limitation of RTD analysis.

$$\mathrm{CoM} = \sqrt{\frac{\sum_i \left(\frac{c_i - \bar{c}}{\bar{c}}\right)^2 \Delta V_i}{V}} \tag{2.50}$$

where c_i is the concentration of degree of freedom i, $\bar{c}$ is the average concentration of the domain, ΔV_i is the volume of degree of freedom i, and V is the total volume of the domain.

As CoM approaches 0, the reactor approaches fully mixed. From this, a target value can be set and the time taken to reach that target is declared the mixing time of the vessel. The benefit of this measure is that it takes into account the spatially non-uniformity throughout the reactor in a way that a single measurement cannot.

2.8.2 Finite Element Method

The Finite Element Method (FEM) is a technique for solving boundary value problems through the use of interpolating polynomials. It is an alternative to techniques such as the Finite Difference and Finite Volume Method (FVM). Some of its benefits are [57]:

- Higher order methods are easier to derive since the interpolants are defined on a per-element basis.
- Irregular/curved domains as well as non-uniform meshes can be easily handled.
- Boundary conditions, especially derivative/flux boundary conditions, naturally arise in the course of deriving the weak form of the FEM and are solved with the same schemes and with the same order of accuracy as the rest of the domain.

The standard version of the FEM, the continuous Galerkin (CG) formulation, requires that the interpolating polynomials be continuous across element facets. Consequently, the CG formulation has numerical stability issues when modelling convection dominated flows or steep gradients/discontinuities. The discontinuous Galerkin (DG) formulation has been proposed to address these issues.

The DG formulation can be considered as a generalization of the FVM and FEM method, which combines the benefits of both [58]. As the name implies, the DG formulation allows for the interpolating polynomials to be discontinuous across element facets. Those discontinuities are then handled through the introduction of numerical flux terms, similar to those appearing in the FVM. As such, it is able to handle convective flows and solutions with very steep gradients better than both the FEM, since it does not result in spurious oscillations, and the FVM, since it uses higher order interpolants [59]. See Figure 2.12 below for a comparison of CG and DG in capturing a curve exhibiting steep gradients for the given mesh size.

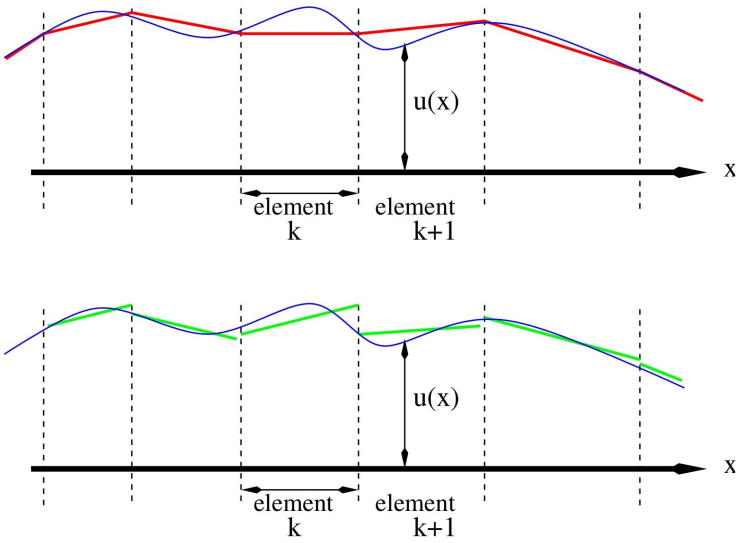


Figure 2.12: Comparison of a CG (top) and DG (bottom) approach capturing the same known solution, each using linear interpolants. The discontinuities in the DG version allow it to better approximate the nonlinear solution, but they will require special treatment for evaluating the value at the element bound. From [60].

Furthermore, because of the discontinuous nature of the formulation, the evaluation of the residual occurs on a per-element-basis. Thus, if suitable numerical flux choices are used, the DG method will also be locally conservative. The drawbacks of the DG method

are: a) the weak form is significantly more complicated to derive and implement than for CG, and b) the memory and computation requirements of DG, for the same number of elements, are generally much higher than for CG or FVM given the same number of elements.

However, the relative higher computational expense of DG versus other method depends on the way that is expense is quantified. For example, Monte compared the FVM performance in `OpenFOAM` to the CG and DG FEM implementations of incompressible Navier-Stokes in `OpenCMP` for a single-phase 2D lid-driven cavity flow [61]. It was found that for the same number of elements, DG takes $30 - 50\times$ longer (210 s) to solve than the FVM (3.94 s) or CG (5.59 s). However, when the simulation is run to obtain a mesh-independent solution, DG is faster (210s) than both the FVM (256 s) and CG (576 s) while also using significantly fewer elements (25,600 for FVM, 14,400 for CG, and 1,600 for DG) *and* resulting in a final error that was approximately 23% smaller as compared to either of the other two methods.

2.8.3 Finite Volume Method

Since two of the main CFD simulations in this thesis were performed with `OpenFOAM`, a brief overview of the numerical solution method employed by this open-source software is provided here. `OpenFOAM` is based on the Finite Volume Method (FVM), which is a widely used technique for solving boundary value problems by dividing the computational domain into discrete control volumes. Within each control volume, the integral form of the conservation equations is solved, ensuring that mass, momentum, and energy are conserved at a local level, and thus in the domain as a whole.

This control volume-based discretization allows it to be inherently conservative, a property that the CG FEM does not have. Further, unlike CG FEM, whose interpolants are continuous across mesh cells, the FVM is better suited for capturing discontinuities and convection-dominated flows. Finally, the conceptual framework used by FVM is more intuitive to understand and straightforward to implement as it directly operates on recognizable control volumes without necessitating the more complex mathematical constructs required by the FEM, such as deriving weak forms or selecting appropriate function spaces.

The advantages of FVM, coupled with the historical availability of more commercial FVM-based CFD software, have contributed to its continued widespread use for fluid dynamics simulations. However, the FVM is not without its drawbacks and limitations. It is heavily influenced by the quality and degree of structure of the computational mesh, to a greater extent than alternative numerical methods. This mesh sensitivity also manifests in

the challenges associated with accurately approximating higher-order derivatives, as gradients must be evaluated by considering multiple mesh cells; a process that becomes more difficult to do accurately on unstructured meshes.

Additionally, while the FVM may provide *a* solution faster and with a lower computational cost than the FEM, achieving a *mesh-independent* solution can often prove more time and resource-intensive. Finally, on a practical level, the difficulty in working with and extending the `OpenFOAM` package itself, while not inherent to the FVM, represents an obstacle specific to the most prominent open-source implementation of this numerical approach.

2.8.4 Finite Difference Method

The proposed compartmental modelling method uses PFRs to create the model for each compartment. The conservation of mass for each PFR is represented by a transient PDE:

$$\frac{\partial C_i(V,t)}{\partial t} = D_{\text{eff}} A_i \frac{\partial^2 C_i(V,t)}{\partial V^2} - Q_i \frac{\partial C_i(V,t)}{\partial V} + R_i(V,t) \tag{2.51}$$

where i is the PFR index, C concentration, V reactor volume, D_{eff} effective Taylor-Arris diffusivity, A reactor cross-sectional area, Q_i volumetric flowrate, and R_i reaction rate per unit volume.

For the initial implementation, mass diffusion will not be considered and the equation simplifies to:

$$\frac{\partial C_i(V,t)}{\partial t} = -Q_i \frac{\partial C_i(V,t)}{\partial V} + R_i(V,t) \tag{2.52}$$

Following the conditions described above, each PDE also required one boundary/inlet condition as follows:

$$C_i(0,t) = \frac{\sum_j^{inlets} Q_j C_j(V_j,t)}{\sum_j^{inlets} Q_j} \tag{2.53}$$

Numerical solution of this system requires spatial discretization, implemented here through the Finite Difference Method (FDM). This approach discretizes the domain into a finite set of points and approximates spatial derivatives using a Taylor series expansion about these points. While such discretization typically transforms coupled PDEs into systems of initial value ODEs, the boundary conditions in this case result in a system of Differential Algebraic Equations (DAEs). The conversion of these DAEs to ODEs is addressed in the following subsection.

The selection of derivative approximation schemes presents various trade-offs that must be carefully considered. The convective nature of the system necessitates particular attention to minimizing numerical diffusion and accurately resolving wavefronts. Although higher-order schemes generally offer superior accuracy, they present specific challenges in this application context. These include the requirement for specialized boundary treatments, which may introduce inconsistencies in discretization order near domain boundaries, and the potential generation of spurious oscillations requiring additional stabilization mechanisms.

A significant practical constraint specific to this application arises from the limited DoFs available per PFR. In contrast to conventional applications where fine spatial discretization is feasible, the CM represents each compartment through multiple PFRs, each necessarily restricted to relatively few discretization points. This constraint diminishes the potential advantages of higher-order schemes and may render them impractical.

A first-order upwind scheme was selected for spatial discretization based on these constraints. Upwind schemes incorporate flow directionality in derivative approximations, making them well-suited for convection-dominated problems. While first-order schemes exhibit lower accuracy compared to higher-order alternatives at equivalent ΔV and Δt, they offer practical advantages for this application: reduced computational overhead for oscillation suppression and better compatibility with the limited degrees of freedom characteristic of the PFR compartments. The discretization of the convection term using a first-order upwind scheme is:

$$Q_i \left.\frac{\partial C_i(V,t)}{\partial V}\right|_{V=V_j} \approx Q_i \frac{C_{i,j}(t) - C_{i,j-1}(t)}{\Delta V} \tag{2.54}$$

where the volume is discretized into K points, $j \in [1, K]$, with uniform spacing ΔV.

This formulation is an upwind scheme since the flow is in the direction of increasing j, thus the gradient at each position is evaluated using only information upwind of it.

2.8.5 Differential Algebraic Equations

While the resulting system of equations for the FDM-discretized PFRs looks like a system of ODEs, the boundary condition used for each PFR actually makes it a system of Dynamic Algebraic Equations (DAEs). A "regular" system of ODEs is a system of explicit ODEs which can be written in the form:

$$y' = f(y,t) \tag{2.55}$$

where y is the dependent variable, t the independent, and $f(\cdot)$ a function of both y and t which provides a unique value for y'.

DAEs on the other hand are part of the class of semi-explicit ODEs; such systems have the general form:

$$y' = f(t, y, z) \tag{2.56}$$
$$0 = g(t, y, z) \tag{2.57}$$

where y are the dependent variables with explicit, isolatable, and unique derivatives and z the dependent variables without such derivatives.

The classic example of a DAE is the unit mass and unit length pendulum swinging under gravity attached to frictionless fulcrum by a rigid rod, as illustrated by Figure 2.13.

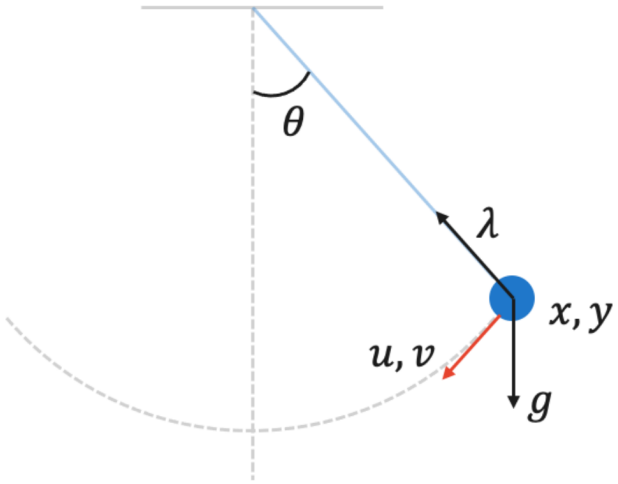


Figure 2.13: Diagram of the pendulum system in Cartesian coordinates. x and y are position, u and v are velocity components in x and y, g is acceleration from gravity, θ is the angle between the rod and the vertical, and λ is the tension per unit length of the rod. From [62].

In Cartesian coordinates the governing equations for this system are:

$$x' = u \tag{2.58}$$
$$y' = v \tag{2.59}$$
$$u' = -\lambda x \tag{2.60}$$
$$v' = -\lambda y - g \tag{2.61}$$
$$0 = x^2 + y^2 - 1 \tag{2.62}$$

where (x, y) denotes the position of the pendulum tip, (u, v) represents its velocity components, and λ is the tension force per unit length along the rod.

This system illustrates two key characteristics of DAEs: the absence of an explicit equation for λ or λ', and the presence of a purely algebraic constraint $x^2 + y^2 = 1$. The structure of this system highlights the fundamental difference from ODEs: while the first four equations are differential in nature, the final equation is purely algebraic and creates a manifold on which the solution must evolve. Two primary approaches exist for solving DAE systems: direct solution methods and index reduction techniques.

Direct solution methods employ specialized initial value problem (IVP) solvers, typically based on implicit multistep methods. At each time step, these methods transform the discretized system into a nonlinear root-finding problem that simultaneously satisfies both the differential equations and the algebraic constraints. While this approach can ensure the algebraic constraints are met, to within specified tolerances, it requires solving large nonlinear systems at each time step, making it computationally intensive.

Index reduction approaches, on the other hand, transform higher-index DAEs into lower-index systems through systematic differentiation of the algebraic constraints. The differentiation index of a DAE system is defined as the minimum number of times the algebraic constraints must be differentiated to obtain an explicit ODE system. This index serves as a measure of the mathematical complexity and numerical challenges associated with the DAE system. For instance, in the pendulum system, differentiating the position constraint $x^2+y^2 = 1$ with respect to time yields a velocity constraint. Further differentiation produces an acceleration constraint, which can be rearranged for an explicit equation for λ. A final differentiation yields an explicit equation for λ'. These equations, and the resulting index of the DAE after each differentiation are shown below.

$$\begin{aligned}
&\text{index 3:} && 0 = x^2 + y^2 - 1 && \text{(length constraint)} && (2.63)\\
&\text{index 2:} && 0 = xu + yv && \text{(tangential motion)} && (2.64)\\
&\text{index 1:} && 0 = u^2 + v^2 + \lambda - gy && \text{(centripetal acceleration)} && (2.65)\\
&\text{index 0:} && \lambda' = -3gv && \text{(ODE)} && (2.66)
\end{aligned}$$

Each differentiation step reveals additional hidden constraints inherent in the system. For the pendulum example, differentiating the position constraint reveals that velocities must be tangential to the motion of the pendulum ($xu + yv = 0$).

Converting the DAE to an ODE by replacing the original algebraic constraint with its differentiated form offers significant computational advantages by eliminates the nonlinear root-finding at each time step and allows for the use of standard ODE solvers. However,

this transformation comes at the cost of the original algebraic constraints no longer being explicitly enforced, and errors in these constraints can accumulate over time. This drift can lead to significant violations of the physical constraints the original DAE modelled.

A hybrid approach offers a practical compromise between direct DAE solution and complete index reduction. By retaining an intermediate constraint equation (such as the tangential velocity constraint) rather than the original position constraint, the system can be solved more efficiently while maintaining better control over overall error than with a full index reduction. While this approach does not guarantee exact satisfaction of the original constraints, it can provide a better balance between computational efficiency and solution accuracy. The resulting lower-index problem typically exhibits better numerical properties while still capturing essential physical constraints.

For this work, the resulting network of PFRs produces, after discretization with the FDM, are each modelled as DAEs because of the boundary condition used for each PFR:

$$C_i(0,t) = \frac{\sum_j^{inlets} Q_j C_j(V_j,t)}{\sum_j^{inlets} Q_j} \tag{2.67}$$

The system is of index 1 since differentiating Equation (2.67) once will produce an explicit equation for $\frac{\partial C_i(0,t)}{\partial t}$:

$$\frac{\partial C_i(0,t)}{\partial t} = \frac{\sum_j^{inlets} Q_j \frac{\partial C_j(V_j,t)}{\partial t}}{\sum_j^{inlets} Q_j} \tag{2.68}$$

The algebraic constraint in Equation (2.67) was transformed into the ordinary differential equation Equation (2.68) through index reduction. This transformation was motivated by both computational efficiency and practical implementation considerations. Direct DAE solution methods, while potentially more robust, would require specialized solvers that which present installation challenges for potential users of the CM.

This index reduction approach also introduces an important mathematical constraint on the simulation: boundary conditions must be differentiable in time since they appear in the right-hand side of Equation (2.67). While this might appear to limit the applicability of the model, realistic boundary conditions do not exhibit true discontinuities. In practice, apparent discontinuities in real systems are better represented by smooth, rapidly varying functions. This smoothness not only better reflects physical reality but also improves numerical performance, particularly for adaptive time-stepping schemes that rely on error estimation. Such schemes can maintain specified error tolerances more effectively when working with smooth functions rather than discontinuous ones.

Chapter 3

Literature Review

This chapter examines the current state-of-the-art in compartmental modelling, focusing on five key aspects: compartmentalization, compartment modelling strategies, dynamic compartments, topology-based methods, and the treatment of zero-convective flux planes. While CMs find applications across diverse fields, from pharmaceutical pharmacokinetics and biomedical systems such as artificial pancreas [63, 64] to building ventilation analysis [65], this review focuses primarily on stirred-tank reactor applications. For a more detailed overview, the reader is directed to the work of Haag et al. [17] and Jourdan et al. [16]. Although significant advances have been made in compartmentalization methods for stirred-tank reactors, important limitations and knowledge gaps persist. These will be critically examined to establish the context for the novel methods presented in this thesis.

3.1 Compartmentalization

The process of creating a CM consists of two main steps: compartmentalization and compartment modelling. Recent literature has predominantly focused on different compartmentalization methods. The first class of approaches to compartmentalization uses a uniform structured grid [53, 54, 66], as shown in Figure 3.1. This approach is notable for its simplicity, requiring only the specification of the number of compartments along each dimension. The work of Le Nepvou De Carfort et al. [66] exemplifies this approach, with the procedure following these key steps:

1. A FVM simulation is performed to solve the Navier-Stokes equations.

2. The resulting velocity vector field is then projected onto a coarser structured grid of compartments, with each compartment modelled as a CSTR.

3. The transient mass conservation equation is solved across the compartmental network using the CFD-derived velocity field.

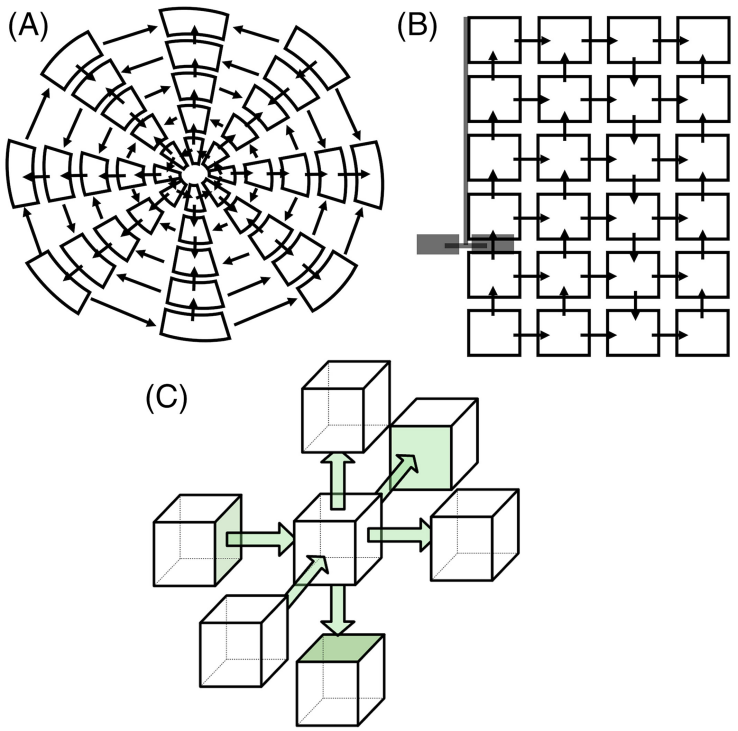


Figure 3.1: Illustrative example of a structured compartmentalization in a cylindrical tank. **(A)** Radial-axial plane. **(B)** Radian-tangential plane. **(C)** Each internal compartment is connected to 6 others. From [54].

Though not framed in these terms by the original researchers, this class of approaches is equivalent to a coarse-mesh, low-order FVM simulation of the mass conservation equations. It is notable that in several implementations, the resolution difference between the structured grid and the original CFD mesh is relatively modest. For instance, Le Nepvou De Carfort et al. [66] employed a CM with 422,500 compartments compared to a CFD simulation with 861,716 mesh elements. While this approach offers an accessible framework for implementation, it presents a significant limitation: the high number of degrees of freedom used make it computationally prohibitive for modelling complex biological reactions.

The analogy between uniform structured grid compartmentalization and uniform structured CFD meshes is apt; both offer straightforward implementation but may require excessive discretization to achieve acceptable accuracy. Alternative compartmentalization strategies employ non-uniform and unstructured approaches, where compartments are formed through the aggregation of mesh cells based on various scalar or vector fields. A notable gap in this literature is the lack of comparative studies between different compartmentalization techniques. Most validation efforts are limited to comparing proposed CMs against one of three benchmarks: single-CSTR approximations, full CFD simulations, or experimental data. Delafosse et al. [47] present a notable exception by conducting a systematic comparison between different velocity-based compartmentalizations, utilizing individual velocity components, velocity magnitude, and uniform spatial discretization. Their findings are particularly significant: the uniform discretization demonstrated superior performance not only at equivalent compartment counts but even when compared to velocity-based approaches using four times as many compartments. Their findings indicated that automated approaches performed less effectively, primarily due to challenges in delineating zones that could reasonably be characterized as well-mixed. These results highlight two critical aspects of compartmental modelling: first, the fundamental importance of the compartmentalization strategy, and second, the non-intuitive nature of performance optimization in these models

The parameters used for compartmentalization include, but are not limited to, fluid velocity, solid/gas fraction, turbulence metrics, reactant concentration, reaction rates, temperature, residence times, streamlines, expert flow knowledge, and residence time distributions [17, 16, 21]. In bioreactor applications, velocity emerges as the predominant parameter for compartmentalization. This preference stems from the velocity field's relative independence from cell metabolism-induced changes, in contrast to parameters such as species concentrations. This independence enables the resulting CM to maintain validity across a broader range of operating conditions. However, these various fields are often used as mere guides in a manual process [51, 44, 48, 21, 50, 67, 68]. The justification for manual approaches varies among researchers, with only a subset explicitly providing a rationale. In their work on a stirred tank reactor, Nauha et al. adopted a manual compartmentalization strategy based on their established understanding of flow patterns, which they characterized as "quiet predictable", arguing that this approach enabled them to minimize compartment count while maintaining consistent compartmentalization across different scenarios [48]. Their compartmentalization scheme is illustrated in Figure 3.2. Their emphasis on compartment count minimization and desire to control compartmentalization-dependent variations in CM predictions aligns with Delafosse et al.'s earlier observations [47], underscoring the necessity for more robust automated compartmentalization methods. In

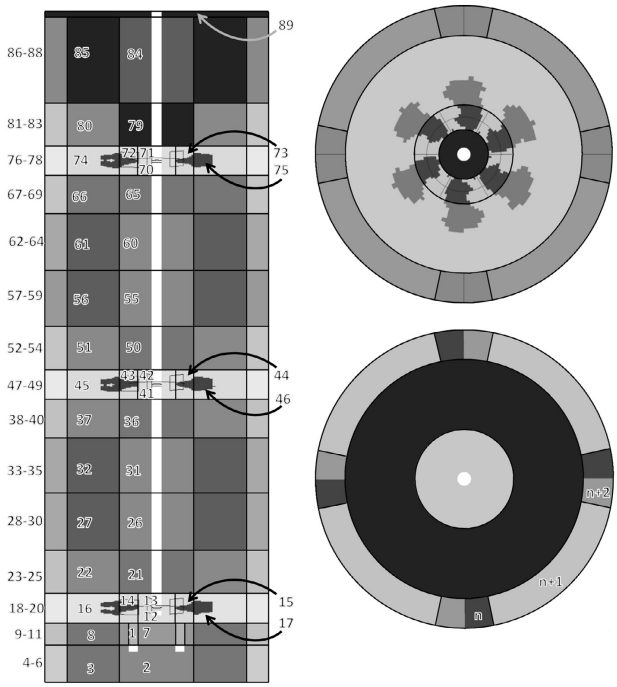


Figure 3.2: Illustration of the manual custom-made compartmentalization of the stirred tank used by Nauha et al. from different angles. **Left**: vertical cross-section of the compartments showing the overall structure. **Top-right:** horizontal cross-section at the plane of the impellers, showing special irregular compartments for each impeller blade. **Bottom-right:** horizontal cross-section in all other parts of the tank. The compartments are large toroids that span the full 360° of the reactor. From [48].

subsequent research, Nauha [49] encountered limitations when attempting to implement their automatic compartmentalization method due to limitations of the proprietary CFD software being used. This experience highlights a crucial consideration in research software selection: the importance of open-source platforms that allow necessary functional modifications. This principle extends beyond CFD to encompass compartmental modelling software, where flexibility and adaptability are essential for research applications.

As previously noted, the rationale for manual compartmentalization is not explicitly stated in all publications. However, two probable explanations can be derived from the literature. Based on the experience gained during this thesis and the criteria cited by various authors for manual compartment definition, the first significant factor appears to be pragmatism. The translation of expert knowledge into compartment boundary placement

through visual inspection of flow fields is considerably more straightforward than codifying this decision-making process into an implementable algorithm.

Secondly, several studies focused not on advancing compartmentalization methods, but rather on demonstrating the broader applicability of CMs as a modelling approach for specific applications [67]. In this context, given the absence of readily available automated compartmentalization tools, manual methods represent a practical choice. Moreover, considering the previously discussed challenges of automatic methods in matching manual compartmentalization performance, the manual approach provides an optimal baseline for evaluating the fundamental utility of compartmental modelling.

Manual compartmentalization approaches present two significant limitations. First, the CM's accuracy becomes heavily dependent on the user's prior knowledge and insights [47], analogous to how the effectiveness of local mesh refinement depends on the criteria used for element modification. This creates a paradoxical situation, as the initial motivation for model development often stems from a need to understand systems where such detailed knowledge is lacking. Second, manual compartmentalization demands substantial user intervention and effort. This parallels the challenges of manual CFD mesh generation; necessary when automatic meshing tools are inadequate, but ultimately a manual time-intensive task that should have been automated. This limitation becomes particularly acute for iterative tasks, such as design and control optimization, where compartment structures may require repeated modification throughout the process.

Existing automatic compartmentalization methods present their own set of limitations. Several previously developed approaches are either constrained by geometric specificity, limiting their broader applicability, or are presented as frameworks that require expert-level knowledge for effective implementation [44, 48, 69, 70]. An effective compartmentalization method must extend beyond mere computational implementability; it should provide sufficient prescription such that users need only specify their desired position on the accuracy-speed trade-off spectrum, rather than requiring them to determine the fundamental basis fields for compartmentalization.

The compartmentalization framework introduced by Tajsoleiman et al. [24] illustrates both the promise and inherent limitations of automatic methods that rely on user-selected basis fields. While theoretically applicable to arbitrary geometries, the framework lacks sufficient prescription for straightforward adaptation without system-specific expertise. Their implementation for a stirred tank geometry, shown in Figure 3.3, employs radial and axial velocity components as compartmentalization fields. However, these represent just one option among many potential choices, including: Cartesian velocity components, vorticity, pressure, turbulent energy dissipation rates, and shear rates.

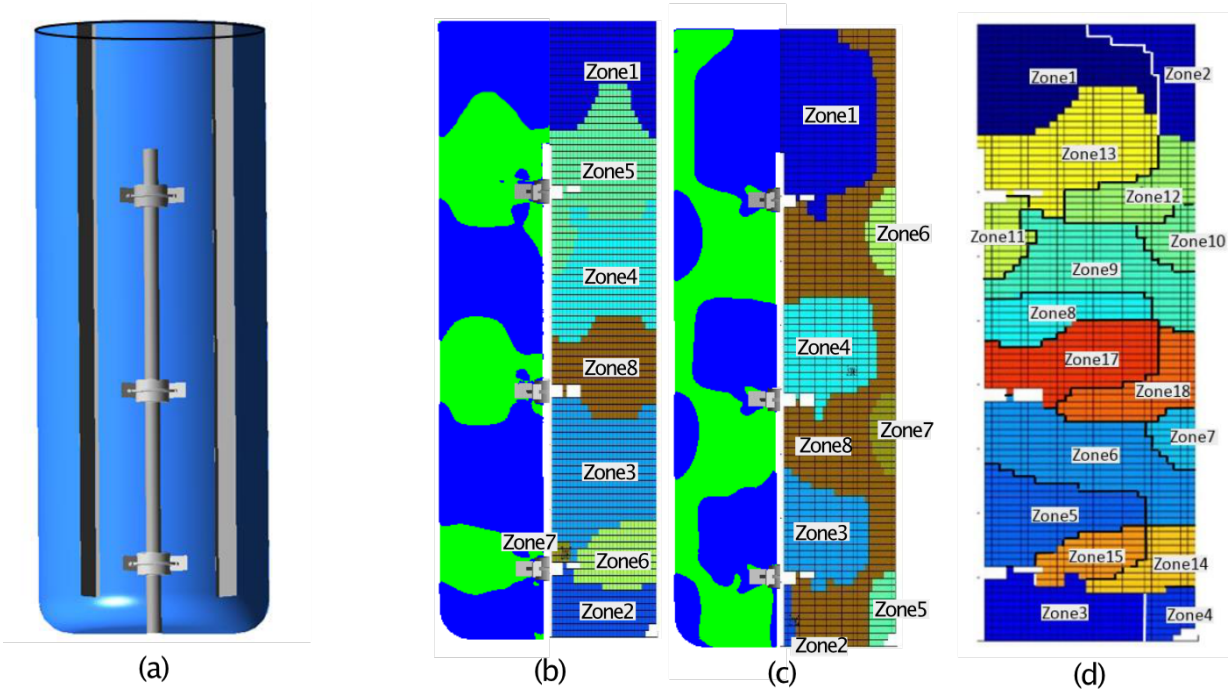


Figure 3.3: Overview of the automatic compartmentalization of a stirred tank reactor using the method of Tajsoleiman et al. **(a)** The geometry of the domain showing 3 bottom mounted Rushton turbines and 4 baffles. **(b)** Zones based on changes in radial velocity. **(c)** Zones based on changes in axial velocity. **(d)** Compartments resulting from intersecting the zones in (b) and (c). From [24]

The framework provides no systematic guidance for selecting among possible alternatives. Given that compartmentalization directly influences CM effectiveness, the framework's utility becomes heavily dependent on the user's ability to select appropriate basis fields. This dependence on user expertise not only increases implementation complexity but restricts the framework's broader adoption in the field.

Savarese et al. [71] proposed an alternative approach using k-means clustering to aggregate mesh elements into spatially contiguous compartments. While innovative, their method requires species concentrations and temperature distributions from combustion reactions as clustering parameters. This requirement necessitates solving complex reaction systems, or at minimum, a simplified representative system, within the CFD simulation prior to compartmentalization. Such an approach presents significant computational overhead and, more importantly, fails to leverage the predominantly one-way coupling between

hydrodynamics and reactions characteristic of many bioreactor systems. This limitation is particularly relevant in bioreactor applications, where the separation of hydrodynamic and reaction time scales typically allows for more efficient strategies.

Several researchers [50, 68, 72] have developed CMs based on one-dimensional input-output relationships, characterizing system behaviour through RTDs and temporal pressure variations. While these models successfully replicate input-output behaviours, they present a significant limitation: the resulting CM lacks clear spatial correspondence with the physical system. This disconnect between model structure and physical reality potentially limits the model's predictive capabilities for local phenomena and its applicability to design optimization.

A fundamental limitation of existing methods is their reliance on user-defined thresholds that significantly influence compartment count and model performance. The selection of these thresholds appears to often require ad-hoc tuning. Some researchers have adopted an iterative approach, adjusting thresholds either to match experimental data [73, 74] or to achieve compartment-independent solutions [46]. Both approaches typically result in models containing thousands of compartments, which undermines the computational efficiency advantage of CMs. While increasing compartment count does improve model accuracy, the global refinement strategy often used is inherently inefficient. This parallels developments in CFD simulations, where local mesh refinement is largely preferred to global refinement due to its superior computational efficiency in achieving improved simulation accuracy. Local refinement strategically allocates computational resources to regions where they provide the greatest benefit. Similar principles could be applied to compartmental modelling, focusing refinement on regions where additional resolution would most effectively improve model performance while maintaining computational tractability for biological kinetics calculations.

3.2 Modelling of Compartments

Contemporary compartmental modelling literature predominantly employs the CSTR approximation for individual compartments [16]. The appropriateness of this approximation is intrinsically linked to both the compartmentalization method and the variables of interest in the analysis. For instance, when compartmentalization is based on the ratio of transport-to-reaction timescale, i.e. the Damköhler number, the CSTR approximation may adequately represent concentration distributions. However, this approximation is less appropriate when compartmentalization relies solely on velocity, currently the most com-

mon approach in the literature. The following discussion examines alternative modelling approaches that move beyond the standard CSTR approximation.

Several researchers have explored hybrid approaches incorporating both CSTRs and PFRs, using a variety of selection criteria. One set of approaches relies on expert knowledge of flow characteristics [75, 70, 67, 76] Another relies on qualitative assessment of flow fields and temporal response data [50, 68, 72]. An example of this second approach is illustrated in Figure 3.4.

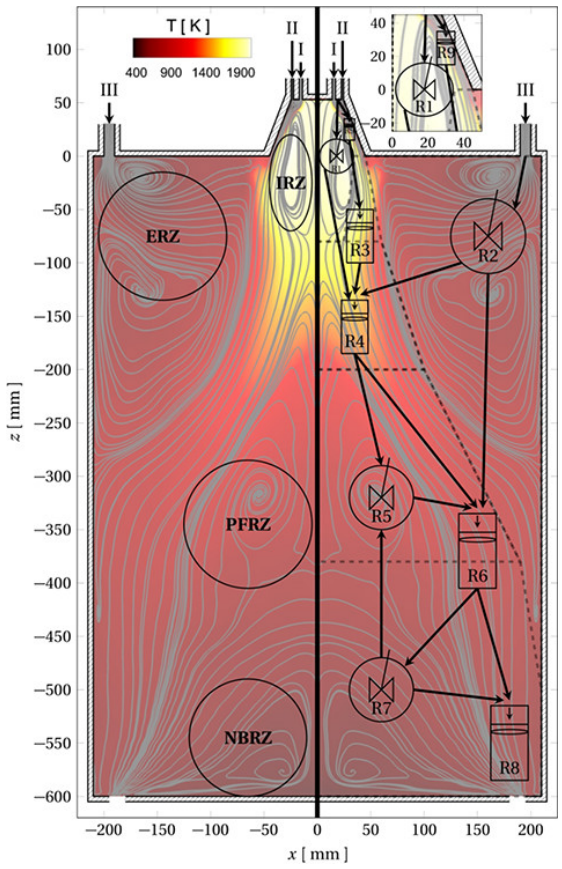


Figure 3.4: Manual compartmentalization of the flow profile in a burner. **Left:** Three main regions are identified based on a visual inspection of the flow profile. **Right:** Each region is split into multiple compartments, and each compartment is modelled by one or more CSTRs/PFRs, with volumes and flowrates fitting to simulated data. From [50].

Although these studies utilize CFD simulations for initial compartment identification, they determine compartment volumes and inter-compartmental flow rates through opti-

mization against experimental RTDs or system outputs. The necessity of this optimization step may indicate a mismatch between the location and connections of the identified compartments as the underlying velocity field should contain sufficient information to directly determine those values. Further, the reliance on ad-hoc optimization suggests a disconnect between compartment identification and compartment modelling which could be addressed through a more integrated approach.

Du et al. [44] conducted a systematic investigation into the effects of switching between PFR and CSTR representations within their CM. Their findings demonstrated that a hybrid approach, utilizing both reactor types, achieved superior performance compared to one which uses only a single reactor type. These results highlight the need for careful consideration of local flow characteristics when choosing a compartment model. This underscores the importance of maintaining consistency between the physical flow structure represented by each compartment and its mathematical model representation.

Notwithstanding the advances described above, most proposed compartmental modelling methods exhibit a fundamental inconsistency between their compartmentalization criteria and the mathematical models employed for the resulting network. This inconsistency manifests in several critical limitations: inability to resolve internal gradients, loss of spatial fidelity, and artificial acceleration of material transport through the domain, particularly evident in under-predicted mixing times. The use of zero-dimensional approximations for compartment modelling introduces significant errors in local reaction rate predictions and fails to capture metabolic responses to extreme concentration values [22]. This limitation becomes particularly problematic in regions of high concentration gradients, such as those surrounding feed ports, or in areas experiencing extreme pH variations near injection points. Resolving this inconsistency constitutes a primary objective of the method presented in this work.

3.3 Dynamic Compartments

A critical limitation of current compartmental modelling approaches is the insufficient investigation of dynamic compartments; those whose geometry, properties, and/or flows evolve in time. This limitation is particularly relevant for fed-batch bioreactors, where both the moving impeller and the changing working volume necessarily alter the underlying compartmentalization structure. However, existing research has made initial steps toward addressing dynamic conditions.

Tajsoleiman et al. [24] evaluated compartmentalization robustness across varying impeller speeds (100–600 RPM) by scaling inter-compartmental flows proportionally to im-

peller power input. Their findings revealed that extrapolated CMs underestimated CFD mixing times by approximately 30%, compared to 20% for CMs specifically constructed at given RPM values. However, their investigation did not extend to volume variations characteristic of fed-batch operations. While Jourdan et al. [16] proposed a theoretical framework for dynamic compartment identification using dimensionless numbers, the concept remains untested.

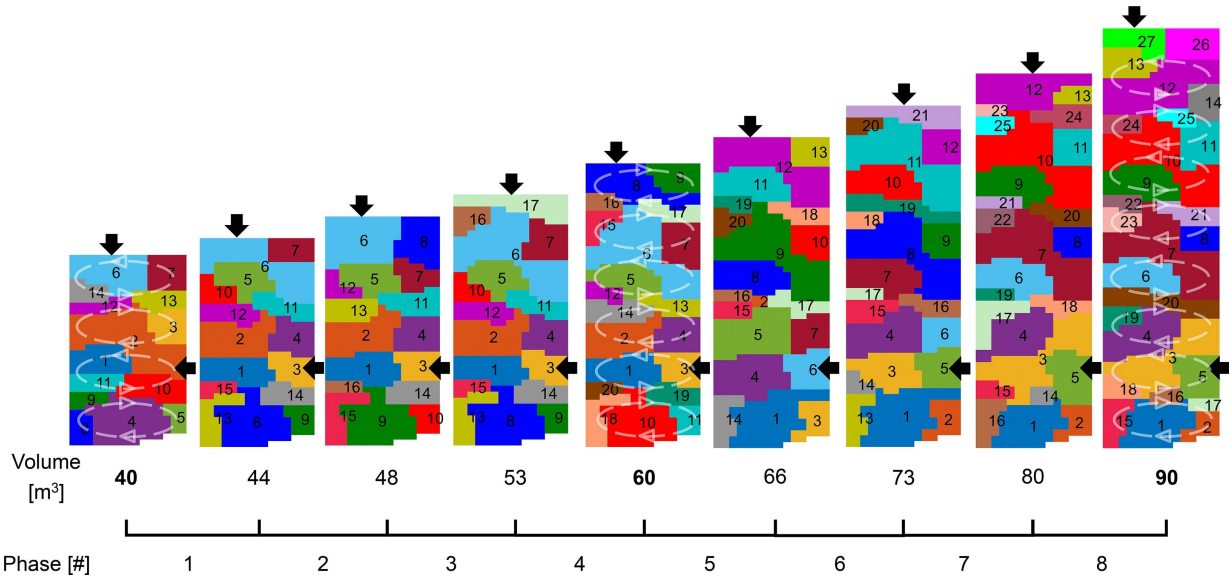


Figure 3.5: Illustration of how changing reactor working volume can be incorporated into a CM method. CFD results are available at 40, 60, and 90 m^3 and used to create CM. For volumes in-between, the CM is created by interpolating between the two nearest CFD-based CM. Current volume is tracked, and the compartmentalization is changed only once the volume has changed by a set amount from its starting value. From [77].

Nadal-Rey et al. [19] appears to have presented the first published implementation of dynamic compartments for fed-batch operations, as illustrated in Figure 3.5 Their approach involved developing discrete CMs at various working volumes and implementing step-wise volume changes to approximate the continuous feeding process. While innovative, their model showed significant limitations, with residence time underestimation up to 72% and yield predictions exhibiting mean absolute percentage errors of 57% (σ = 22%). These substantial deviations were attributed to necessary model simplifications, including limited compartment numbers and assumptions of rotational symmetry, particularly problematic for systems with off-centre feed ports.

Recent work by Morchain et al. [78] investigated dynamic flow field evolution by implementing CFD-informed modifications to inter-compartmental flow rates. While this approach accounts for temporal variations in flow patterns, it constrains the compartmentalization to a static uniform grid of CSTR-compartments. This limitation becomes particularly significant when flow field modifications alter the fundamental structure of the system, as the fixed compartment boundaries cannot adapt.

These studies collectively highlight a significant research opportunity: the development of CMs capable of continuous temporal evolution. Such advancements would particularly benefit applications involving variable working volumes or changing operating conditions where current static or discretely updated models show limitations. Further, the same techniques could be used to model time-varying flow, such as caused my long-timescale oscillations, without restrictions on the location and shape of the compartments.

3.4 Other Topology-Based Methods

Existing compartmental modelling approaches can be broadly categorized into three dominant groups: expert-driven manual compartmentalization, uniform grid-based divisions, and similarity-based mesh element clustering. While these approaches have demonstrated utility, they often lack systematic consideration of fundamental flow characteristics. This thesis presents a distinct approach by employing flow topology analysis, specifically focusing on unidirectional alignment, as a basis for compartment identification.

The use of topological features for compartmentalization represents a significant departure from conventional methods. Unlike traditional methods that rely on expert intuition, uniform spatial division, or similarity-based clustering, topology-based approaches leverage fundamental flow characteristics to inform compartment creation, potentially leading to more physically meaningful reduced-order models. By explicitly considering unidirectional alignment in the flow field, the method provides a more mechanistic basis for compartment definition compared to traditional geometric or statistical clustering approaches. While not entirely novel, as evidenced by a small number of previous topology-based studies, this approach remains relatively unexplored.

The pioneering work in this approach appears to be that of Fichet et al. [79], who introduced a two-stage compartmentalization method. Their method first employs inert tracers, uniquely valued at each domain inlet, to delineate regions of influence associated with distinct inlet streams. This initial partitioning is then refined using fluid age distribution analysis, where each compartment is subdivided into a series of CSTRs based on fluid age uniformity, as illustrated in Figure 3.6

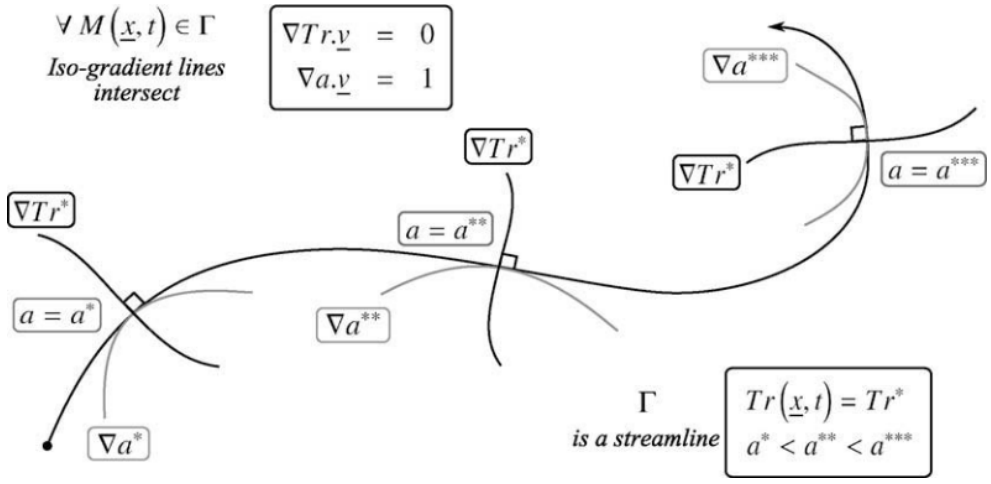


Figure 3.6: Inert tracers a and Tr are used to identify compartments and choose how many CSTRs-in-series to model each compartment with. Tr is injected with a unique value from each domain inlet, allowing for regions of influence to be identified. Streamlines of constant value Tr are used to delineate compartments. a is used to quantify the age of the fluid, and split each streamline into pieces of roughly uniform age. From [79].

The approach presented in this thesis, while sharing the fundamental principle of using flow topology for compartment identification, differs in its specific focus on unidirectional alignment. This represents an evolution in topology-based compartmentalization methods, moving from inlet-influenced zones and age-based segregation to a more comprehensive consideration of flow field structure, allowing for the modelling of closed and semi-closed systems where the concept of fluid age has less descriptive power.

Next is the work of Donnelly [21], who developed a novel approach drawing parallels to liquid crystal modelling. This method analyzes non-local topological features through the examination of a velocity direction field, $\boldsymbol{n}$, derived from steady-state velocity data. The approach characterizes flow structure through three fundamental deformation metrics:

$$\text{Splay} = (\boldsymbol{\nabla} \cdot \boldsymbol{n})^2 \tag{3.1}$$

$$\text{Twist} = (\boldsymbol{n} \cdot \boldsymbol{\nabla} \times \boldsymbol{n})^2 \tag{3.2}$$

$$\text{Bend} = (\boldsymbol{n} \times \boldsymbol{\nabla} \times \boldsymbol{n})^2 \tag{3.3}$$

These deformation metrics, through their magnitudes and relative proportions, provide a basis for identifying distinct flow regions and their corresponding compartmental

representations, as illustrated in Figure 3.7. The method implements a hierarchical modelling approach where compartments are represented either as single CSTRs or, in regions of strong unidirectional flow, as CSTRs-in-series. This work serves as an inspiration for the method developed in this thesis, which builds upon and extends these concepts of topological flow characterization.

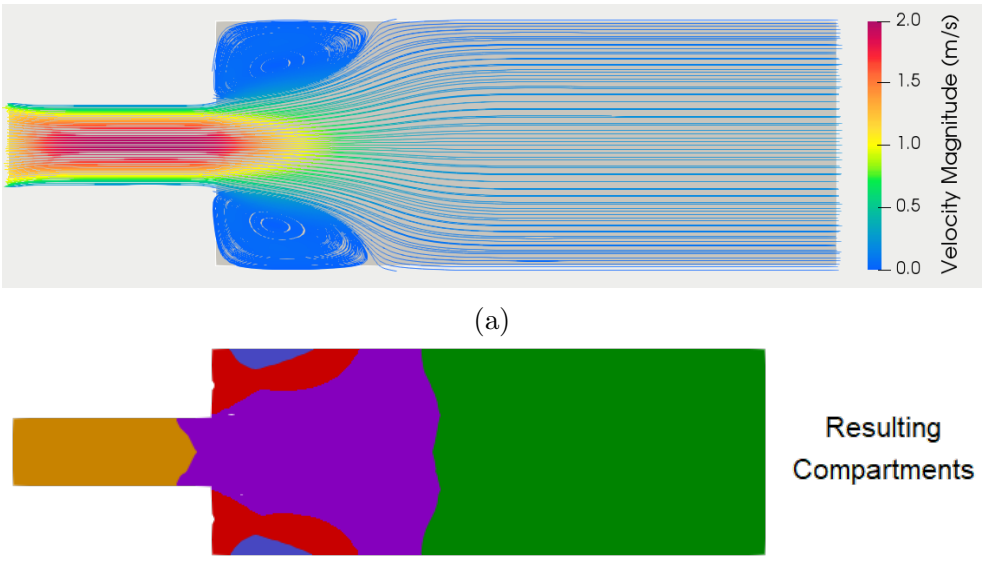


Figure 3.7: Cross-section of the CFD and compartmentalization results for flow in a pipe with a sudden expansion. **(a)** Flow topology visualized through velocity streamlines coloured by velocity magnitude. **(b)** Resulting compartments: orange - inlet, green - outlet, purple - splay, red - recirculation, blue - stagnation zone. From [21].

The most recent significant contribution to topology-based compartmentalization comes from Karches [80], who developed a method based on critical point analysis of steady-state flow fields. This approach identifies critical points within the flow domain and uses their interconnecting streamlines to define compartment boundaries, as depicted in Figure 3.8. By using critical points and their connecting streamlines, the method captures fundamental flow features that often correspond to important mixing boundaries or flow separation regions. However, as most existing work, each compartment is modelled as a single CSTR.

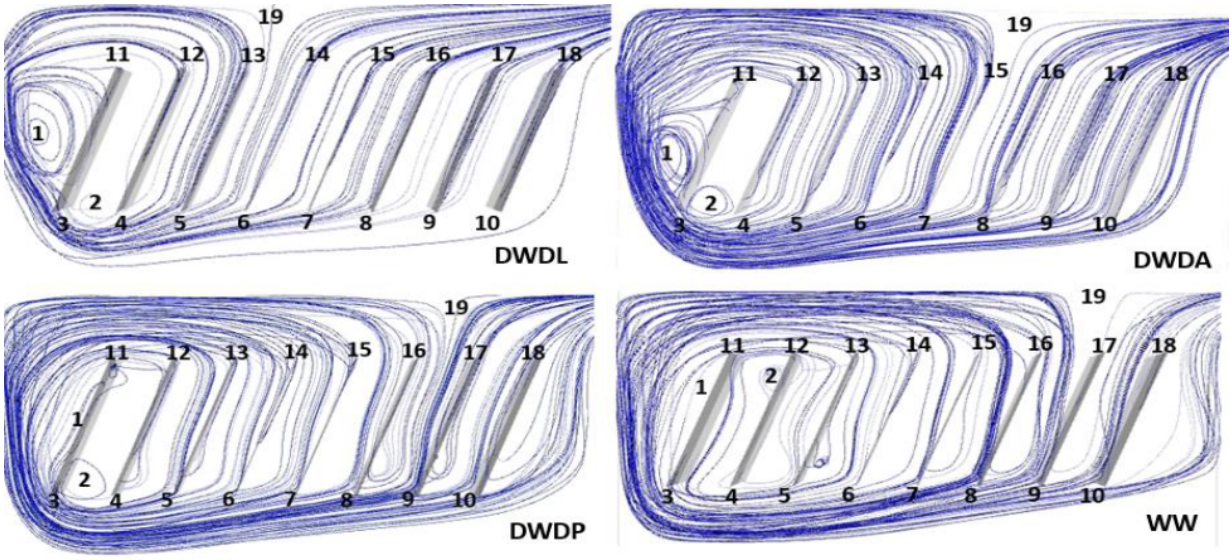


Figure 3.8: Streamlines visualizing flow topology of the same geometry under 4 different conditions. Critical points used for compartmentalization are labelled as 1–19. From [80].

The work of Karches, along with the previously discussed approaches of Fichet et al. [79] and Donnelly [21], establishes a growing body of literature supporting the utility of topology-based compartmentalization methods. These approaches collectively demonstrate the advantages of using fundamental flow characteristics, rather than arbitrary geometric divisions or purely statistical clustering, as the basis for compartmentalization.

3.5 Handling of Zero-flux Planes

Recent studies into compartmental modelling of stirred tank reactors identified limitations in accurately representing key flow features. Studies by Bai et al. [55] and Promma et al. [12] highlighted challenges in predicting mixing times and attribute these inaccuracies to the inadequate representation of horizontal planes at recirculatory region boundaries generated by Rushton turbines. As illustrated in Figure 3.9, these critical planes occur at two locations: at the impeller outlet planes and at the midpoint between impellers.

Promma et al. [12] observed that, without manual intervention, their CM under-predicted mixing times. They identified compartments extending across these horizontal planes, as illustrated in Figure 3.9, as the root cause of this issue. These spanning compartments, the orange and yellow regions, created artificial "short-cutting" effects, allowing

unrealistically rapid mass transfer between the upper and lower recirculatory zones. Their solution was to manually introduce cut planes at critical heights and subdividing compartments that intersect the planes. This modification substantially improved mixing time predictions, increasing them from 60.05 s to 189.34 s, in-line with expected values.

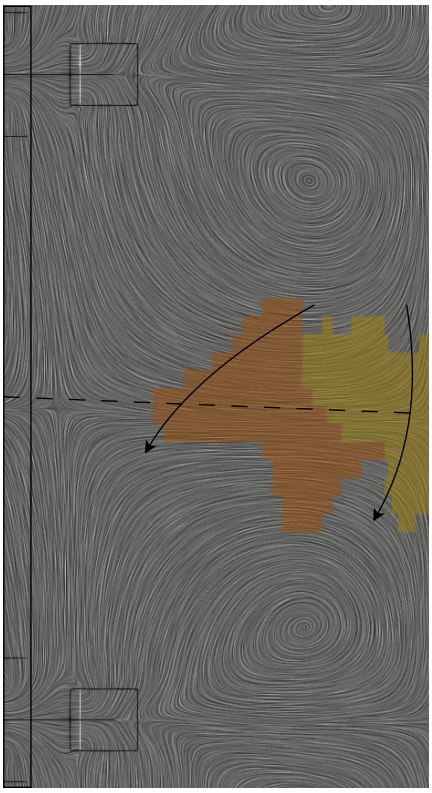


Figure 3.9: Visualization of the flow topology, using a line integral convolution, of the region between two Rushton impellers from a stirred tank MRF-based CFD simulation. The dashed horizontal line in the middle of the figure indicates the horizontal plane that forms at which the velocity, and thus convective transport, is in the horizontal direction. The yellow and orange regions represent two different CSTR-based compartments which span either side of this plane. The use of the CSTR model for each compartment causes "short-cutting", i.e. causes mass to move artificially quickly across the plane. From [12].

Bai et al. [55] encountered similar challenges with horizontal planes in their CM, though with notably different outcomes from Promma et al. While both studies identified these planes as critical to mixing time accuracy, Bai et al. observed mixing time over-prediction

rather than under-prediction. Their focus was on the planes aligned with impellers, rather those in the inter-impeller region‘. Most significantly, their solution directly contradicted Promma et al.’s approach: they emphasized the importance of avoiding compartment boundary alignment with these planes, rather than enforcing such alignment.

The contradictory mixing times and solutions can be attributed to fundamental differences in inter-compartmental flow calculations. Bai et al. employ a simplified approach using net flows between compartments, while Promma et al. implement a more detailed method that preserves bidirectional transport. When compartment boundaries align with horizontal planes, Bai et al.’s net-flow approach results in artificially reduced transport rates due to flow cancellation. Conversely, Promma et al.’s preservation of gross flows maintains local mixing effects across boundaries. This difference explains why Bai et al. avoid boundary alignment with these planes, while Promma et al. specifically enforce it.

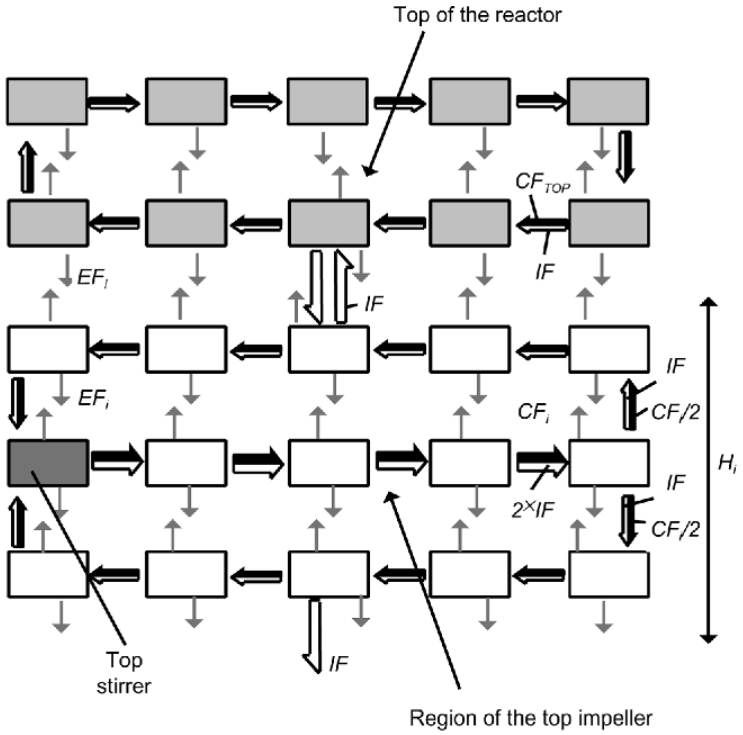


Figure 3.10: Structured compartmentalization of the recirculatory region around the top impeller. The distribution of compartments, their connections, and flowrates are obtained from design equations and expert knowledge of the system. From [3].

The work of Vrabél et al. [3] and Pigou et al. [22] appears to have anticipated key as-

pects of compartmental modelling strategies that would later be formally identified by Bai et al. and Promma et al. Their compartmentalization scheme, illustrated in Figure 3.10, incorporates both aligned compartments at the inter-impeller region and offset compartments at the impeller plane, features that align with later recommendations. However, these choices appear to have been made through expert intuition rather than systematic analysis, as the explicit consideration of horizontal plane effects is absent from their discussion. Significantly, their approach demonstrates both the strengths and limitations of expert-driven compartmentalization. While their model effectively captured key flow features, its reliance on prior knowledge of specific tank-impeller configurations limits its adaptability to design modifications without substantial expert intervention.

The issue discussed above highlight critical limitations of current compartmental modelling methods. First, there exists a fundamental disconnect between the well-mixed assumption used in compartment models and the criteria used for compartment identification, leading to overestimation of mixing within compartments. Second, existing methods fail to adequately account for directional information flow. As illustrated in Figure 3.9, compartments that span across directionally-distinct flow regions, such as the orange compartment, create artificial mixing paths. While mass transport in these regions is predominantly horizontal, with minimal vertical convective transfer across the marked plane, the compartment structure incorrectly suggests significant vertical mixing. Third, the treatment of inter-compartmental flows, either as net or gross flows, significantly impacts how CSTR compartments interact with no-flux planes, requiring case-specific solutions. This issue exemplifies a broader limitation of current approaches: most rely on global element clustering without considering actual flow connectivity, resulting in compartments that may not reflect true mass transport pathways.

While manual interventions can address these issues in simple geometries through explicit plane definitions, e.g. $z = z_{\text{plane}}$, such solutions become impractical for complex, irregular geometries. This limitation emphasizes the need for automated methods that inherently account for directional flow characteristics.

Chapter 4

Compartmental Modelling Methods

This chapter proposes a novel compartmental modelling method for general multiphase three-dimensional geometries which addresses the limitations identified in Chapter 3. Two distinct approaches are developed to handle systems whose time-averaged flows are not representative of their time-varying flows.

4.1 Method for Single-Phase 2D Systems

This section[1] presents a CFD-based compartmental modelling method which addresses several of the previously identified gaps in the literature. The key differences of the proposed method compared to previously reported methods are: (1) the compartmentalization criteria is automated, (2) the compartments can be approximated with one-, two-, or three-dimensional models, (3) the reduced-order models for compartments are chosen based on the compartmentalization criteria rather than expert knowledge, and (4) the proposed method allows for the results of the compartmental model simulations to be easily projected back onto the physical domain, greatly improving the interpretation of model predictions. The high-level steps in this proposed method are as follows:

1. **Pre-processing:** Convert data into a standardized intermediary format.
2. **Compartmentalization:** Split the domain into a network of compartments, each representing a region of unidirectional flow.

[1] Adapted from [81].

3. **Compartment Modelling:** Model each compartment as a series of 1D PFRs.

4. **Simulation:** Discretize and solve the system of coupled mass balance equations.

The method creates compartments by merging mesh elements through the use of a CFD-derived velocity field. Best practice for performing CFD simulations is to obtain a mesh-independent solution in order to ensure that all relevant phenomena are accurately resolved. This thesis argues, and the results presented herein back up this argument, that a mesh-independent solution is not required for the purpose of building a compartmental model. The model error introduced by the simplifying assumptions of the CM, e.g. the commonly used assumption that each compartment is well mixed, is significantly higher than the error associated with fine, but not mesh-independent, resolution CFD simulation.

4.1.1 Pre-processing

To work with the simulation results from different CFD packages, a standardized representation of the mesh and of the resulting field is used. In this format, each field is defined as having a single value within each mesh element. In order to accommodate results from both high-order FEM packages, e.g. `OpenCMP` [82], an additional step is required to capture the spatial variation within each mesh element. For low-order computational methods, such as the FVM, this pre-processing step is not needed. Consider Figure 4.1; the left side shows the higher order FEM solution being projected onto a zeroth-order interpolant.

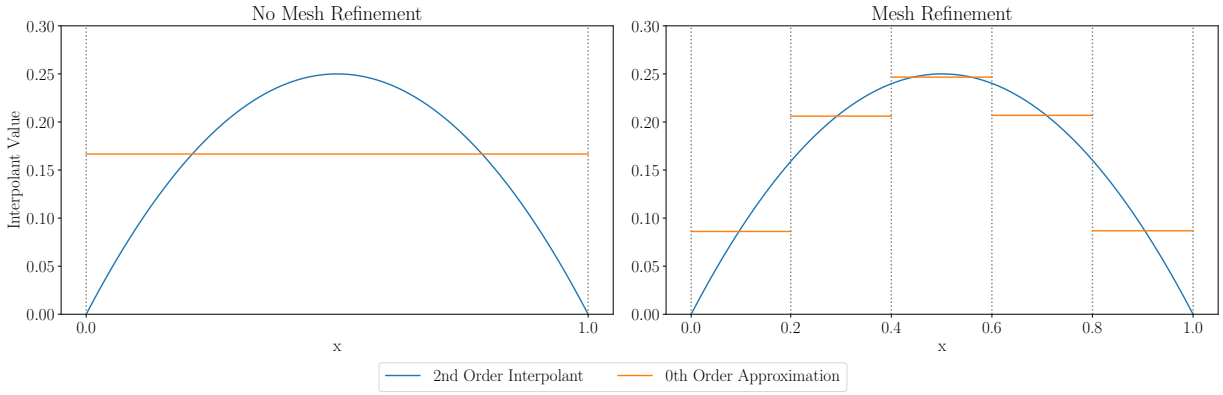


Figure 4.1: Projection of the higher-order FEM result onto a zero-order interpolant for a one-dimensional mesh. The original higher-order FEM interpolant is shown in blue while the final zero-order interpolant is shown in orange. **Left:** Projection on the original mesh. **Right:** Projection on the refined mesh.

The zero-order interpolation result, shown in orange in Figure 4.1, inherently loses detailed spatial variation information within each element. To preserve these spatial characteristics, mesh refinement techniques can be employed by subdividing elements to enhance the resolution of the zero-order interpolant, as shown on the right side of Figure 4.1. This refinement process can be implemented through either global- or local-refinement.

Global refinement involves uniform subdivision of all mesh elements, creating an equal number of new elements throughout the domain. While straightforward to implement, this approach can be computationally inefficient as it applies the same level of refinement to all elements, regardless of whether such refinement is necessary. Local refinement offers a more selective approach by utilizing a scalar labelling function. In this application, the function could be based on the curvature of the interpolant: elements whose interpolant has minimal curvature require little to no refinement for accurate zero-order representation, while those with high curvature require finer discretization. For this work, global refinement was selected despite its higher computational overhead in pre- and post-processing stages, primarily due to its implementation simplicity. The development of local refinement strategies remains a promising direction for future work to optimize computational efficiency.

4.1.2 Compartmentalization

Using the data in the standardized format, the compartmentalization step involves grouping mesh elements, mesh cells, into compartments with locally unidirectional flow. This process begins with a set of elements, whose locations are user-specified at the domain inlet(s), e.g. feed ports. These elements are used as the "seeds" for generating new compartments through a sequentially iterative procedure involving the following 6 steps.

Step 1: all neighbouring elements of the current compartment which are not in another compartment are grouped into a set of available neighbours.

Step 2: the angle, θ, between the direction vector of the seed element and that of each available neighbour is calculated.

Step 3: all neighbours with an angle $\theta > \theta_t$ are removed from the set of neighbours and added to the set of seeds. Comparing to seed element direction guarantees that the difference in flow directions between any two elements in a compartment is $\leq 2\theta_t$, enforcing a user-specified degree of unidirectional flow. Step 3 is illustrated in Figure 4.2 for two elements.

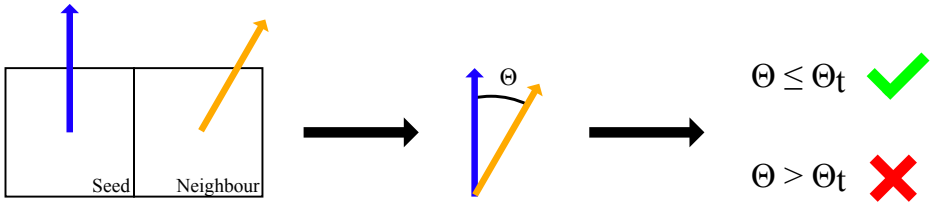


Figure 4.2: Rejecting a neighbour based on the angle between the direction vector of the compartment's seed element and that of the neighbour element. The blue and orange arrows represent the direction vectors of the seed and neighbouring elements, respectively.

Step 4: the remaining neighbours are checked for two criteria. First, all elements which are upstream of the compartment are excluded, as shown in Figure 4.3.

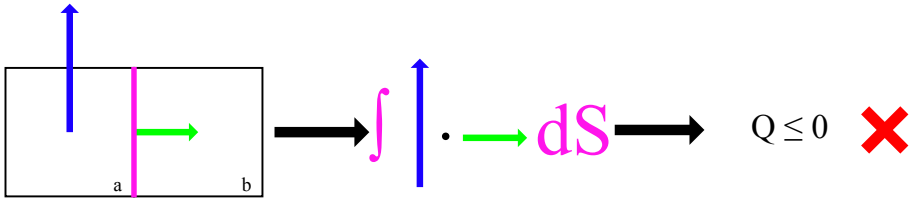


Figure 4.3: Rejecting a neighbouring element, B, because it is upstream of the compartment. The blue arrow is the velocity vector of an element a in the compartment, and the green arrow is the outward facing unit normal to the purple facet between a and b.

Second, all elements which have an inadequate amount of flow between them and the compartment are excluded. An element b is considered to have an adequate amount of flow from compartment A if there exists an element a in A where the flow from a to b is non-negligible based on the following inflow into element b:

$$\frac{Q_{a\to b}}{\sum_i Q_{i\to b}} \geq threshold \tag{4.1}$$

and the outflow requirement for element a is,

$$\frac{Q_{a\to b}}{\sum_j Q_{a\to j}} \geq threshold \tag{4.2}$$

where *threshold* is user specified.

An example of this constraint is illustrated in Figure 4.4.

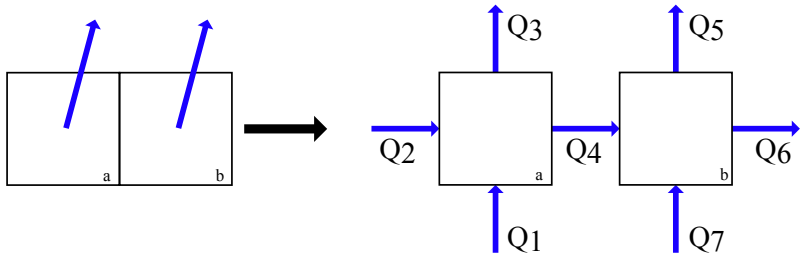


Figure 4.4: Example showing if the flow from a to b is adequate. Blue arrows on the left represent the velocity vector for each mesh element. On the right, the blue arrows represent the flowrate through each facet of the elements, labelled as Q1–Q7.

Following the example in Figure 4.4, the above equations become:

$$\frac{Q_4}{Q_4 + Q_7} \geq threshold \tag{4.3}$$

$$\frac{Q_4}{Q_3 + Q_4} \geq threshold \tag{4.4}$$

Any neighbours which satisfy these two constraints are added to the current compartment. Neighbours that do not satisfy these constraints are appended to the list of seeds for use in creation of new compartments.

Steps 1–4 above are repeated until no further neighbours can be added to the current compartment. The process iterates by selecting subsequent seeds from the list until all elements are assigned to a compartment. The algorithm prioritizes the selection of the most upstream unassigned element that has been previously seen. This upstream-biased selection is strategically important because compartment growth only occurs in the downstream direction. Selection of upstream elements maximizes the potential extent of each compartment, preventing unnecessary fragmentation that would occur if downstream elements were selected as seeds. Such fragmentation would result in multiple smaller compartments where a single larger compartment would have been possible.

Step 5: convert the resulting compartments into a network. The first step is to identifying the mesh facets which make up the bounding surface of each compartment. Each element is bound by, and connected to other elements through, their facets. For two-dimensional meshes, facets are line segments and for three-dimensional mesh they are most

often triangles, quadrilaterals, or hexagons. Figure 4.5 shows an illustrative example with five neighbouring compartments ($A - E$) and how the bounding facets of compartment A are determined, i.e. by finding all facets which are shared by only one element in the compartment.

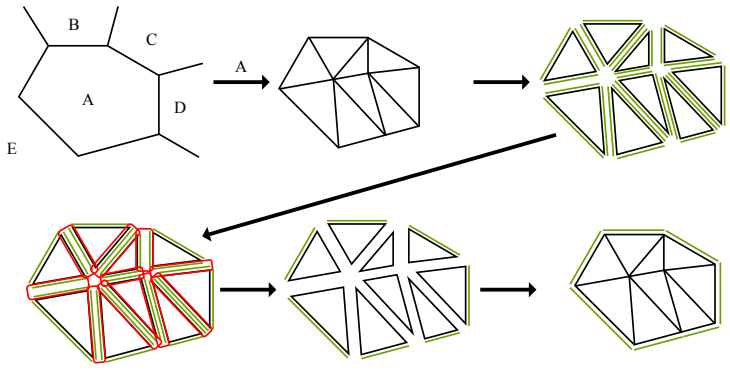


Figure 4.5: Identifying the bounding facets for compartment A. The facets of all elements in A are collected. Internal facets will appear twice, once for each element in the compartment it shares. Bounding facets are identified as those appearing only once.

The bounding surface is then split into contiguous sub-surfaces based on the adjacent compartments, as shown in Figure 4.6.

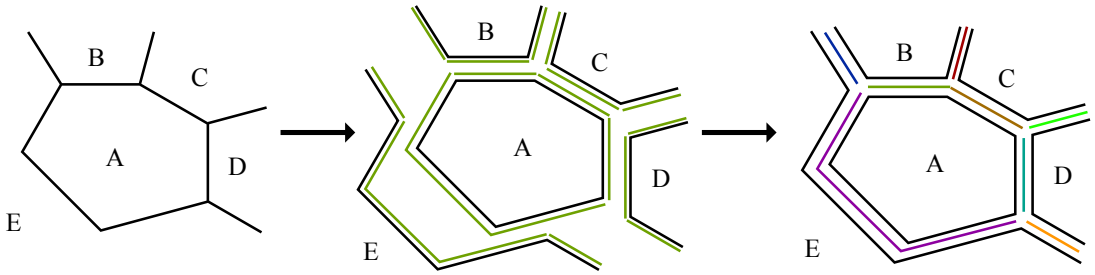


Figure 4.6: Splitting the bounding surface of compartments A through E based on the compartment on the other side. **Left:** Compartments identified. **Middle:** Bounding facets for each compartment identified in green. **Right:** Bounding facets split grouped, shown using colour, based on which two compartments they are between.

These sub-surfaces are then used to calculate connections between compartments and

and build the graph. Figure 4.7 contains the graph representation of the compartments shown in Figure 4.6.

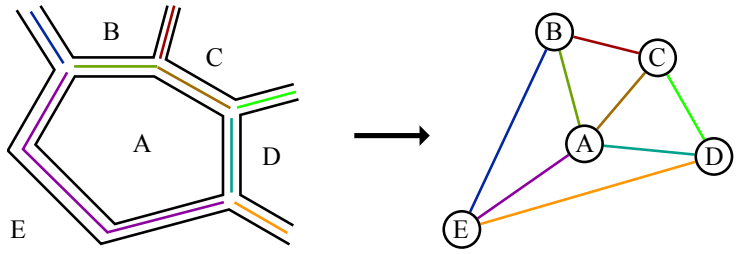


Figure 4.7: **Left:** Compartment network shown as labelled regions in physical space. **Right:** The same compartment network represented as an undirected graph.

Step 6: merge compartments in the network to ensure certain guarantees. First, any compartment smaller than a user-specified threshold is merged into neighbouring compartments with the most similar average direction. Further, if possible, only downstream compartments will merged into. Upstream neighbours are considered only if no downstream ones are available. This is done to reduce the computational cost of the resulting CM since small compartments are expected to have negligible influence on the accuracy of model relative to their additional computational cost. Then, the network is checked to ensure that no compartment has only one connection, either to another compartment or a domain inlet/outlet. If any are found, they are merged using the merging procedure described above. This is done to ensure that mass can be balanced around each compartment, something which would be impossible if a compartment had only one inlet or one outlet. This process is iterated until all compartments pass both checks. This is illustrated in Figure 4.8 where three neighbouring compartments are merged into one.

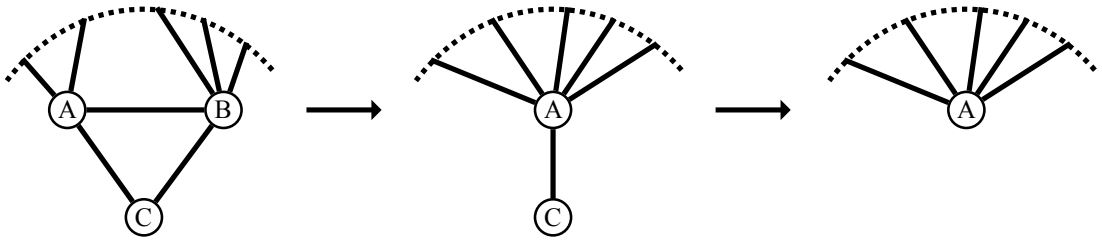


Figure 4.8: Illustrative example of the merging of compartments. After compartment B is merged into A, compartment C is left with a single connection, with A. This is unphysical as mass cannot be balanced around C, thus compartment C is then also merged into A.

4.1.3 Compartment Modelling

For the purposes of this work, each compartment is modelled as a series of 1D PFRs. Note that the use of 1D PFRs was a choice and not an inherent part of the proposed method. The only restriction imposed by the method is that the model must represent unidirectional flow. As long as unidirectional flow assumption is maintained, the model for each compartment could be a one-, two-, or three-dimensional, allowing for spatial variations perpendicular to the main flow direction. Dispersion perpendicular to the main flow direction can also be included.

Furthermore, to provide better resolution of the concentration profile, each compartment is approximated as a series of PFRs. Accordingly, for each compartment, the number of PFRs, the size of each PFR, and the location of inlets and outlets for each compartment is calculated as per the following steps:

Step 1: the surface connecting two compartments is split into sub-surfaces, one facet at a time, based on the flow through each facet. This is done to capture the location of side flows into compartments, as well as to avoid zero net flows between compartments. All facets which have a flow through them smaller than a user-specified threshold are neglected. The remaining facets are grouped together into spatially contiguous sub-surfaces which all have flow in the same direction. If any of the sub-surfaces are below a user-specified threshold, they too are neglected. If all sub-surfaces are removed based on these criteria, then the corresponding connection between the two compartments is removed from the compartment network. This is illustrated in Figure 4.9.

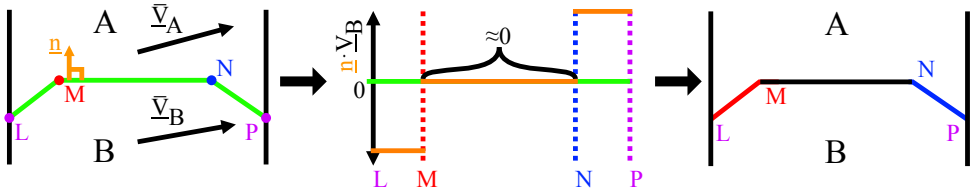


Figure 4.9: Illustrative example of splitting the surface between compartments A and B based on the sign of the flux. The sub-surface between points L and M has net flow from $A \rightarrow B$, the portion from M to N has no net flow and was removed, and the portion from N to P has a net flow from $B \rightarrow A$.

Step 2: the spatial ordering of the connections for each compartment is calculated. For each sub-surface, a position vector is computed by calculating the flow-weighted average position vectors of all facets within it. These sub-surface position vectors are then

projected onto the average direction vector of the compartment. The resulting scalar values are normalized between $[0.0, 1.0]$, with 0.0 corresponding to the first, most upstream, connection to the compartment and 1.0 to the last, most downstream, connection. This is illustrated in Figure 4.10.

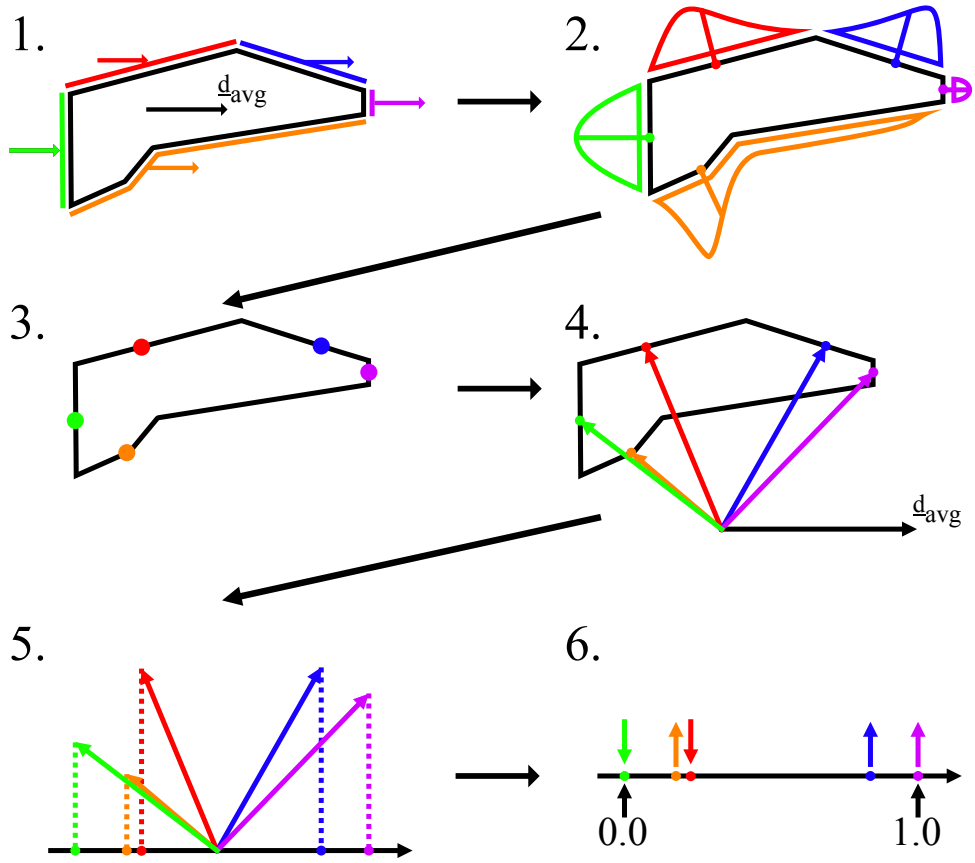


Figure 4.10: Illustrative example of the computation of spatial ordering of compartment connections. **1.** Compartment with each sub-surface represented by a different colour. Arrows indicate the direction of flow through the surface. $\underline{d}_{avg}$ is the average direction vector of the compartment. **2.** Spatial distribution of flow through each surfaces. **3.** Flow-weighted representative position of each surface. **4.** Position vector of each surface connected to average direction vector of the compartment. **5.** Position vectors projected onto the average direction vector. **6.** Dot-product results from 5. normalized from 0.0, most upstream, to 1.0, most downstream. Upward facing arrows represent inlets into the compartment, and downward facing arrows represent outlets.

Step 3: adjust flowrates to preserve conservation of mass. The pruning of sub-surfaces as well as the approximate nature of the CFD solution, requires adjusting the flowrates through each connection to ensure mass is conserved around each compartment. Further, a subsequent step, step 5, requires that no compartment have a net outflow of mass. Thus a preliminary flowrate optimization routine is required to ensure that no compartment has a net outflow; later a second optimization will be performed which ensures that mass is perfectly balanced around each PFR. This essentially involves imposing conservation of mass on the compartment with flowrates adjusted as follows,

$$Q'_i = Q_i + \Delta Q_i \tag{4.5}$$

where Q' is the adjusted flowrate, Q the original flowrate, and ΔQ is the non-negative adjustment to the flowrate. A non-negative adjustment is used in order to keep the optimization problem simple. Allowing adjustments to be negative either requires a non-linear objective function to keep the solution space bounded, or a mixed integer problem with different adjustment variables for positive and negative flowrate adjustments.

These adjusted flowrates are based on the mass conservation constraint for compartment k,

$$\sum_{i}^{inlets} Q'_i - \sum_{j}^{outlets} Q'_j \geq 0 \tag{4.6}$$

$$\sum_{i}^{inlets} (Q_i + \Delta Q_i) - \sum_{j}^{outlets} (Q_j + \Delta Q_j) \geq 0 \tag{4.7}$$

$$\underbrace{\sum_{i}^{inlets} Q_i - \sum_{j}^{outlets} Q_j}_{\varepsilon_k} + \sum_{i}^{inlets} \Delta Q_i - \sum_{j}^{outlets} \Delta Q_j \geq 0 \tag{4.8}$$

$$\sum_{i}^{inlets} \Delta Q_i - \sum_{j}^{outlets} \Delta Q_j \geq -\varepsilon_k \tag{4.9}$$

where ε_k is the net inflow into compartment k. Equation (4.9) is formulated as an inequality to ensure that there cannot be net outflow from the compartment, which is needed by a subsequent step of the method.

Conservation of mass is imposed through the solution of a linear optimization problem

where the *total amount of adjustment* is minimized as follows,

$$\begin{aligned} \underset{\Delta Q_i}{\text{minimize}} \quad & \sum_{i}^{N} \Delta Q_i \\ \text{subject to} \quad & \sum_{i}^{inlets} \Delta Q_i - \sum_{j}^{outlets} \Delta Q_j \geq -\varepsilon_k, \quad k \in [1, M] \\ & \Delta Q_i \geq 0, \quad i \in [1, N] \end{aligned} \tag{4.10}$$

where M is the number of compartments and N is the number of compartment connections.

Step 4: Adjusting connection locations. Following this optimization, the location of the connections found in Figure 4.10 are checked against the following criteria and adjusted if they do not satisfy them collectively:

- An inlet is the first connection and an outlet is the last connection.
- Domain inlets are the first connections, and domain outlets the last.
- All domain inlet/outlet connections occur sequentially.
- Connections are ordered such that backflow does not occur at any point in the compartment. If a compartment has a net outflow, it cannot be guaranteed that this requirement can be satisfied, hence the need for step 3. Mor information on this point is provided in step 4.

An additional constraint is imposed to reduce the computational cost of the resulting compartment model; sequences of connections which are closer to each other than a user-specified tolerance are moved to be at the same location. Figure 4.11 shows an example of this process where the creation of a sequence of PFRs with and without merging of connections is shown. In this example, moving of connections reduced the sequence of PFRs from four to two, halving the computational complexity without significant loss of accuracy.

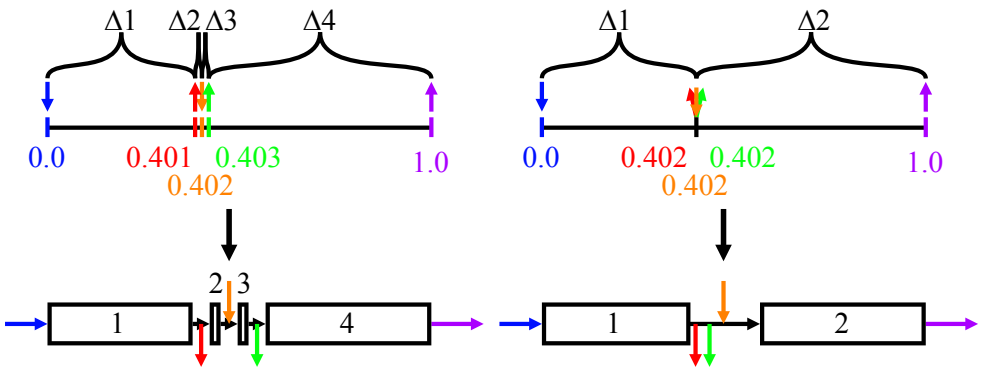


Figure 4.11: Illustration of the resulting PFR-in-series from a compartment with 5 connections, 3 of which are spatially close to each other. **Left:** Without merging, this configuration produces 4 PFRs, 2 of which are significantly smaller than the rest. **Right:** With merging, this configuration produces 2 PFRs, halving the number of degrees of freedom for this compartment.

Step 5: Use the connections to split the compartments into a series of PFRs. The size and location of each PFR in the sequence used to model a compartment is determined by the number of connections and their spatial locations. Specifically, the size of each PFR is the volume of the compartment scaled by the fractional distance between connections. Figure 4.12 shows an example of the PFR sequence used to model a sample compartment.

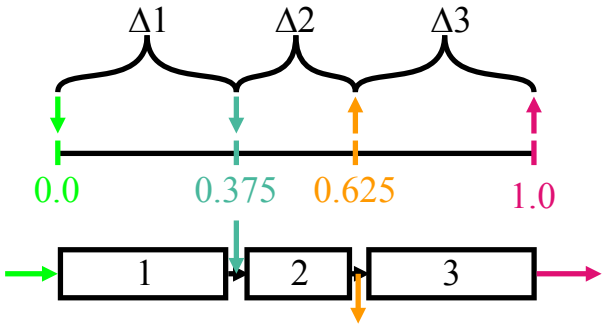


Figure 4.12: The construction of a PFR sequence used to model a compartment with 4 connections, 2 inlets and 2 outlets, with different spatial locations. The result is a sequence of 3 PFRs.

The intra-compartment flow rates, represented by black arrows between PFRs in the Figure 4.11 and Figure 4.12, are determined through mass balance calculations at each

junction point. For instance, the flow rate between PFR 1 and PFR 2 is calculated as the flow entering PFR 1 plus the flow of the dark green side feed. This calculation specifically relies on the guarantees provided by step 4, and thus the corrections performed by step 3, in order to produce valid results. When examining the connection between PFR 2 and 3, the flow rate is calculated as the difference between the inlet flow to PFR 2 and the orange side outlet. To maintain positive flow directionality from PFR 2 to PFR 3, this difference must be positive, requiring that the orange flow be less than the flow into PFR 2. Step 4 ensures that connections are reordered to prevent this flow reversal, but can only be guaranteed to succeed if the compartment does *not* have a net outflow.

Step 6: Final adjustment of flows for overall mass conservation. Unlike step 3, whose goal was to prevent net outflows from any compartment, the goal of this optimization is to perfectly balance mass around each PFR. It is derived the same as the previous optimization problem, with the only difference being that equality, instead of inequality, constraints are used:

$$\begin{aligned} \underset{\Delta Q_i}{\text{minimize}} \quad & \sum_{i}^{N} \Delta Q_i \\ \text{subject to} \quad & \sum_{i}^{inlets} \Delta Q_i - \sum_{j}^{outlets} \Delta Q_j = -\varepsilon_k, \quad k \in [1, M] \\ & \Delta Q_i \geq 0, \quad i \in [1, N] \end{aligned} \tag{4.11}$$

where M is the number of compartments and N is the number of compartment connections.

4.1.4 Validation

In this work, the proposed compartmental model is validated against experimental results reported in the literature and against corresponding CFD simulations.

Residence Time Distribution

To quantitatively compare the different compartmental modelling methods and to validate against CFD simulations, the residence time distribution (RTD) of an inert tracer was used. Given a step change injection of an inert tracer at the inlet of a domain at $t = 0$, the RTD function $E(t)$ describes the fraction of total tracer mass that exits the reactor versus time [83, Chapter 16],

$$E(t) = \frac{dF(t)}{dt} \tag{4.12}$$

The cumulative distribution (CD) function $F(t)$ describes the fraction of total tracer mass that has exited the reactor at time t,

$$F(t) = \int_0^t E(t)\ dt \tag{4.13}$$

For a step change of tracer injected at the inlet, the CD function can be rewritten as,

$$F(t) = \frac{M_{out}(t)}{M_{in}} \tag{4.14}$$

where M_{in} is the mass flow rate into the reactor, assumed constant in time, and $M_{out}(t)$ is the mass flow rate out of the reactor at time t.

Scalar measures that can be computed from the RTD and CD functions include the first three moments:

mean residence time

$$t_m = \int_0^\infty tE(t)\ dt \tag{4.15}$$

variance

$$\sigma^2 = \int_0^\infty (t - t_m)^2 E(t)\ dt \tag{4.16}$$

skewness

$$s^3 = \frac{1}{\sigma^{3/2}} \int_0^\infty (t - t_m)^3 E(t)\ dt \tag{4.17}$$

4.1.5 Kullback-Leibler Divergence

In addition to these RTD moments, the Kullback-Leibler divergence [84], a measure of relative entropy, is also calculated between the two CMs (the PFR-based and the CSTR-based) and the CFD results. This is a measure which enables direct comparison of distribution functions. It is calculated based on the expectation of the logarithmic difference between two distributions:

$$D_{KL}(A||B) = \sum_x A(x) \log\left(\frac{A(x)}{B(x)}\right) \tag{4.18}$$

where $A(x)$ and $B(x)$ are the two distributions. The larger the magnitude of $D_{KL}(A||B)$, the larger the difference between distribution functions.

4.2 Extension to 3D Multi-Phase Closed Systems

The extension of the base method to 3D multiphase flows required only incremental refinements. Several enhancements were also implemented to improve performance for industrially relevant applications. This section details the most significant of these modifications, presenting the most current iteration of the compartmental modelling method.

4.2.1 Convection-Guided Compartment Merging

A key refinement of the method extends the compartment merging criteria to also consider the volumetric flow between compartments, not just the angle between the average flow direction. This modification aligns the merging step with the initial compartmentalization process, which accounts for propagation of information when merging neighbouring mesh elements into an existing compartment. Figure 4.13 demonstrates the advantages of the revised merging method through a simplified three-compartment scenario.

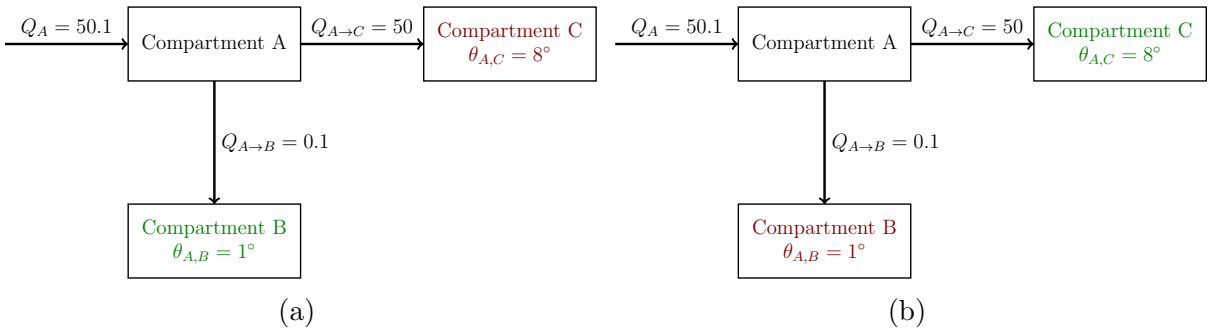


Figure 4.13: Simplified example showing the different merging result between (a) the original method and (b) the new method. Green represents the compartment chosen for merging and red represents the rejected one. **(a)** The original method chooses compartment B despite the insignificant amount of flow between A & B. **(b)** The new method chooses compartment C, despite the higher angle difference, because of the much higher flow between them.

The previous method, which relied solely on flow alignment, would merge compartments A and B due to their nearly parallel average direction vectors. However, this approach overlooked the critical factor of inter-compartmental mass transport. The new method merges compartments A and C, despite their larger directional difference, because on their

significantly higher inter-compartmental flow rate, approximately 500 times greater than the A-B exchange. This choice better reflects the mass transport behaviour of the system, as the dominant flow path from A to C can be more accurately represented by the PFR-in-series model in the resulting merged compartment. This revision ensures that the merged compartments better preserve the underlying transport physics of the original system. IT is particularly important for the capturing of flow behaviours which previous CMs failed to model accurately, as discussed in Section 3.5.

4.2.2 Switch to Facet Fluxes

When calculating the flow through the bounding surface of a compartment, the original method used the mesh elements inside of the boundary to calculate the flowrate through each facet. Applying this method had two problems. Consider Figure 4.14 which depicts two mesh elements, the velocity vector in them, and the surface they share highlighted.

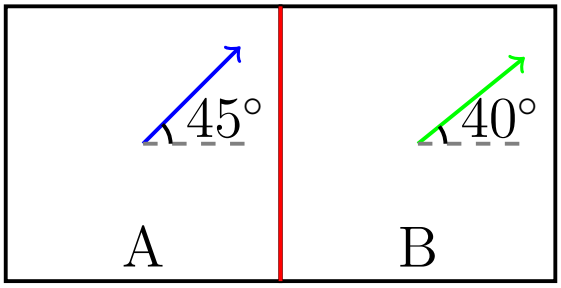


Figure 4.14: Illustration of two adjacent mesh element with different velocity vectors in each. If the element's velocity vector is used for calculating the flow through the red facet, then different values are obtained if calculated from element A versus element B.

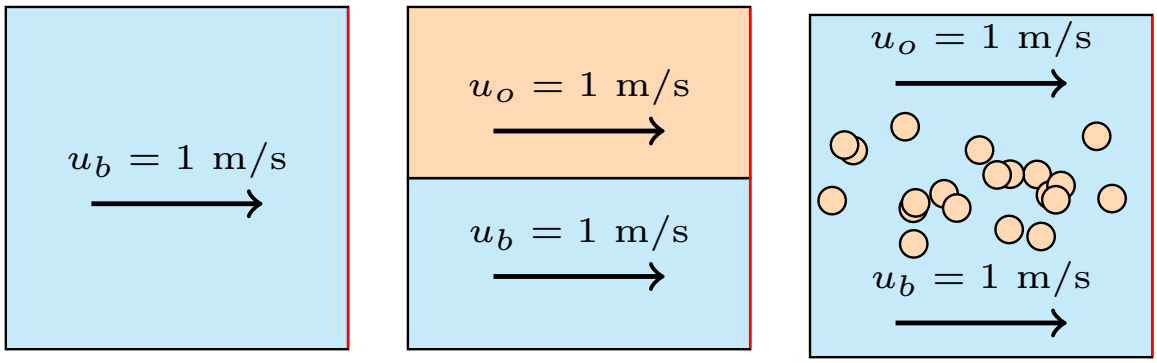
First, volumetric flowrate calculations exhibit directional dependence, producing inconsistent results when evaluated from opposite sides of an interface. Second, this method inadequately represents convective transport mechanisms. Consider element B: mass transport across the reference surface occurs along the blue velocity vector, not the green vector. This discrepancy introduces systematic errors in mass flux calculations. To address this issue the method was changed to use facet-centred velocities. This was done using a first order upwinding scheme; the velocity used for each facet is that of the mesh cell upstream of the facet. This approach ensure consistent flowrate calculations since each facet has a single velocity associated with it, and the upwinding method used better captures the convective transport characteristics of the flow.

4.2.3 Extension to Multi-phase Flows

For single-phase flows, the volumetric flow rate through a surface is calculated as:

$$Q = \int_{\Omega} \boldsymbol{u} \cdot \boldsymbol{n} \, \mathrm{dx} \tag{4.19}$$

where Q is the volumetric flowrate in units of volume/time, $\boldsymbol{u}$ is the velocity in units of length/time, and $\boldsymbol{n}$ is the unit normal to the surface Ω.



(a) Single phase. (b) Stratified two-phase. (c) Dispersed two-phase.

Figure 4.15: Calculation of volumetric flowrate for different flow regimes through the red facet on the right side of each domain. u_o represents the local velocity of the orange phase, and u_b the local velocity of the blue phase. **(a)** For a single phase system the calculation is straight forward. **(b)** For a stratified two-phase system the calculation requires an adjustment to the surface area used in the calculation since the orange phase does not **(c)** For a dispersed two-phase system the area must again be adjusted to account for the presence or absence of individual bubbles/droplets.

The volumetric flux at any point on the surface is given by the dot product $\boldsymbol{u} \cdot \boldsymbol{n}$. For a surface oriented perpendicular to the flow direction, $\boldsymbol{u} \cdot \boldsymbol{n} = |\boldsymbol{u}|$, making the velocity magnitude equal to the volumetric flux, and thus the volumetric flux vector and the velocity vector are equal. However, this equivalence does not hold for multiphase systems. Figure 4.15 illustrates this distinction through three scenarios in unit square domains (1 meter $\times$ 1 meter). The single-phase system, subfigure (a), demonstrates a straightforward case where the volumetric flow rate Q_b through the red surface is 1 m^3/s. For the stratified two-phase flow shown in subfigure (b), Q_b reduces to 0.5 m^3/s through the same surface. This reduction can be represented mathematically in two equivalent ways: by

restricting the integration domain to the blue phase region only, or by using superficial velocities that account for the phase fraction. The dispersed two-phase flow in subfigure (c) follows the same principle, where the volumetric flow rate must account for the spatial distribution of phases. Despite the different flow morphology, the fundamental concept of phase-dependent flow rates is the same as for the stratified case.

The dispersed two-phase simulations in this work employ the Euler-Euler framework, where individual bubbles are not resolved explicitly. Instead, each mesh element contains a phase fraction α_i, representing the volume fraction occupied by phase i. For a surface oriented perpendicular to the flow, the volumetric flux becomes $\alpha_i|\boldsymbol{u}_i|$, where $\boldsymbol{u}_i$ is the phasic velocity and $\alpha_i\boldsymbol{u}_i$ is the volumetric flux vector. Consequently, the volumetric flow rate for phase i through any surface can be expressed as:

$$Q_i = \int_\Omega \alpha_i \boldsymbol{u}_i \cdot \boldsymbol{n} \, \mathrm{dx} \tag{4.20}$$

The compartmental model adopts volumetric flux rather than phasic velocities to maintain mathematical consistency between single- and multi-phase systems. Likewise, the volume of a compartment is calculated using the volume of each mesh element multiplied by the value of α in that element. This approach provides an additional advantage: it inherently compartmentalizes only those regions which contain the phase of interest without requiring additional modifications or steps. Recall that the method employs a two threshold-based filtering steps: excluding computational cells with superficial velocities below a tolerance, and removing inter-compartmental connections with flowrates below a tolerance. Since regions without phase i have zero superficial velocity by definition, they are automatically excluded during the compartmentalization process.

4.2.4 Handling of Closed Systems

The optimization procedure described in Step 3 of Section 4.1.3 ensures that no compartment has a net outflow of mass. This means that all compartments must have mass perfectly balanced, or have a slight inflow. Because compartments cannot have a net outflow, a net inflow in a compartment must be balanced by a net inflow into the system through domain boundaries. This poses a problem for closed systems which lack domain inlets/outlets. In closed systems, conservation of mass requires that any compartment with net inflow be balanced by a corresponding net outflow from other compartments. As illustrated in Figure 4.16(a), if compartment A exhibits net inflow, compartment B must necessarily experience net outflow. However, this net outflow condition conflicts with the

requirements of the subsequent PFR-in-series decomposition outlined in Section 4.1.3, it means that the optimization routine would fail to converge since it would be impossible to satisfy all of the constraints.

To resolve this issue, a temporary virtual domain inlet is introduced and connected to each compartment, as shown in Figure 4.16(b). Following the initial optimization, the virtual connections are removed, and the final optimization is performed on the network of PFRs, ensuring proper mass conservation within each reactor unit.

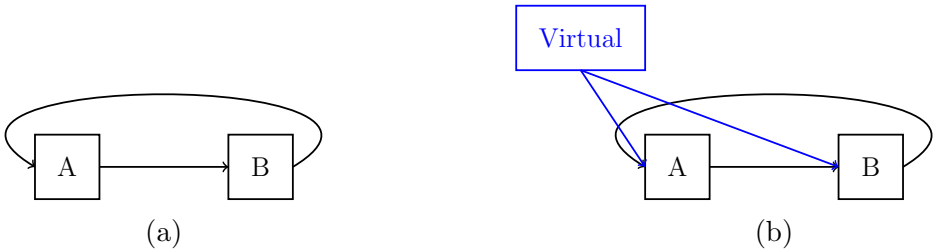


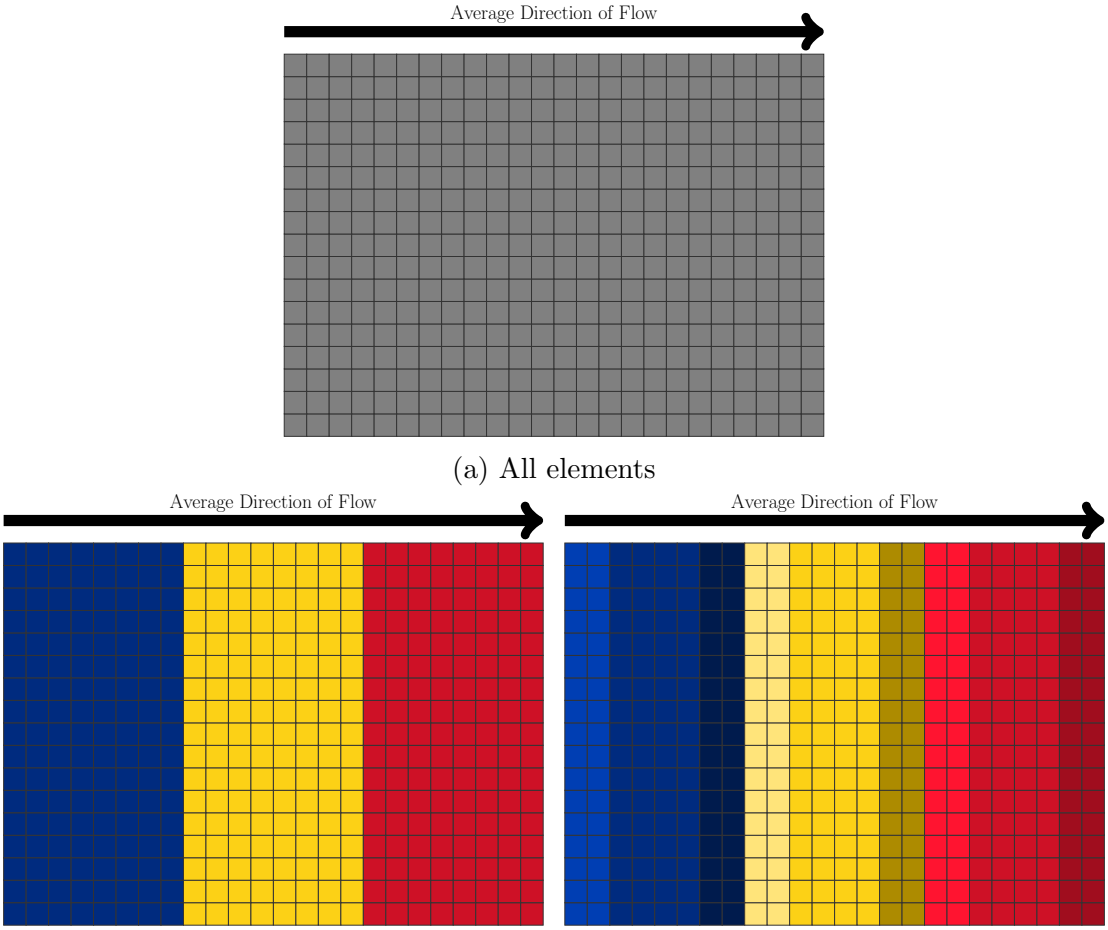
Figure 4.16: Input Compartment networks of a closed system to the optimization in Step 3 in Section 4.1.3. **(a)** Original closed system. **(b)** With virtual connections added it is possible for both compartments to have no net outflow.

4.2.5 Element-to-DOF and DOF-to-Element Mappings

A mapping between CM DoFs and the mesh elements of the original three-dimensional domain enables complex spatially-varying reactions, visualization capabilities, and compatibility with various CFD software formats. Figure 4.17 illustrates the elements of a mesh colours by which PFR and which DoF within the PFR they represent.

A DoF-to-element mapping establishes a piecewise linear interpolation scheme for projecting compartmental model results onto the original mesh. For each DoF, the mapping stores a list where each entry is comprised of: (1) a mesh element, (2) the second-nearest DoF within the same PFR, and (3) an interpolation weight factor. This structure enables both zeroth-order interpolation, using only the nearest neighbour value, and first-order interpolation using the nearest neighbours and the weight factor. The mapping is generated by applying the projection and splitting operations described in Section 4.1.3 to element position vectors, determining their PFR association and specific DoF location. An element-to-DoF mapping, constructed by inverting the DoF-to-element mapping, projects values from the CFD mesh onto the compartmental model. This mapping facilitates local parameter calculations, such as kLa values from CFD simulation results (see Section 6.4),

enables the use of spatially-varying initial conditions, and allows for placement of virtual sensors at precise locations.



(a) All elements

(b) Coloured by PFR

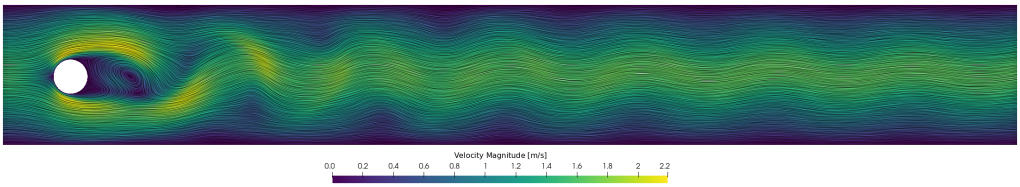
(c) Coloured by DoF

Figure 4.17: Labelling the mesh elements of a compartment made of 3 PFRs with 3 DoFs per PFR. The mapping, (c), allows for CM simulation results to be projected onto the mesh, or spatially varying CFD results on the mesh to be projected onto the CM.

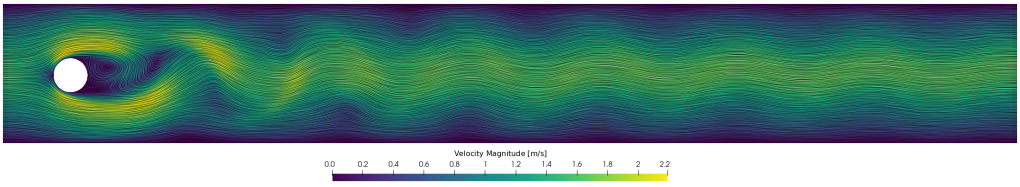
However, because each element can only be mapped to a single DoF, it is possible to obtain CM which have DoFs that do not map to any mesh element. This can occur if there are too many PFRs or too many DoFs per PFR resulting in certain regions having more DoFs than mesh elements. These DoFs are handled by having them share the mapping of the closest DoF and producing a weighed average of the two.

4.3 Multi-CM Methods for Transient Flows

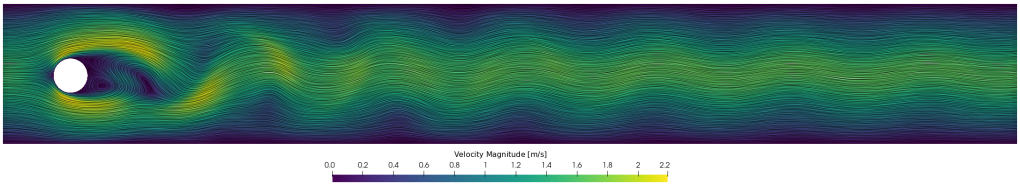
The previous sections established a method for developing CMs based on a single representative velocity field. However, certain systems lack a singular representative velocity field, whether steady-state or time-averaged. Consider the velocity field for the flow past a cylinder simulation shown in Figure 4.18; the vortex shedding is stable and periodic, no fixed steady-state exists. Such cases require an alternative approach to CM development. This section presents two methods for addressing scenarios where traditional single-profile approaches prove inadequate.



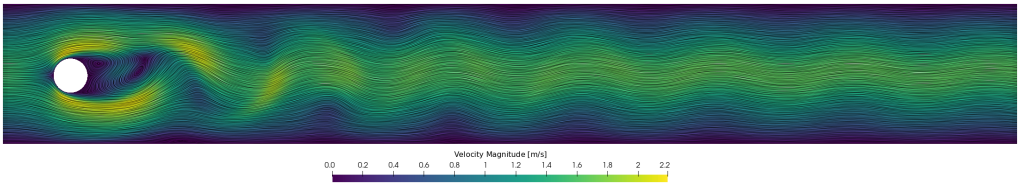
(a) Time = 10.504 s



(b) Time = 10.896 s



(c) Time = 11.680 s



(d) Time = 12.480 s

Figure 4.18: Snapshots of the flow profile showing the characteristic periodic steady-state of the system. Colouring based on velocity magnitude with an LIC filter overlayed.

4.3.1 Flow Snapshot Selection Method

Both of the two methods to be covered by this section are based on using multiple temporal snapshots of the flow. The selection of snapshots requires special care as they will directly influence the quality of the results. In particular, this work focuses on scenarios with a periodic steady-state whose flow topology also changes in time, such as vortex shedding from flows past a cylinder or rotating impeller systems in baffled tanks. These flows exhibit multiple characteristic frequencies. Rather than relying on correlations or assumptions to obtain them, the method analyzes the CFD simulation results to identify the relevant frequencies. It begins with a CFD simulations using sufficient temporal resolution to capture multiple velocity snapshots per period, thereby avoiding aliasing effects. The process involves:

1. Sampling velocity magnitude to establish temporal variations.
2. Performing a Discrete Fourier Transform (DFT) to identify dominant frequencies.
3. Identifying the highest frequency requiring explicit CM representation.
4. Selecting timesteps that ensure multiple samples per period of the highest frequency while spanning a complete period of the lowest frequency and being uniformly spaced.

The selection of frequencies to capture depends on case-specific requirements for model accuracy; there is a direct trade-off between the number of timesteps used and the computational cost of the resulting model. Figure 4.19(a)–(b) illustrates the frequency analysis technique applied to an oscillatory flow. The frequency decomposition reveals two distinct components: a primary frequency at 1 Hz and a secondary frequency at 5 Hz.

Figure 4.19(c)–(d) demonstrates profile reconstruction using multiple samples per minimum period for two cases: (c) incorporating both identified frequencies and (d) using only the dominant frequency for sampling. The subsampling rate is first constrained by the highest frequency selected. Although including both frequencies would likely enhances the accuracy of the CM, it substantially increases computational cost.

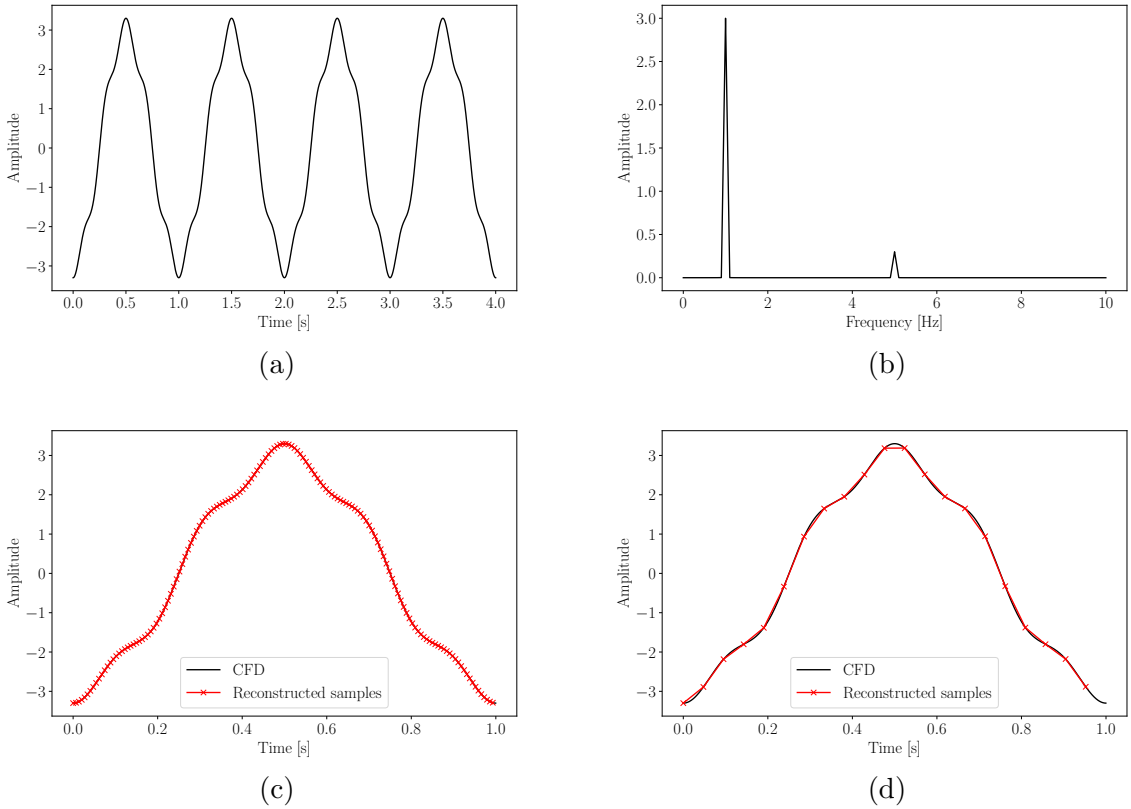


Figure 4.19: Illustration of how a period signal is sampled and its frequencies used to dictate snapshots to use for compartmentalization. **(a)** Compound sine wave showing data for multiple periods. **(b)** Power spectrum of the data. **(c)** The sampled profile based on all identified frequencies. **(d)** The sampled profile based on only the most important frequency.

The second, and equally important, consideration for picking subsamples is that they must remain equally spaced. This requirement imposes constraints on which samples can be selected. Within each period of $K+1$ snapshots, the first and last snapshots (labelled 0 and K) represent identical points in the cycle. The problem of picking every T^{th} point out of $K+1$ snapshots, where 0 and K represent the same point, for a total of N points is equivalent to factoring K into two integers: $K = T \times N$. To obtain the valid options for T all prime factors $\{P_1, P_2, ...\}$, including repeats, of K are calculated. A value for T can then be chosen by multiplying together a subset of the prime factors, and N is found by multiplying together the remaining factors. The resulting N snapshots are renumbered

consecutively from 0 to $N-1$ and are used in both of the proposed approaches.

4.3.2 Snapshot Ensemble CM Method

The first approach generates a CM for each snapshot and runs an independent simulation on each of them. The results are then combined together into an ensemble as a post-processing step. This method aims to establish concentration envelopes that quantify local flow variations and predict expected value bounds. These bounds are designed to encompass oscillatory behaviour while producing mean values that correspond to time-averaged concentrations of the oscillatory flow.

Figure 4.20 illustrates this concept, demonstrating that regions experiencing smaller flow oscillations produce narrower envelopes where mean values better approximate the actual flow behaviour.

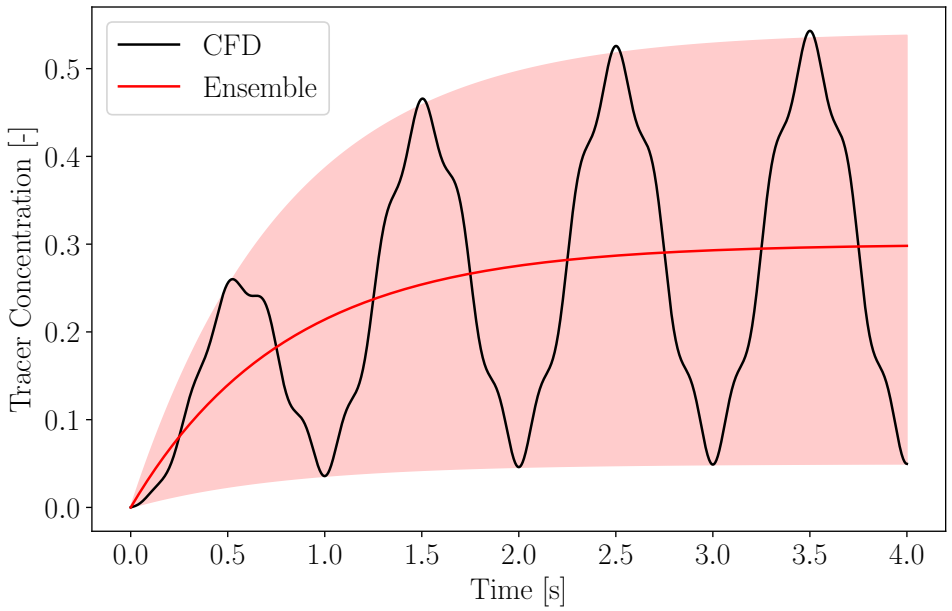


Figure 4.20: Illustration of a hypothetical tracer concentration based on the oscillatory flow described previously, showing how the ensemble creates an envelope that encompasses the oscillations.

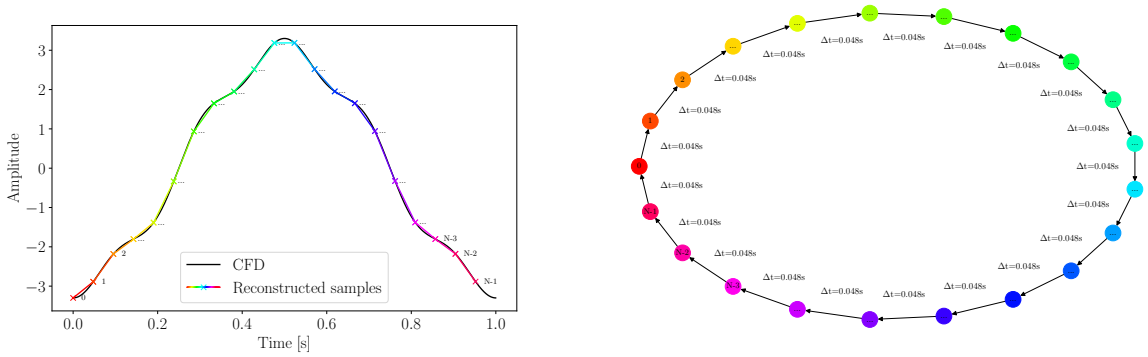
The primary advantage of this approach lies in its parallel execution capability and the

elimination of inter-compartmental mapping requirements. However, it requires the execution of multiple complete CM simulations. In scenarios where error tolerances necessitate small timesteps, the computational cost increases significantly.

4.3.3 Dynamic CM Projection Method

The second approach extends the recent work of Morchain et al. [78], which demonstrated that time-varying connections between CSTR compartments could capture flow oscillations using CFD data. However, their method imposed restrictive assumptions: an unchanging and uniform grid-based compartmentalization with CSTR-based compartments. The proposed method captures transient velocity profiles while eliminating these constraints.

The method is built upon the snapshots identified by the method described in Section 4.3.1, each sequential snapshot separated by a time interval Δt. Because of the periodic nature of the flow being compartmentalized, the last and first snapshots are also sequential. This forms a closed loop of CM, as illustrated in Figure 4.21.



(a) (b)

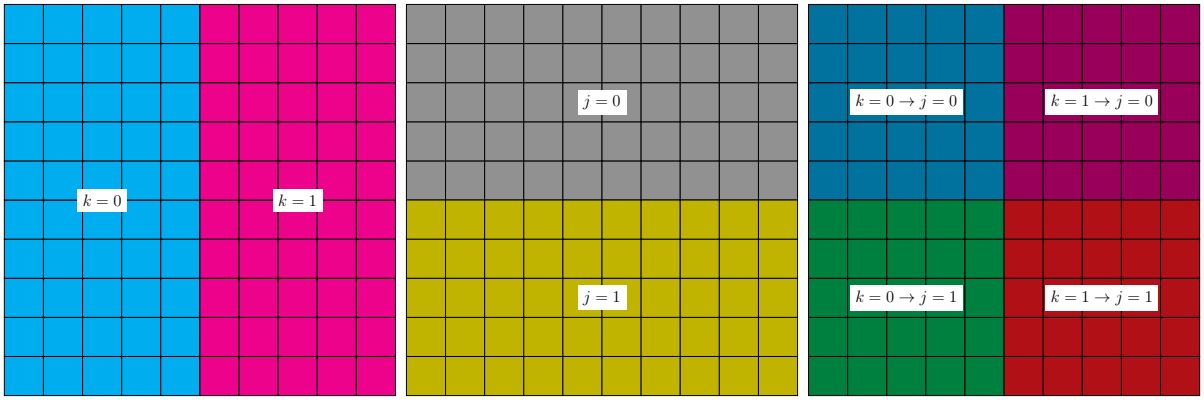
Figure 4.21: Illustration of how a sampled periodic signal is used to created a period set of CMs. **(a)** Periodic signal with snapshots labelled and coloured according to their corresponding CM. **(b)** The resulting compartments and their projection connections to form a closed loop.

A CM is constructed for each snapshot following the method described in Section 4.2. For each CM, DoF-to-element and element-to-DoF mappings are generated according to Section 4.2.5. These mappings are used to create transition matrices, $T^{i \to i+1}$, which project the state of CM i onto CM $i+1$. The use of a transition matrix allows for removing the

restrictive requirement that each CM have the same compartmentalization. This allows the method to be extended to time-varying flows where the topology also changes, as demonstrated in Figure 4.18. This transition matrix has dimensions $n_{\text{states,i+1}} \times n_{\text{states,i}}$. Each row j contains weights for computing the concentration of DoF j in CM $i+1$, and each column k represents the weight corresponding to DoF k from CM i. Each entry is calculated as:

$$T_{jk}^{i \to i+1} = \frac{V_{k \to j}}{V_j} \tag{4.21}$$

Where V_j is the volume of DoF j from CM $i+1$ and $V_{k \to j}$ is the volume of DoF k from CM i that ends up in DoF j. This formulation preserves mass conservation during projection by effectively mapping mass onto mesh elements and back onto DoFs. Figure 4.22 illustrates this process for 4 overlapping DoFs.



(a) (b) (c)

Figure 4.22: Illustration of the projection matrix's role in moving the state between two different CM. **(a)** Mesh elements labelled by DoF k from CM i. **(b)** The same elements labelled by DoF j from CM $i+1$. **(c)** The same elements labelled by source DoF in CM i and destination DoF in CM $i+1$.

However, as noted in Section 4.2.5, situations arise where DoFs become sufficiently clustered that there are not enough mesh elements for all of them. To maintain mass conservation, these cases require special handling. Such unmapped DoFs are projected onto the same elements as their nearest neighbour within the same PFR. In cases where an entire PFR or CSTR lacks mapping, the mapping of the nearest DoF is utilized. The

projection weights are adjusted to account for multiple DoFs corresponding to individual mesh elements, thereby preserving mass.

Two further adjustments are needed. The first, and most critical, involves ensuring satisfaction of the algebraic equations governing inlet DoFs for each PFR:

$$C_{\text{inlet}}^{\text{target}} = \frac{\sum_{k}^{inlets} Q_k C_k(V_k, t)}{\sum_{k}^{inlets} Q_k}$$

For all inlet DoFs, the deviation between their actual values and those required by the governing equation is determined. Two similar correction approaches are then implemented, depending on the inlet DoF connections. For DoFs connected to other PFRs, adjustments are made to the inlet DoF, downstream DoFs, and connected DoFs, while conserving mass to satisfy the algebraic constraint. For DoFs connected to domain boundary conditions, only downstream DoFs are modified, as the inlet DoF values are fixed by boundary conditions and no upstream DoFs exist.

Finally, concentration values are constrained within minimum and maximum bounds. This constraint is applied both independently and as part of the previous correction steps. The objective is to ensure that all non-boundary condition DoFs in CM $i+1$ remain within the minimum and maximum values of CM i. Boundary condition DoFs are exempt from these constraints as their values are governed by boundary condition equations. This approach, analogous to slope limiters, helps maintain physical validity of the results.

The simulation process of this methods has the following 4 steps:

1. Initial conditions are applied to the first CM, with i set to 1 and $t = t_{\text{start}}$.
2. CM i is simulated for duration Δt using an IVP solver, advancing t by Δt.
3. States from CM i are projected onto CM $i+1$, and i is incremented, with i cycling from $N+1$ back to 1.
4. If $t = t_{\text{end}}$, the simulation terminates; otherwise, go to step 2.

Chapter 5

Application to Representative 2D Domain

This section[1] validates the proposed method, as described in Section 4.1, on a representative 2D geometry. This particular geometry was chosen because it features many of the important flow topologies found in 3D systems, but is more interpretable, easier to visualize, and faster to simulate.

5.1 CFD Simulation

In order to evaluate the performance of the proposed compartmental modelling method, it was compared to a 2D CFD simulation. The flow geometry used consists of a channel with an additional re-circulatory region close of the inlet, as shown in Figure 5.1, whose presence induces several different flow regions including: unidirectional, recirculatory, and transitional. This geometry was selected because the resulting flow profile includes the important 2D flow topologies, unidirectional flow, vortex, and saddle points, seen in typical engineering applications, e.g. stirred tanks, packed bed reactors, etc. The use of a simple geometry for this comparative study was also ideal for interpretation, analysis, and visualization.

[1] Adapted from [81].

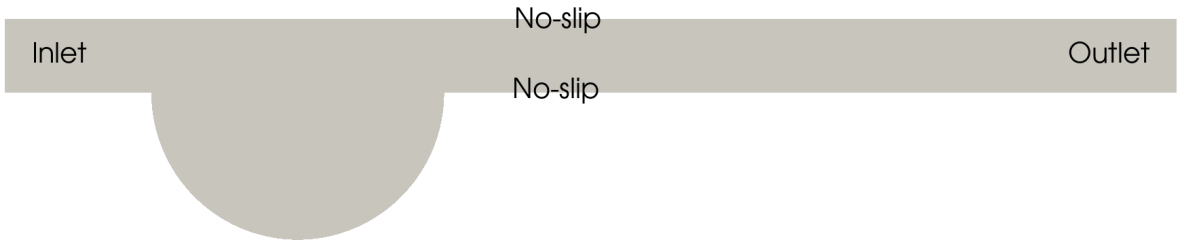


Figure 5.1: Schematic of the simulation domain with boundary conditions with the inlet, outlet, and no-slip boundary conditions labelled. The domain has a height of 1, the inlet length of 2, the outlet length is 10 and the radius of the semi-circle is 2 where all units are non-dimensional.

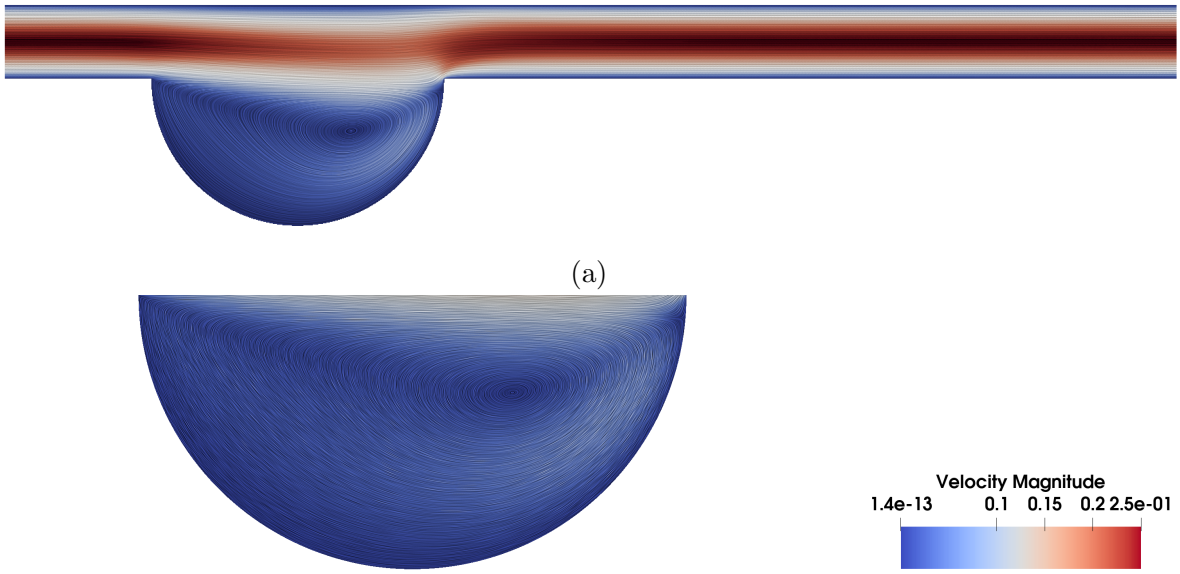
Simulations boundary conditions include no-slip walls and a fixed parabolic inlet velocity profile, $V_x = 4y(1-y)$, which is compatible with no-slip conditions on adjacent boundaries and zero normal-stress on the outlet. CFD simulations were performed using the open-source finite element-based multiphysics package `OpenCMP` [82]. The hydrodynamics of the system were modelled using INS [85, Chapter 19] with Reynolds number of $Re = 125$ and neglecting gravity. Under these conditions, a steady-state velocity profile exists, and was solved for directly.

That steady-state velocity profile was then used as a constant velocity vector field in the convection-diffusion equation for tracer simulations. This was done to reduce the computational time of the CFD simulations for performance comparisons to CM simulations. The unit step change in tracer concentration at the inlet was modelled using a cosine ramp function rather than a perfect unit step in order to improve stability. That unit step was approximated using the following inlet boundary condition:

$$C_{inlet}(t) = \frac{C_0}{2}\left(1 - \cos\left(\pi \frac{t}{t_{ramp}}\right)\right) \tag{5.1}$$

where C_0 is the dimensionless inlet concentration, in this case 1, and t_{ramp} is the time interval over which the concentration goes from 0 to C_0, for these simulations, both CFD and CM, it was set to 1.

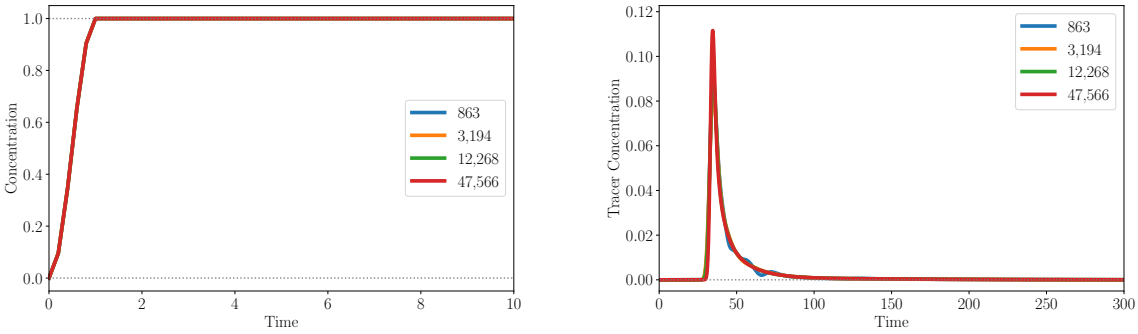
The steady-state velocity profile is shown in Figure 5.2 which was found from simulations using a range of mesh of increasing densities, ranging from 863 elements, coarse, to 47,566 elements, fine.



(a)

(b)

Figure 5.2: Line integral convolution (LIC) visualization of the computed velocity field coloured by velocity magnitude. **(a)** The full domain. **(b)** The recirculatory region.



(a)

(b)

Figure 5.3: Effect of mesh refinement levels on CFD simulation results. **(a)** Inlet tracer concentration as a function of time. **(b)** CFD RTD.

Tracer simulations were performed to determine the reference RTD and CD functions, shown in Figure 5.3(b), using the CFD velocity profiles, as described in Section 4.1. The

velocity field results for the range of mesh densities were used to both determine the effect of mesh density on the accuracy of resulting CM models and, subsequently, to find a suitable trade-off between accuracy and computational complexity. Figure 5.3(a) shows the cosine ramp inlet boundary condition used and Figure 5.3(b) shows the computed RTDs for the range of meshes densities. While not discernible from the plots, the main difference between RTDs predicted by CFD simulations using the range of meshes is in the height and narrowness of the RTD peak. Table 5.1 shows the computed mean residence time, variance, and skewness of the RTDs for the range of mesh densities, which gives more quantitative understanding of the their similarity, with the variance and skewness being significantly different only between the coarsest mesh and the others.

Table 5.1: Summary of the mean residence time, variance, and skewness of the RTD from the CFD simulation across the range of tested mesh densities.

Mesh density	t_m	σ^2	s^3
863	47.86	863.51	651.19
3,194	48.31	855.04	631.12
12,268	48.17	843.52	631.53
47,566	48.15	843.61	635.08

5.2 Compartmentalization

The flow alignment-based compartmentalization was applied to the CFD-computed velocity field in Figure 5.2(a) which resulted in the compartments shown in Figure 5.4(a). Regions of the domain adjacent to the no-slip boundaries, yellow, were excluded from the compartmentalization since the velocity direction becomes degenerate as the velocity magnitude approaches zero, see Section 4.1. Qualitatively, the compartmentalization correctly identified high aspect ratio compartments in the constant cross-section regions, i.e. before and after the re-circulatory region. The transitional regions in the neighbourhood of the recirculatory region were compartmentalized with intermediate aspect ratio compartments which are aligned well with the expected two primary directions of flow: (i) flow bypassing the recirculatory region and (ii) flow entering the recirculatory region.

Finally, the compartmentalization of the recirculatory region, shown enlarged in Figure 5.4(b), was also qualitatively as expected; based on the unidirectional flow assumption, the compartments appear as circular sectors, with increasing deviation from sector-like shape as they get closer to the primary chord. Another intuitive observation was that

the sector-like compartments intersect at a point within the semi-circle which was shifted downstream. That point, as can be seen from Figure 5.2(b), represents the centre of the vortex which forms in the semi-circle.

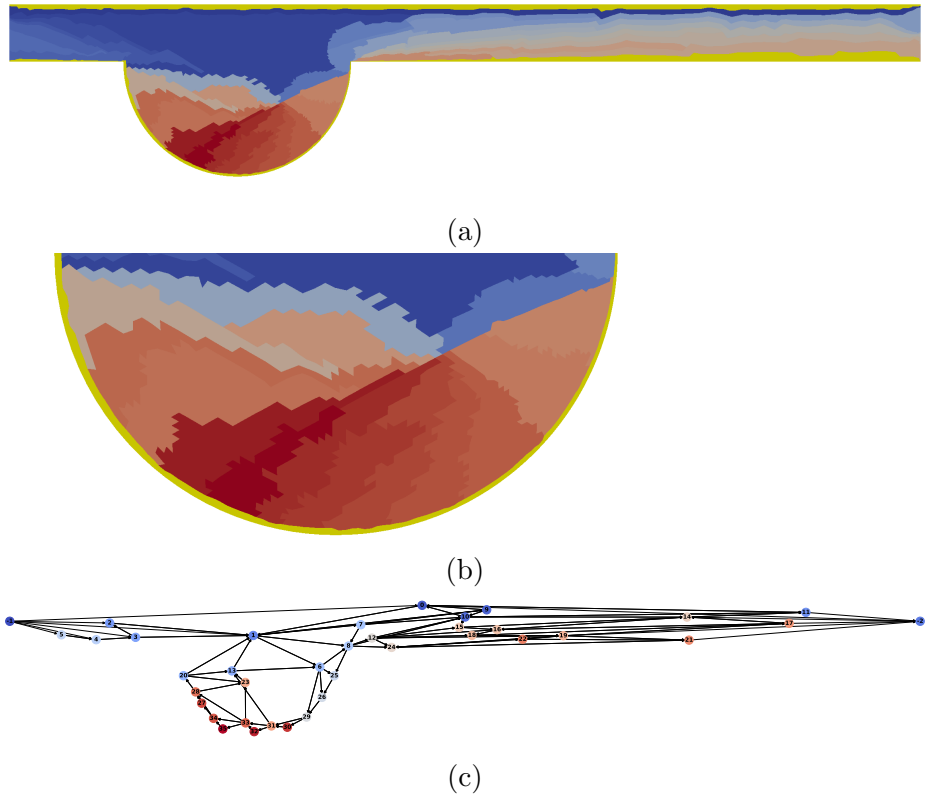


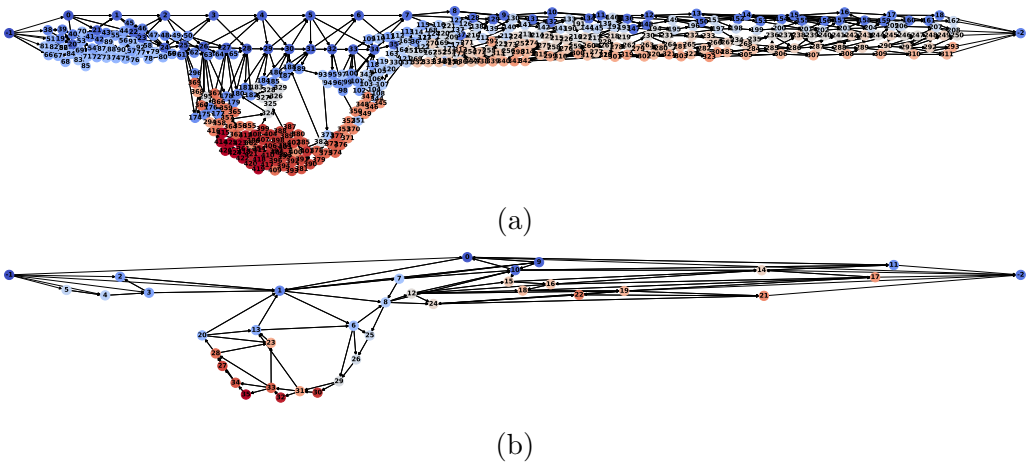
Figure 5.4: Visualizations of the compartmentalization of **(a)** the full reference domain, **(b)** the recirculatory region, and **(c)** the directed spatial graph of centroids for each compartment, nodes, and connections between them, edges. Yellow represents mesh elements not compartmentalized, while the remaining colours, red to blue, signify different compartments. Note that the spatial position of the nodes have been modified slightly to decrease overlap and increase readability.

Figure 5.4(c) shows the directed spatial graph of the compartmentalization with nodes corresponding to compartment centroids and edges corresponding to connections between them. The network appears qualitatively correct, based on the CFD-predicted flow profile; in unidirectional flow regions the nodes are located approximately along the direction of the flow, likewise in the recirculatory region the node positions and edges are in the

radial direction. The single large compartment above the recirculatory region shown in Figure 5.4 has a combination of edge directions, including horizontal and vertical edges which represent flow coupling the unidirectional and recirculatory regions.

5.3 Compartment Modelling

The compartmentalization computed from above was used to create compartment models using either the zero-dimensional CSTR-based approach or the one-dimensional PFR-based approach as described in Section 4.1. The resulting networks for the two approaches are shown in Figure 5.5. The PFR-based compartmental model, for the same compartmentalization, clearly introduces significantly more degrees of freedom. For example, for the coarsest mesh, the zero-dimensional approach, CSTR-based compartments, produced 30 degrees of freedom, while the one-dimensional approach, PFR-based compartments, produced 406. The spatial placement of the additional degrees of freedom followed from the consideration of variations within compartments and is expected to increase accuracy.



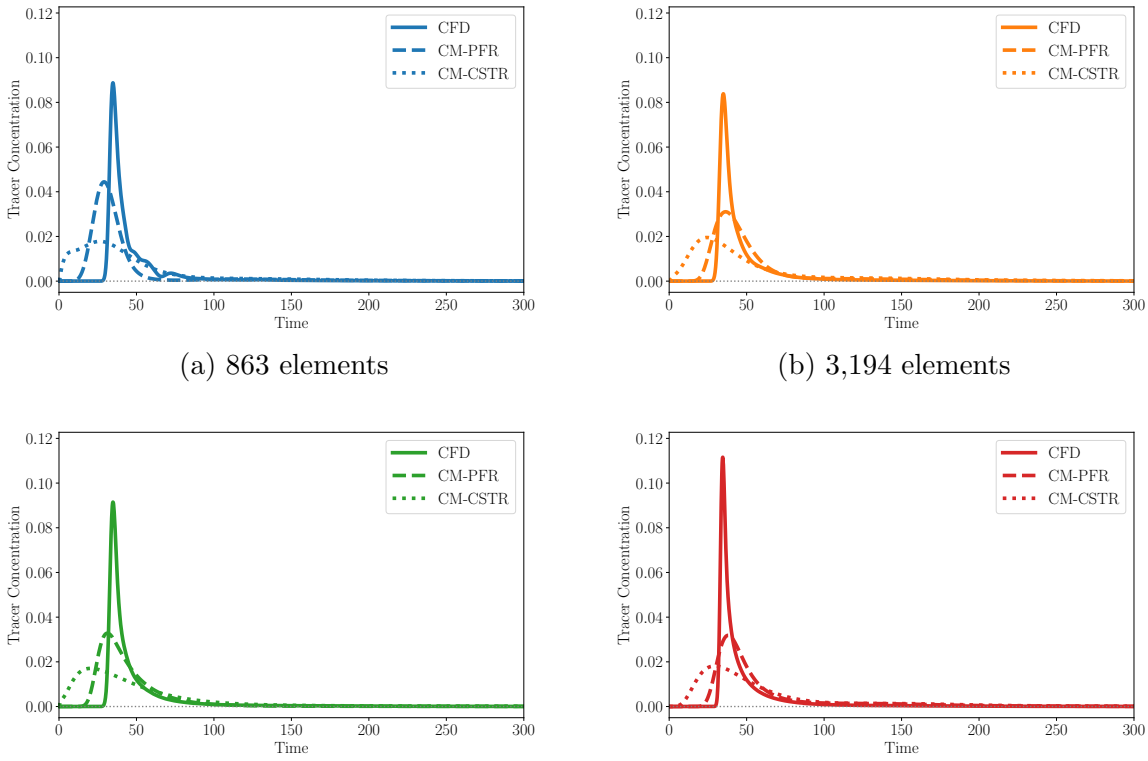
(a)

(b)

Figure 5.5: Spatial graphs showing the different compartment networks with nodes representing each CSTR, or PFR, and edges representing the connections between them. **(a)** The PFR-based CM. **(b)** The CSTR-based CM. Note that the spatial position of the nodes have been modified slightly to decrease overlap and increase readability.

The compartmental models were compared to each other and to the CFD simulation, over the range of mesh densities, using the RTD results from the inert tracer simulation.

Figure 5.6 shows RTD results for both compartmental models at different mesh refinement levels.



(a) 863 elements

(b) 3,194 elements

(c) 12,268 elements

(d) 47,566 elements

Figure 5.6: RTD functions computed using PFR/CSTR-based CM and directly using CFD for the range of mesh refinements.

For all tested meshes, the CSTR-based CM yields RTD functions which are significantly more diffuse than those from the PFR-based CM. This was expected due to the use of the zero-dimensional CSTR approximation compared to one-dimensional PFRs. Analyzing the RTD functions, the first three moments of the RTD functions are shown in Table 5.2, 5.3, and 5.4. There, quantitative measures of the behaviour shown in the graph can be seen. While the mean residence time may appear to be similar between the two CM approaches, the variance and skewness measures quantify what is shown in the figures above, the CSTR-based CM has a wider distribution and different shape distribution than the PFR-based CM and the CFD.

Table 5.2: Mean residence time from the CFD, PFR-CM, and CSTR-CM simulations at the different tested mesh densities.

	t_m		
Mesh density	CFD	CM-PFR	CM-CSTR
863	47.86	42.49	47.57
3,194	48.31	55.23	54.66
12,268	48.17	45.60	46.78
47,566	48.15	58.73	60.46

Table 5.3: Variance of RTD from the CFD, PFR-CM, and CSTR-CM simulations at the different tested mesh densities.

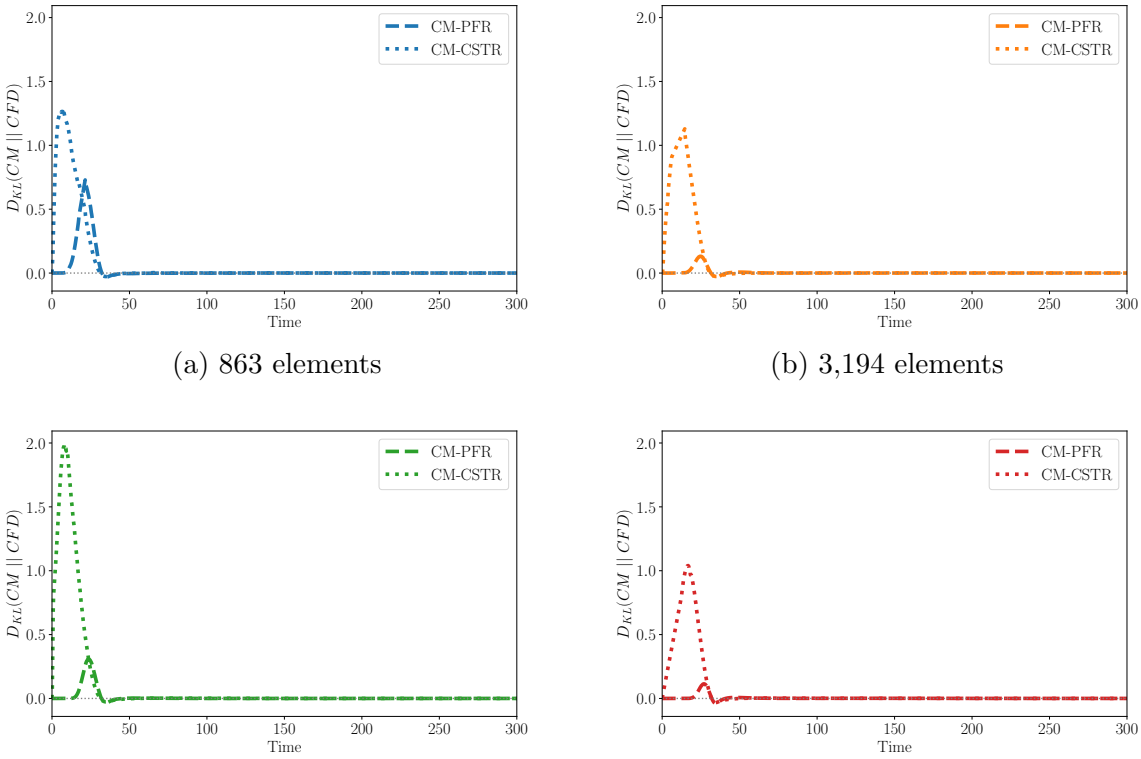
	σ^2		
Mesh density	CFD	CM-PFR	CM-CSTR
863	863.51	1,509.69	2,165.05
3,194	855.04	1,584.84	2,633.57
12,268	843.52	824.72	1,720.58
47,566	843.61	1,725.21	2,479.00

Table 5.4: Skew of RTD from the CFD, PFR-CM, and CSTR-CM simulations at the different tested mesh densities.

	s^3		
Mesh Density	CFD	CM-PFR	CM-CSTR
863	651.19	839.17	740.76
3,194	631.12	716.34	766.15
12,268	631.53	622.74	667.86
47,566	635.08	707.83	702.23

Furthermore, because of the zero-dimensional approximation, concentration immediately propagates from the inlet of the domain to the outlet, thus inaccurately capturing the time delay of the system. The effects of this can be seen by looking at the RTD curve between $t = 0$ and when the CFD peak starts; the higher the RTD curve for the CM model in that region, the worse it captures to true time delay of the system. It can be seen that the proposed PFR-based CM method performs significantly better than the CSTR-based

CM in this regards. This behaviour is also captured by the KL divergence in Figure 5.7, which show that the KL divergence between the CSTR-CM and the CFD is significantly larger than between the PFR-CM and the CFD.



(a) 863 elements

(b) 3,194 elements

(c) 12,268 elements

(d) 47,566 elements

Figure 5.7: KL divergence between RTDs computed using the CSTR/PFR-based CMs and directly through CFD for the range of mesh refinements.

It is important to note that the results of the 12,268 element mesh appear to differ significantly from those of the mesh directly finer and coarser than it, this is true for both the PFR-based and CSTR-based CM results. The CFD simulation, CM creation, and CM simulation were all run several times but the results did not change by an appreciable amount. A reason for its significantly different performance could not be confirmed, though it is hypothesized that it is caused by a combination of two factors. The first is the long-tail nature of the RTD. The mean residence time appears to be in agreement with the results shown in Figure 5.6, but the variance and skew do not; specifically, the variance for this

mesh density does not appear to be *half* of that of the others. The other hypothesized factor is a dependence on the specific orientation of the mesh elements at this mesh resolution due to the element-based method for compartmentalization.

Since computational cost was a key driver for developing the CM, the computational costs of the different methods were calculated and summarized in Table 5.5. The times for the CM approaches were separated into an offline portion, the one-time cost of creating the CM, and the online portion, the cost of solving the model for a set of operating conditions. The computation times shown in the table are in units of cpu–s, which accounts for the multithreaded nature of many of the steps. The time required to solve the hydrodynamics was *not* included in neither the CFD times nor the offline CM times as it is common to all approaches. These results show that the PFR-based CM more accurately approximates the CFD behaviour than the CSTR-based CM while still being up to 3 orders of magnitude faster that the CFD.

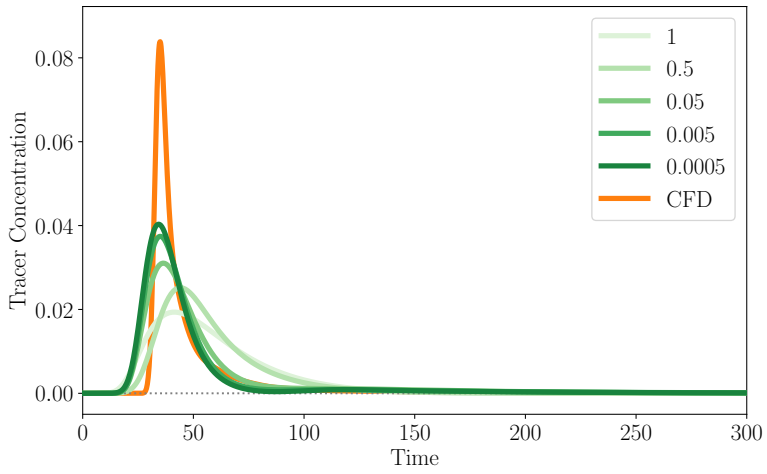
Table 5.5: Computational costs, in cpu-s, of the CFD and CM RTD simulations over the range of tested mesh densities. The CM times are split into Offline, one-time cost for creating the CM, and Online, running the simulation, costs. Neither the CFD nor Offline CM times include the cost of solving for the steady-state velocity profile.

Mesh Density	CFD [cpu-s]	CM-PFR [cpu-s]		CM-CSTR [cpu-s]	
		Offline	Online	Offline	Online
863	6,000	1.86	3.70	1.17	0.05
3,194	6,120	4.31	10.11	4.24	0.06
12,268	64,800	17.72	22.04	17.57	0.10
47,566	1,512,000	80.10	44.98	79.03	0.08

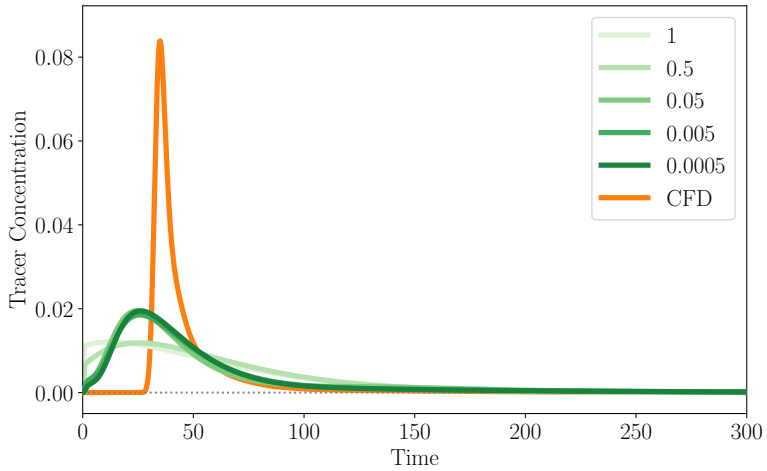
As the number of mesh elements is increased, both the offline and online computational costs increase. While the mesh-element-based nature of the method naturally implies that CM construction costs scale with mesh density, the unexpected increase in online computational cost warrants attention. This increase is particularly noteworthy given that the fundamental system topology remains constant across different mesh densities, suggesting that the resulting CM should remain relatively constant.

Investigation of the results indicates that this behaviour stems from the treatment of subsurface divisions when flux direction changes across surfaces. As mesh facets typically do not align perfectly with streamlines, increased mesh density leads to more frequent flux direction changes. This results in a proliferation of connections and PFRs representing the same spatial region. Thus, while the fundamental structure of the CM does not change

with increasing mesh density the minimum spatial resolution does, which in turn increases the minimum computational cost of the model. This limitation represents a significant drawback of the current method and highlights a critical area for future development.



(a)



(b) CSTR-based CM

Figure 5.8: Effect of minimum compartment size on the RTD predicted from the resulting CM. **(a)** The PFR-based CM proposed in this work. **(b)** Previous state-of-the-art CSTR-based CM.

Finally, while the proposed compartmentalization method is largely automated and based on the CFD simulation results, it requires several user-specified values, such as the minimum compartment size. To understand the effect of the compartment size on accuracy, the compartmental modelling was performed for a range of minimum compartment sizes. The range starts at a size of 1, the whole domain has a size of ≈ 5.5, and decreases several orders of magnitude to a minimum of 0.0005. Figure 5.8 shows the RTD results for both the PFR or CSTR-based CM over the range of compartments sizes specified, and Table 5.6 shows the computed statistics for the RTD.

Table 5.6: Effect of minimum compartment size on the RTD statistics of the PFR-based and CSTR-based CM. The CFD statistics included are those obtained from the same mesh used for creating both types of CMs.

Min Size	t_m		σ^2		s^3	
	PFR	CSTR	PFR	CSTR	PFR	CSTR
CFD	**48.31**		**855.05**		**631.12**	
0.0005	49.02	51.87	1,473.24	2,139.52	830.11	778.71
0.005	51.35	52.01	1,591.93	2,254.56	831.55	803.94
0.05	55.20	54.66	1,580.84	2,633.57	714.21	766.15
0.5	57.37	57.73	543.10	1,962.72	195.57	803.94
1	57.83	56.81	669.02	2,204.00	166.69	778.71

As mentioned above, the CSTR-based CM, due to its zero-dimensional nature, is significantly less accurate at capturing the time delay of the system. Visually this can be seen from Figure 5.8(b) where the peak of the CM RTD curve occurs, for all minimum compartment sizes, before any tracer has exited the reactor in the CFD. The proposed CM is more accurate, as can be seen by observing the same location in Figure 5.8(a). This added accuracy is directly related to the compatibility of the compartmental model and the compartmentalization criteria. The CSTR-based CM has a disconnect between these two steps, as is often found in literature; the compartments are found on the basis of having unidirectional flow, but the use of a zero-dimensional compartment model, a CSTR, neglects the spatial variation and speed at which concentration, information, propagates through the compartment. The PFR-based CM, on the other, hand does not have this disconnect. Even more important is that apparent convergence of the PFR-based CM towards the CFD result, behaviour absent in the CSTR-based CM. By using PFRs-in-series the CM was able to capture the major spatial variations while also better approximating the rate at which mass moves through the system.

5.4 Conclusions

A flow alignment-based compartment modelling method was developed which enables high-order compartmentalization and modelling of generalized flow geometries. This method was applied to a representative 2D geometry and compared to both a traditional CSTR-based CM and to a reference CFD simulation.

The method identified compartments of unidirectional flow by analyzing the local flow alignment and modelled each compartment as a set of PFRs in series. This allowed for the creation of more accurate compartment models by capturing spatial variations that occur along the average flow direction. The use of the presented method was shown to increase the accuracy of the RTD predictions compared to the traditional CSTR-based CM. Specifically, the presented method was found to more accurately capture the time delay of the system. This additional accuracy came at the cost of longer online computation time, as compared to the simpler CSTR-based CM, but still provided a reduction in computation time of about 3 orders of magnitude when compared to CFD simulation. This reduction in computational complexity also enables the use of more standard computational hardware, for example, desktop computing as opposed to high-performance computing required for most CFD simulations.

The local flow alignment-based method presented in this work provides researchers focused on a range of chemical engineering applications with a comprehensive method for creating reduced-order compartmental models of general flow topology. An implementation of this method has been developed for research community as an open-source software package `OpenCCM` [86].

Chapter 6

Application to a 3D Two-Phase Stirred Tank Bioreactor

This chapter presents an application of the compartmental modelling method, described in Section 4.2, to an industrially-relevant 30 m^3 stirred tank bioreactor, a 3D two-phase system. The choice of this bioreactor design and operating conditions was motivated by availability of previously published data from experimental and simulation-based study of this process by several research groups [3, 22, 12].

Past compartmental modelling approaches failed to properly capture the flow behaviour present in such stirred tank reactors, requiring manual intervention, as discussed in Section 3.5. The proposed method was evaluated on its ability to capture this behaviour and was found to do so effectively without manual intervention. The resulting CM is further validated using experimental data [3, 22] including direct time-varying measurements of tracer concentration and mixing times. Finally, the CM was found to enable simulation of complex spatially varying interphase transport.

6.1 Stirred Tank Bioreactor

The 3D multiphase multiphysics process selected to illustrate and validate the compartmental modelling method is the stirred tank bioreactor. Industries widely deploy stirred tank reactors for chemical, pharmaceutical, and food production processes. These reactors play a crucial role in manufacturing vaccines, medications, ethanol, and various value-added

chemicals. The fermentation processes occurring in these bioreactors generally involve aerobic growth, which requires a continuous supply of both substrate and oxygen.

Scaling up these systems presents significant challenges, particularly when mixing and/or oxygen dissolution rates are unable to meet the demand of microbial growth, resulting in significant spatial gradients within the reactor and complex overall growth rates. These spatial gradients result in the formation of zones of nutrient and oxygen depletion, scenarios to be avoided [87, 10]. Accurate prediction of these gradients proves critical for successful scale-up.

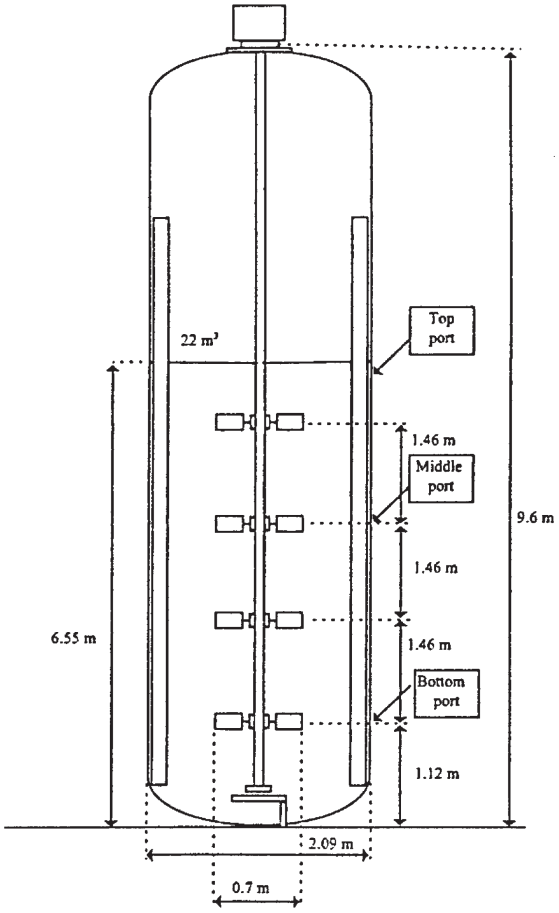


Figure 6.1: Cross-sectional view, with dimensions labelled, of the stirred tank used for the simulations. From [3].

One such scaled-up industrial scale reactor is shown in Figure 6.1. It has a total volume of 30 m^3, with a working volume of 22 m^3 and is known to form significant gradients [10].

This particular reactor was chosen because of the availability of published experimental, CFD, and CM results for it [3, 22, 12], which will act as validation data for the current work. Experimental tracer profiles, and mixing time calculations, in the tank were obtained through a florescent tracer [3]. To limit the number of times the tank had to be drained, multiple tracer runs were performed between draining, each with increasingly higher tracer concentrations. Since the measurement error is proportional to the tracer concentration, this procedure introduced systematic error, which was kept to below 20%.

To obtain the hydrodynamic results needed for compartmental modelling, a multiphase CFD simulation is performed using the two-fluid model `multiphaseEulerFoam` solver in `OpenFOAM` with the Ishii-Zuber correlation for drag [88], the Tomiyama lift model [89], and a constant coefficient virtual mass model. The simulation was run for 70 seconds to allow for the flow field to fully develop, followed by an additional 30 seconds of simulation time to calculate the time-averaged fields. The physical properties use for the simulations are listed in Table 6.1. Air was sparged at a rate of 2.22×10^{-2} kg/s, where the sparger was approximated as a source term rather than a boundary condition in order to reduced simulation complexity. The sparger also acts as a source term for the other governing equations, adding $-1 \times 10^{-6} \frac{\text{kg·m}}{s^2}$ in the z-direction for the conservation of momentum as well as $90 \frac{\text{kg·m}^2}{s^3}$ to the turbulent energy dissipation. The initial condition used for the simulation assumed a stationary liquid and gas phase with the sparger initially not in operation.

Table 6.1: Physical properties used for the two-phase `OpenFOAM` simulation.

Property	**Air**	**Water**	**Units**
Molecular mass	2.89×10^{1}	1.8×10^{1}	g/mol
Dynamic viscosity	1.84×10^{-5}	3.645×10^{-4}	Pa·s
Prandtl number	7×10^{-1}	2.289×10^{0}	–
Surface tension	7×10^{-2}		N/m

The impeller motion, rotating at 13.9487 rad/s, was modelled using the MRF method, despite the considerations in Section 2.5, for two reasons:

1. Since the CM is formulated using the time-averaged flow field, any transient topology features resolved by the SM method would be averaged out. The desire to capture these transient features, discussed in more detail in Section 6.5, motivated the investigations of Chapter 7.

2. A SM-based simulation was attempted but suffered from numerical stability issues. Additionally, SM simulations would have required an impractically long runtime due to the small time steps needed and the limitations of the available software and hardware.

A structured mesh with 1.34-million-elements was used, shown in Figure 6.2. This mesh does not result in a mesh-independent solution, but rather a compromise between speed and accuracy further justified by the inherent approximate nature of the resulting CM. As discussed in Section 4.1, once the important flow geometry has been resolved by a given mesh, improving local error from 10^{-3} to 10^{-6} is not a good use of computational resources when those CFD results are then used to create a CM which has $> 5\%$ error.

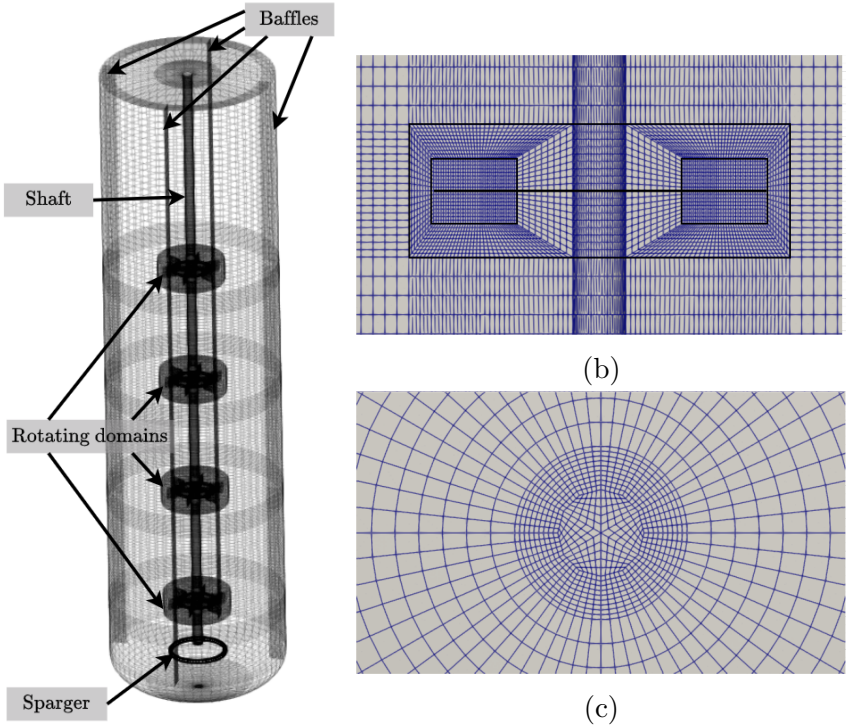
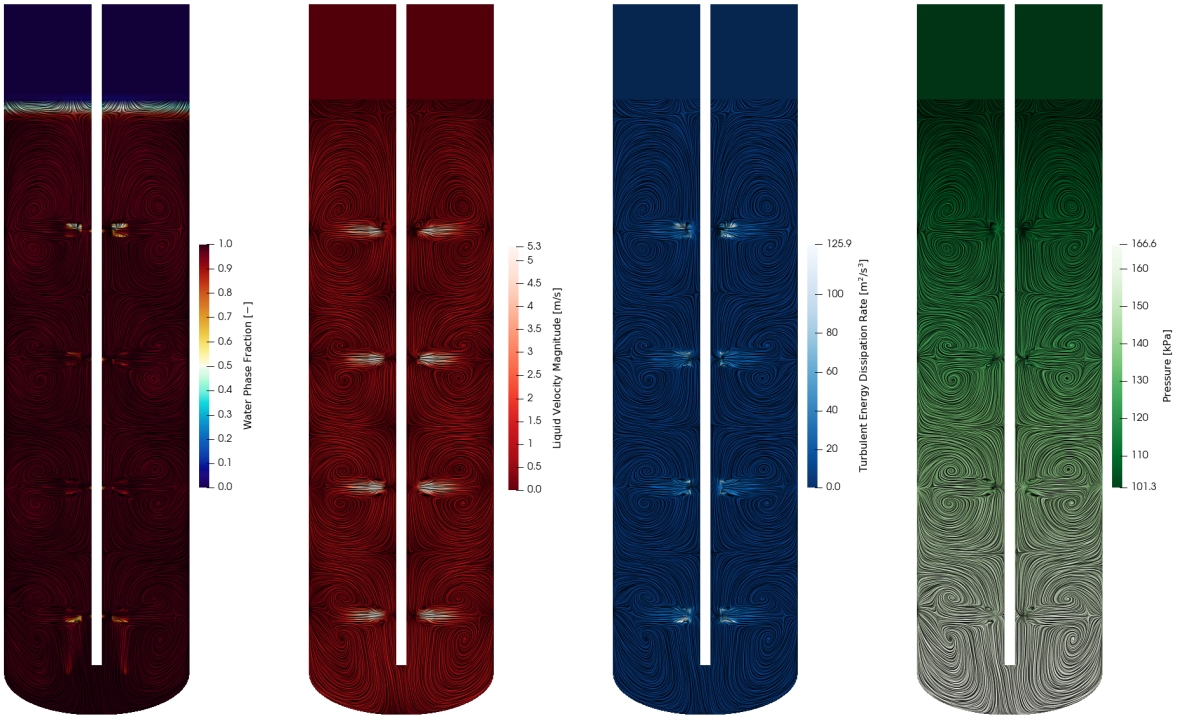


Figure 6.2: Overview of a stirred tank reactor mesh that has been prepared for either an MRF or an SM simulation. **(a)** Overall view of the tank and its mesh. **(b)** Close-up view of the meshing used for the rotating impeller domain. **(c)** View of the bottom of the tank and core of the structured mesh. From [12].

6.2 Compartmentalization

Time-average results of the CFD simulation are shown in Figure 6.3. The 100 s simulation required approximately 280 hours of computation time to complete, corresponding to a simulation speed of about 0.36 s/hr using 32 CPU cores. Although this computational cost is incurred by both the CFD simulation and the CM, the ongoing cost of running simulations using this flow field differs significantly for the two methods.



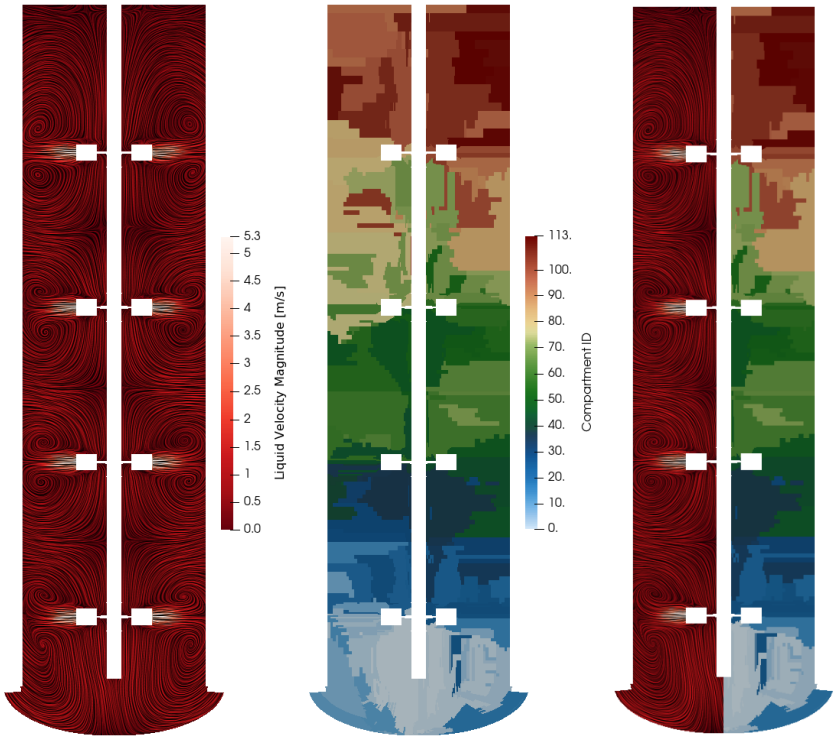
(a) (b) (c) (d)

Figure 6.3: Cross-section of the tank showing time-averaged fields: **(a)** Gas phase volume fraction. **(b)** Liquid velocity with LIC. **(c)** Turbulent energy dissipation rate. **(d)** Pressure.

Based on these time-averaged results, the method generates the CM shown in Figure 6.4, where compartments are differentiated by colour. Since only the liquid phase was compartmentalized, the upper boundary in each cross-section corresponds to the liquid surface rather than the tank ceiling.

A key advantage of the method is its ability to automatically capture behaviour that

past methods failed to. Specifically the horizontal planes with close to zero convective flux across them. These planes were located inline with the impellers and in the inter-impeller region discussed in Section 3.5. The CM method presented in this thesis correctly identifies these surfaces, eliminating the need for manual *ad hoc* intervention, thus providing a more accurate and robust model. These planes can be observed in Figure 6.4, specifically in the side by side provided in (c).



(a) (b) (c)

Figure 6.4: Cross-section of the tank demonstrating the the developed compartmentalization method correctly captures the horizontal no-flux planes. **(a)** Velocity profile with LIC. **(b)** Compartmentalization of the entire liquid phase of the reactor, colouring by compartment ID. **(c)** Side by side of velocity and compartments.

Generating the CM from the velocity field required 202 seconds; a negligible cost when

compared to the 280 hours required to obtain the velocity field. This cost is incurred only once since subsequent simulations in this chapter use the same CM.

6.3 Mass Transfer Simulation

To validate the resulting CM, a inert tracer simulation is run to produce concentration profiles and mixing time predictions. Figure 6.5 illustrates a cross-section of the CM results in time.

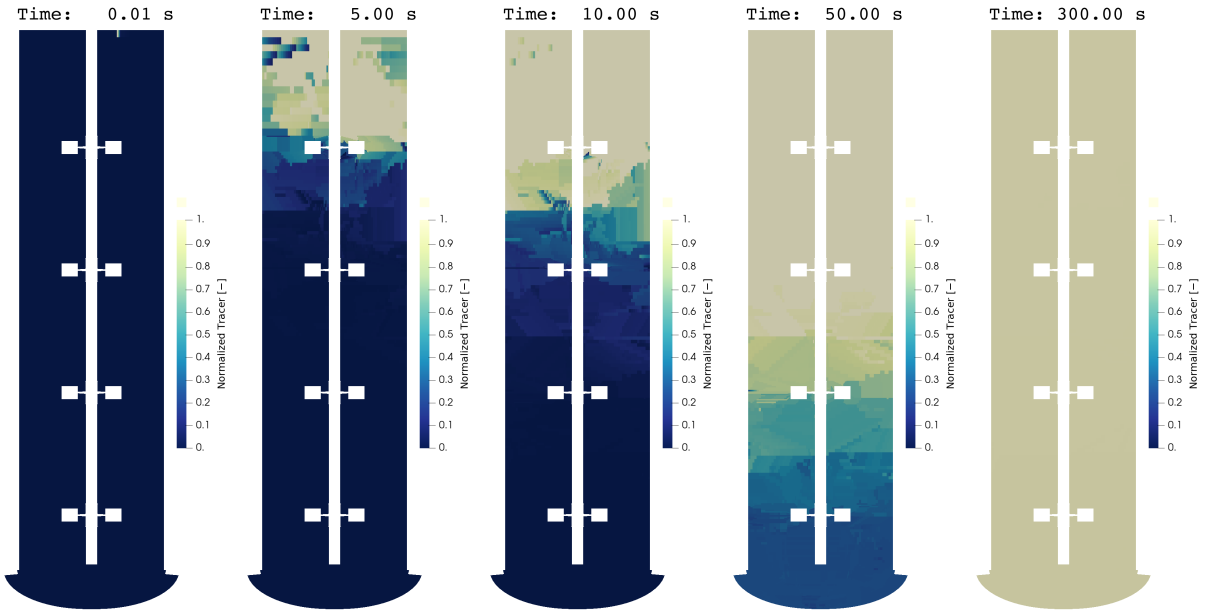


Figure 6.5: Inert tracer concentration along the x-z plane as a function of time. Concentration normalized by the steady-state well-mixed concentration.

The first validation data set consists of concentration measurements from three sensor ports. The sensors were positioned at 3 different heights and thus the probes will be referred to as top, middle, and bottom probes. While most of the geometric details were reported in the original paper, the relative angle between the feed port and the sensor probes was not documented, but is needed for rigorous model validation. To mitigate for

the absence of this information, several angular orientations were simulated with the CM and compared to experimental data in order to determine if the resulting set of simulations could align with the experimental data. Placement of the sensor directly downstream of the feed port was excluded based on two factors:

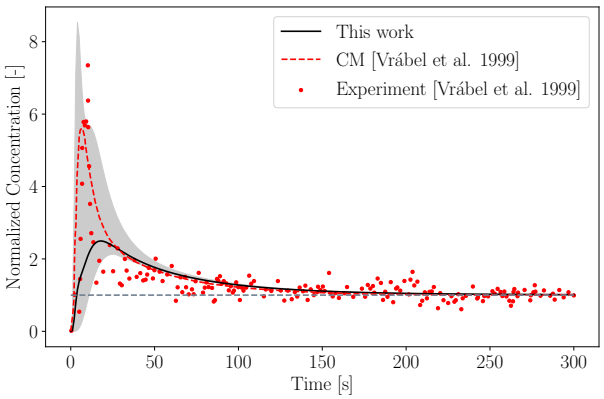
1. The experimental configuration was deemed improbable.
2. The size of the concentration spike and the slight time delay before the occurrence of the spike in concentration at the top probe indicate a degree of spatial separation between the feed port and the sensors.

CM concentrations predictions at the 3 sensors were obtained for each feed-port-sensor angle tested. The set of predictions were used to calculate min, mean, and max values at each point in time. The min, mean, and max values were used instead of selecting a single best-fitting measurement. Using only the best-fitting trajectory would effectively amount to overfitting the model, as there was no *a priori* justification for selecting any particular feed-port-sensor angle over the others. The min and max value were used to create an envelope, show as the grey region in Figure 6.6, while the mean is shown as the black line. The plot also includes experimental data and the CM results reported by Vrabél et al. [3].

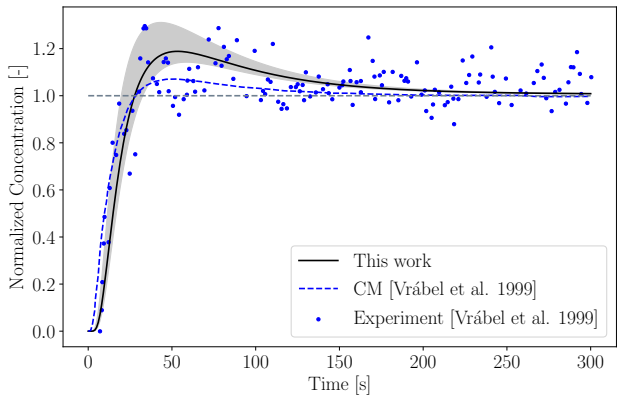
The three main insights from the comparisons in Figure 6.6 are as follows:

1. There is strong agreement between the experimental data and the CM simulation results across all ports. The expected trend is observed including the occurrence of overshoots at the top and middle ports and no overshoot at the bottom port.
2. There is good agreement between the results of the CM developed in this thesis and those of Vrabél et al.'s CM.
3. The proposed method predicts slightly higher overshoots for the top and middle ports, and exhibits a small time delay at the bottom port compared to Vrabél et al.'s approach. Though, it is not clear that either method provides a consistently superior estimation of the results overall.

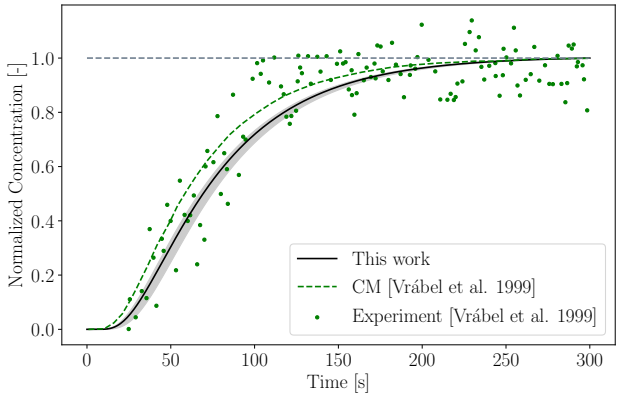
However, it is crucial to emphasize that the CM used Vrabél et al. [3] was manually created specifically for this scenario based on expert knowledge of the flow field. Such a method is very time intensive and inherently not general. In contrast, the CM used in this work is the result of a generalized automatic method being applied to a specific scenario. The proposed method can be, and has been, successfully applied to different geometries without requiring expert knowledge of the systems.



(a) Top port



(b) Middle port



(c) Bottom port

Figure 6.6: Comparison of normalized tracer concentration of the three sensor ports. **(c)** bottom. Values normalized by the steady-state well-mixed concentration. Coloured dots represent experimental data from Vrabél et al. [3]. Coloured dashed lines are CM results from Vrabél et al. [3]. Black line represent average concentration from angle-sweep results of the CM simulation from this work, shaded grey regions represent bounding envelope for the results. Grey dashed line represented steady-state concentration.

Finally, the mixing time was calculated from the CM results and compared to previously reported values. The mixing time was quantified using the CoM described in Section 2.8.1. The comparison with values available in the literature was done based on the time required for the CoM to reach a target value of 0.05. Based on this threshold, a mixing time of 181 seconds is predicted, with the full trajectory of the CoM shown in Figure 6.7. This predicted mixing time agrees well with the experimental measurements reported by Vrabél et al. [3] and falls within the expected prediction range shown by a blue line [3, 90, 22, 12].

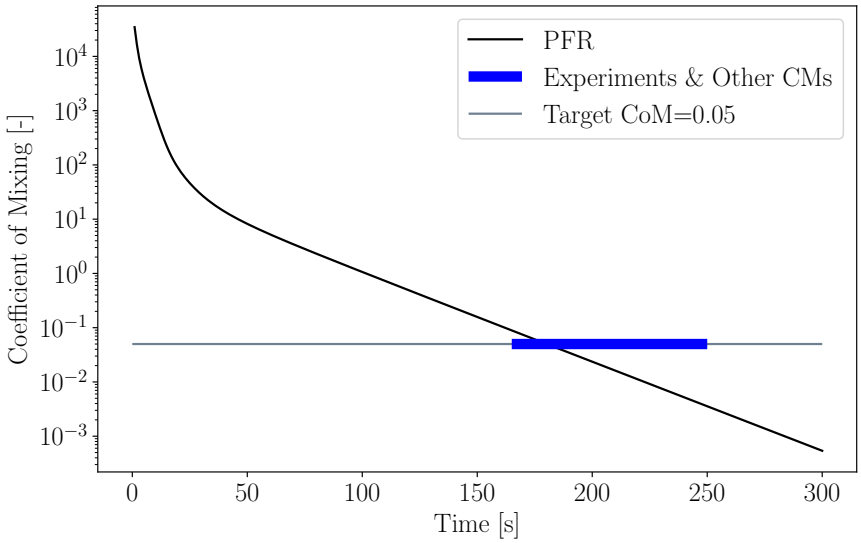


Figure 6.7: Trajectory of CoM as a function of time. Black line is the trajectory of the CM from this work. Grey horizontal line is the target CoM of 0.05. Thick blue line is the range of values from literature [3, 90, 22, 12].

The online portion of the CM simulation required 268 seconds of computation time to simulate 300 seconds, corresponding to a runtime ratio of 1.12. The equivalent CFD simulation, assuming the same 0.36 s/hr, would require approximately 830 hours, four orders of magnitude longer. The computational requirements of the CM are governed by the number of compartments, degrees of freedom per PFR, and the specific scenario under investigation. In this case, the tracer dispersion from a localized injection point necessitated reduced time-steps to maintain numerical accuracy. For comparison, the spatially-varying interphase mass transport presented in the following section, which uses the same CM and numerical tolerances, required significantly less time.

6.4 Interphase Mass Transport Simulation

To evaluate the practical value of the CM, oxygen transport across the gas-liquid interface was simulated. This simulation demonstrated the model's capability to capture complex, spatially varying multiphysics phenomena. The dissolution of oxygen is very important for bioreactor operation, and is often the limiting step for growth, requiring that cell loading and feeding schedules be designed around it [87].

The interphase mass transport across the bubble-liquid interface is modelled by:

$$r = k_L a \left(\mathrm{O}_2^* - \mathrm{O}_2\right) \tag{6.1}$$

where k_L is the mass transfer coefficient, a is the specific surface area of the gas bubbles, O_2 is the local oxygen concentration, and O_2^* the local saturation concentration. The mass transfer coefficient k_L is modelled using Danckwert's surface renewal model [91]:

$$k_L = 0.4 D_L^{1/2} \left(\frac{\varepsilon_L}{\nu_L}\right)^{1/4} \tag{6.2}$$

where D_L is the diffusivity of oxygen dissolved in water, ε_L is the turbulent eddy dissipation rate, and ν_L is the kinematic viscosity of water. Under the conditions used for the CFD simulation, perfectly spherical mono-dispersed air bubbles, the specific surface area of the bubbles is:

$$a = \frac{6\alpha_G}{D} \tag{6.3}$$

where α_G is the gas phase fraction and D is the bubble diameter.

The saturated oxygen concentration is estimated using Henry's law:

$$\mathrm{O}_2^* = \mathrm{P}_{\mathrm{tot}} y_{\mathrm{O}_2} k_{\mathrm{O}_2} \tag{6.4}$$

where $\mathrm{P}_{\mathrm{tot}}$ is the system pressure, y_{O_2} is the oxygen gas fraction, and k_{O_2} is the Henry's law constant.

The values of $\mathrm{P}_{\mathrm{tot}}$, α_G, and ε_L are directly obtained from the CFD simulation with distinct values for each mesh element. The remaining parameters are specified in Table 6.2.

Table 6.2: Physical properties used for the oxygen dissolution simulation.

Parameter	**Value**	**Units**
Diffusivity	2.42×10^{-9}	m^2/s
Bubble diameter	3×10^{-3}	m
Kinetic viscosity	3.645×10^{-7}	m^2/s
Henry's law constant	1.296×10^{-2}	$mmol/m^3/Pa$
Oxygen gas phase fraction	2.09×10^{-1}	–

Figure 6.8 shows the CM predictions of the time evolution of dissolved oxygen concentration throughout the tank.

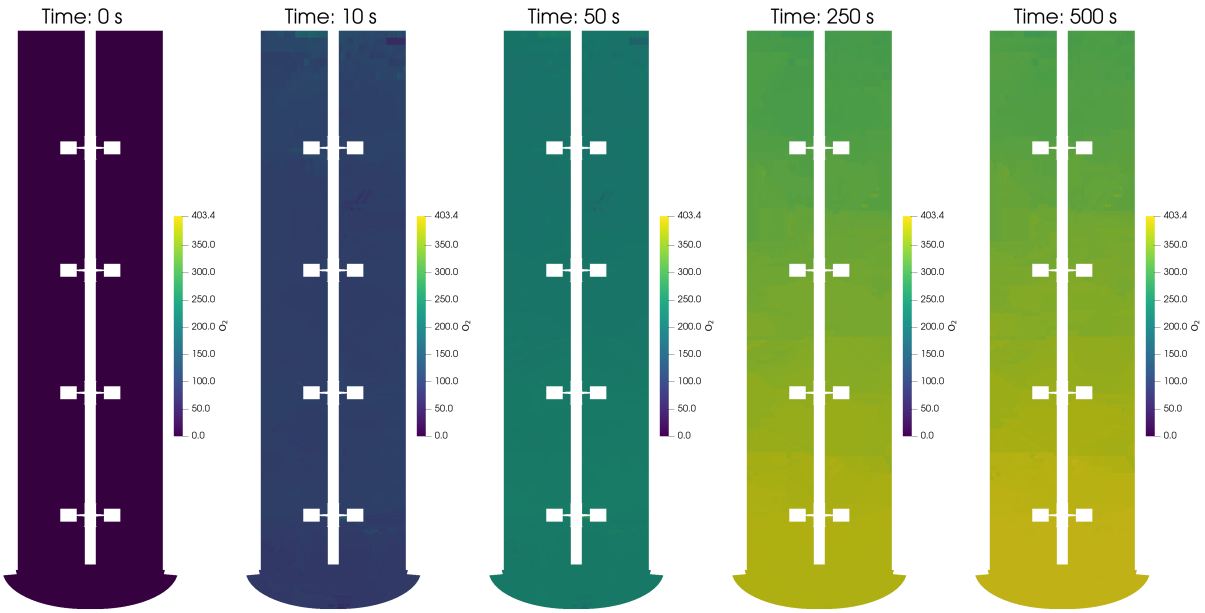
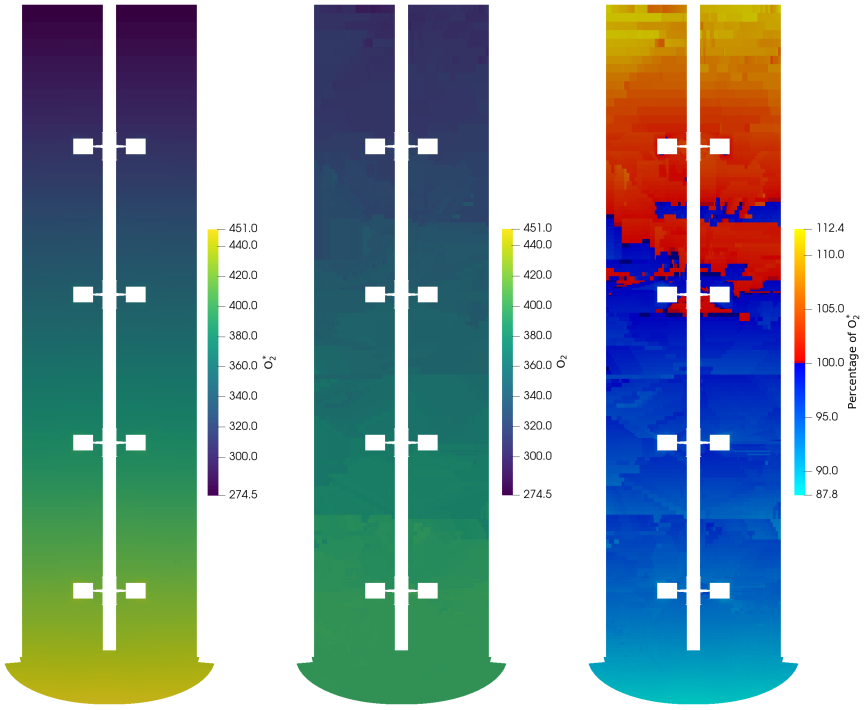


Figure 6.8: Dissolved oxygen concentration along the x-z plane as a function of time.

The steady-state concentration ranges from 294.5 $mmol/m^3$ to 399.2 $mmol/m^3$, a difference of 104.7 $mmol/m^3$ or 35.5% from top to bottom. These values fall within the expected range of O_2^*, ≈451 $mmol/m^3$ at the bottom and ≈274.5 $mmol/m^3$ that at the top.

A concentration gradient was anticipated due to the difference in characteristic time scales: ≈180 s for mixing compared to ≈30 s for the interphase transport based on the

average $k_L a$. While the difference in time scales suggests that a gradient could form, a difference in driving force is also necessary. This difference emerges from the substantial gradient in pressure; a 70% difference between the liquid surface and bottom of the vessel. The resulting concentration profile relative to saturation is illustrated in Figure 6.9.



(a) (b) (c)

Figure 6.9: Steady-state CM simulation results. **(a)** Oxygen saturation concentration. **(b)** Dissolved oxygen concentration. **(c)** Dissolved oxygen concentration as a percentage of oxygen saturation concentration.

The oxygen concentration matches the local saturation concentration near the middle of the tank, the level of the second impeller from the top. Above this level, the reactor operates predominantly above local saturation, reaching up to 13.4% excess, while below this level, concentrations fall below local saturation by up to 14.1%.

This simulation was performed without modelling any microorganism growth or reac-

tions. If those were to be added, it is expected to increase the resulting gradient. The characteristic time of the microorganism's substrate consumption approaches 10^0 seconds, an order of magnitude smaller than interphase oxygen transport and two orders of magnitude smaller than convective transport [22].

In this system, substrate replenishment occurs through a localized feed point and subsequent convective transport throughout the reactor. While oxygen transport also involves convection, it benefits from additional replenishment through interphase mass transport which occurs throughout the bulk of the reactor and at a rate which is an order of magnitude faster than convection. The microbial consumption of both substrate and oxygen occurs at rates two orders of magnitude faster than substrate replenishment and one order of magnitude faster than oxygen replenishment. These significantly different time scales inevitably generate significant concentration gradients within the reactor, substantially impacting process performance.

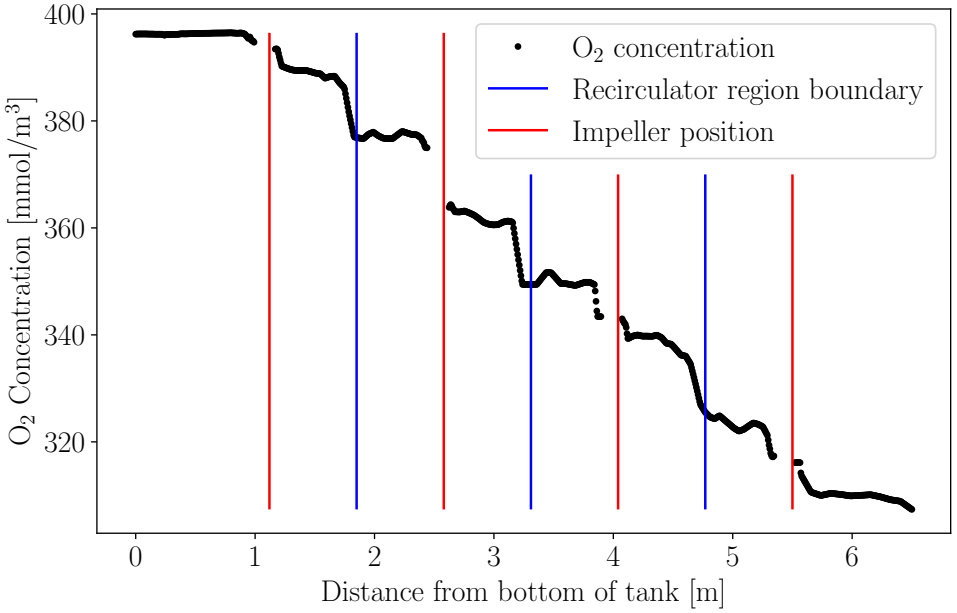


Figure 6.10: Plot of dissolved oxygen concentration as a function of height from the bottom of the reactor along the vertical line at $x = -0.3$ and $y = 0$. Red lines represent the location of impellers, blue lines represent the location of the horizontal planes halfway between impellers. The formation of the characteristic Rushton impeller recirculatory regions are visible as as regions with small gradients separated by abrupt changes in concentration.

Figure 6.9 shows significant concentration differences between adjacent recirculatory zone. These recirculatory regions are characteristic of Rushton turbines and result in horizontal flow occurring at the impeller levels and in the mid-planes between impellers. As discussed in Section 3.5, accurately representing these flow patterns is critical for correct prediction of mixing times and concentration distributions. To better visualize this phenomenon, the steady-state concentration along a vertical line at $x = -0.3$ and $y = 0$ is shown in Figure 6.10. This location was selected specifically to intersect with the impeller positions for clear spatial reference. The boundaries between inter-impeller recirculatory regions are denoted by blue lines, while red lines indicate impeller locations. The recirculatory zones are clearly visible as relatively flat concentration regions separated by sharp step changes. These regions however do not act as CSTRs, while having small concentration gradients within them they are not well mixed. And more importantly, how mass is transported through them, a transient process, and if that is similar to an idealized CSTR, cannot be deduced from this steady-state concentration snapshot.

Computational simulations are essential for evaluating operating conditions to minimize such gradients. By requiring only 58 seconds to simulate 500 seconds of process time, compartmental modelling methods such as the one proposed here provide a computationally efficient way of performing that minimization. Further, this computational efficiency provides sufficient headroom for incorporating more complex reaction kinetics and control systems while maintaining the applicability of these models for real-time applications. In contrast, the equivalent CFD simulation, operating at a ratio of 0.36 s/hr, would have required approximately 1,400 hours, five orders of magnitude longer, without even considering the extra time required for more complex reaction mechanisms. This substantial speed advantage underscores the usefulness of the compartmental modelling method presented in this work.

6.5 Flow Oscillations

One final observation of note is that neither the methods presented by Vrabél et al. [3], Pigou et al. [22], Promma et al. [12], nor the method of this thesis capture the oscillations present in the measurement data, which are most visible in Figure 6.6(a) and (b). Similar oscillatory behaviour has been documented in previous experimental studies, as illustrated in Figure 6.11.

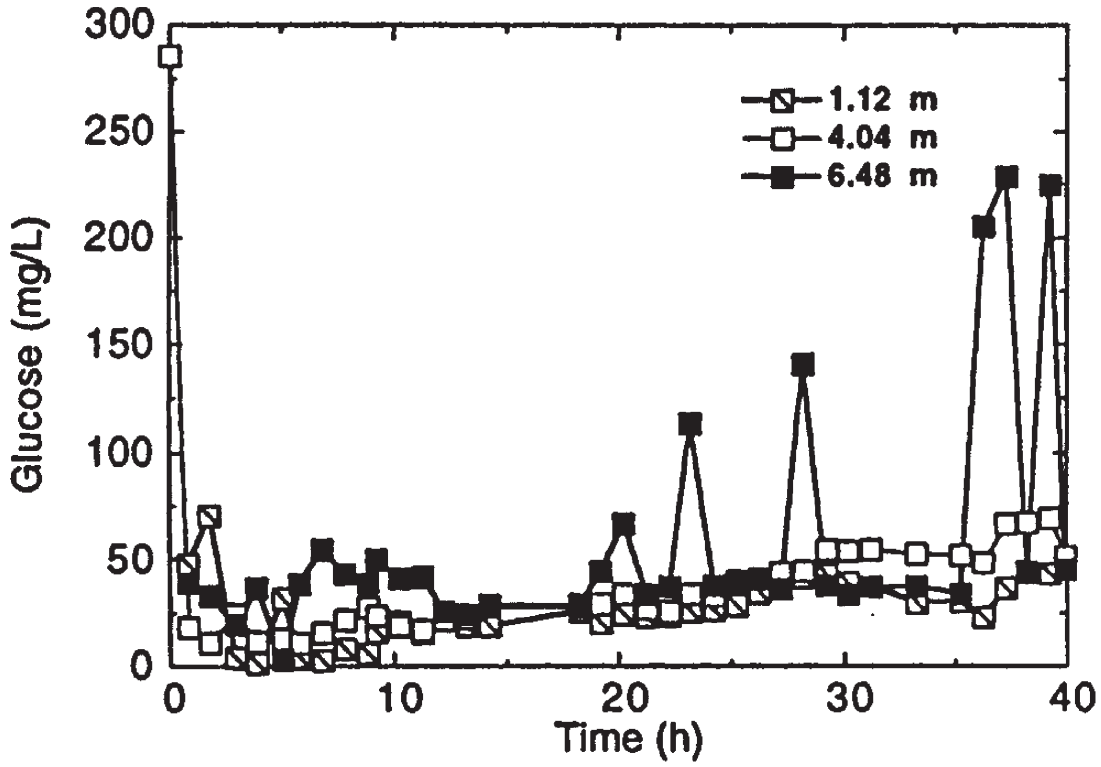


Figure 6.11: Glucose concentrations at three different sample ports shows the presence of large-scale oscillations in the measure values. The black squares, "6.48 m", represent the sensor at the top of the reactor, close to the glucose feed port. From [92].

The observed oscillations are primarily attributed to flow fluctuations [92]. These can be decomposed into two contributions: small-scale turbulent fluctuations and large-scale fluctuations induced by impeller motion. The use of RANS modelling, rather than LES or DNS, eliminates turbulent fluctuations, while the MRF method removes fluctuations associated with impeller movement. Moreover, the current time-averaged flow field method inherently removes any fluctuations present in the CFD simulations. However, for the CM to capture these oscillations the CFD simulations must itself model the oscillations.

The observed oscillations reveal two significant implications for bioreactor performance:

1. Substantial spatial heterogeneities exist within the reactor, even when localized to specific regions, which may not be detected by current sensor configurations. These localized variations significantly influence overall bioreactor productivity; glucose heterogeneities can trigger metabolic pathway shifts in microorganisms. Similar critical

effects are observed around pH control ports, where extreme pH values result in cell death.

2. While the CM accurately predicts the *time-averaged* behaviour, it cannot capture the oscillations observed in the data. The predictions of these oscillations should be based on time-varying CFD results rather than time-averaged results.

This limitation underscores the fundamental relationship between CFD fidelity and compartment model accuracy: the CM is an approximation of a CFD model, thus behaviour that the CM seeks to capture must in the CFD results themselves. Extending the proposed CM methods to such time-varying scenarios is described in Chapter 7. However, to make this approach feasible for large-scale simulations, a different CFD simulation framework, such as one based on the Lattice Boltzmann method [6], is required in order to obtain the higher fidelity hydrodynamics results in a reasonable time.

6.6 Conclusions

The developed method demonstrates significant advances in bioreactor simulation capabilities. This approach automatically models hydrodynamic features that previous compartmental modelling methods were inherently incapable of capturing.

Validation through mixing time and tracer concentration studies shows excellent agreement between the proposed CM predictions and experimental data. This accuracy is achieved without requiring any scenario-specific model customization or manual *ad hoc* intervention, while executing four orders of magnitude faster than CFD simulations. The tracer concentration analysis reveals the need for higher-fidelity CFD simulations to develop time-varying CMs, necessitating alternative CFD approaches, such as the Lattice Boltzmann method [6].

Finally, the CM is used to simulate complex spatially-varying interphase mass transport, enabling predictions of critical gradient formation while maintaining a computational advantage of five orders of magnitude over CFD simulations. These achievements demonstrate the potential of this method as a practical tool for applications that require fast computation, such as design optimization and real-time process monitoring and control.

Chapter 7

Accounting for Time-Varying Flows

CMs commonly rely on a single representative velocity field, typically either steady-state or time-averaged simulation data. However, this approach assumes that a single velocity field adequately characterizes process hydrodynamics for the entire time duration of interest. Analysis of the stirred tank system, presented in the previous chapter, demonstrated that while CM predictions captured general mixing trends they failed to reproduce the oscillatory behaviour observed in experimental measurements, see Section 6.5. This chapter examines a representative 2D case study where the time-averaged velocity field cannot adequately approximate the time-varying velocity field. To deal with this situation, two different compartmental modelling approaches are proposed to capture time-varying hydrodynamics.

Promma et al. [12] demonstrated improved agreement between their stirred tank CM and experimental results by manually introducing oscillations into the flowrates. However, their approach maintained a fixed compartmentalization while flow oscillations were artificially introduced to fit reported experimental data rather than incorporating data from CFD simulations. More recently, Morchain et al. [78] presented a time-varying compartmental model capable of capturing dynamic flow behaviour, though their method relied on a fixed, regular grid of compartments with only the inter-compartment flowrates varying over time. In contrast, this chapter proposes and evaluates methods that account for the time dependent flow topology through dynamic compartmentalization. The following 3 different methods were investigated:

1. Using the time-averaged flow field to produce a single CM, which fails to accurately predict the time-averaged CFD concentration profile.

2. Using an ensemble approach, proposed in Section 4.3.2, which captures the time-averaged concentration profile but fails to capture the full range of concentration oscillations.

3. Using a novel projection-based dynamic CM approach, outlined in Section 4.3.3, which resolves the dynamics of the system but provides inaccurate concentration estimates.

7.1 Chosen Geometry

For this preliminary study, the system chosen to evaluate each of the dynamic compartmentalization methods is the classical two-dimensional vortex shedding resulting from flow past an immersed cylinder in the laminar flow regime. A two-dimensional system was chosen because it is easier to simulate and its results are easier to visualize and interpret. This particular geometry, the Schäfer-Turek benchmark [93], was chosen because it produces a periodic steady-state without the need for moving boundaries or turbulent flow.

Figure 7.1 illustrates the computational domain. The boundary conditions and operating conditions were taken to be the same as those used by Schäfer et al [93]. The walls have no-slip conditions, while the outlet uses a zero gradient condition. The inlet velocity profile is defined by a parabolic function $V_x = 6y(0.41 - y)/0.41^2$, yielding a maximum velocity of 1.5 m/s. The immersed cylinder has a diameter of 0.1 m and the fluid has a kinematic viscosity of 1×10^{-3} m^2/s, resulting in a Reynolds number of 100.

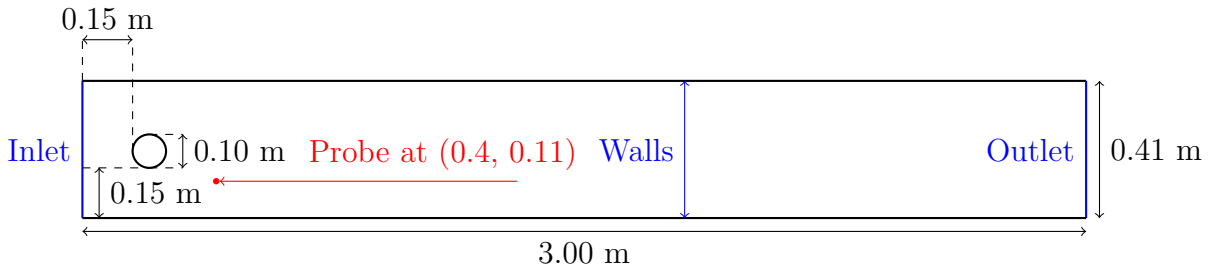
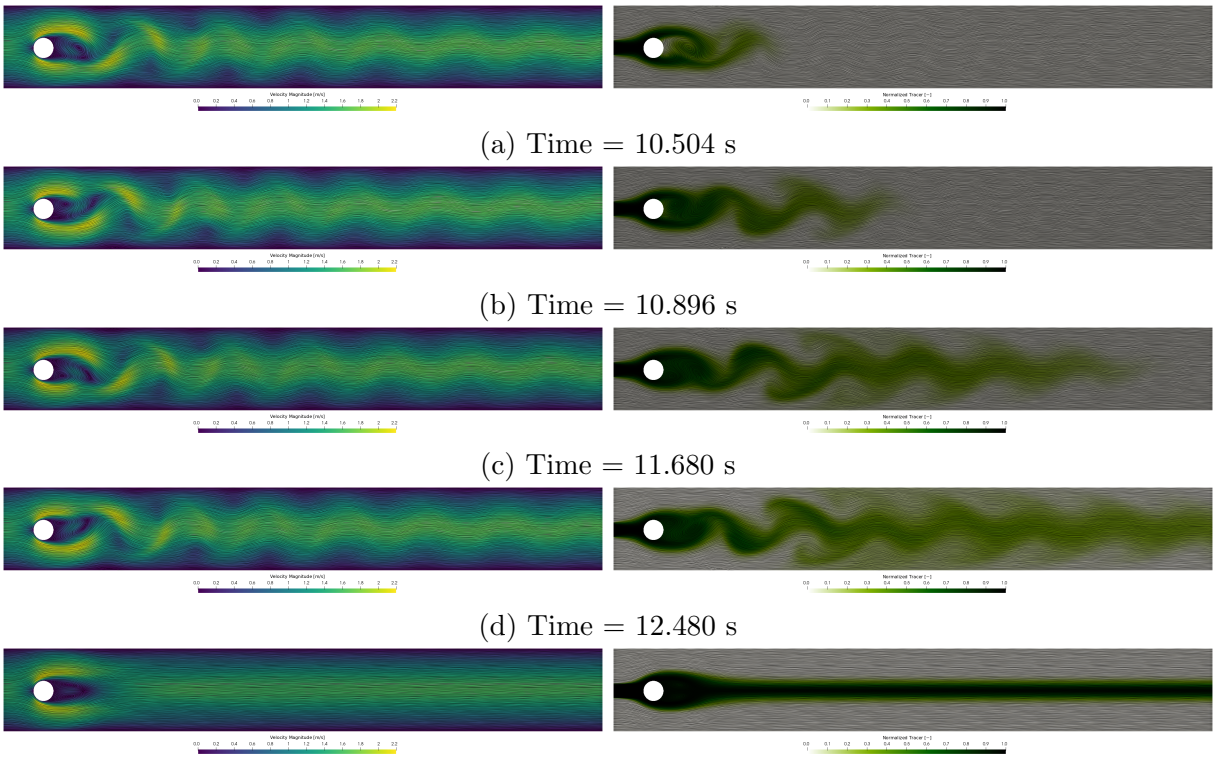


Figure 7.1: Labelled 2D vortex shedding domain with boundary names, domain dimensions, and probe location labelled.

7.2 CFD Simulation

The hydrodynamic simulations performed using `OpenFOAM` consisted of a 10-second initialization period to achieve periodic steady-state followed by a 10-second time-averaging period. Two inert tracer simulations were performed using those results. The first used the time-varying flow beginning at $t = 10$ s to determine the time-varying tracer profile, the second used the time-averaged flow field to compute the steady-state tracer profile. Figure 7.2 presents the results of all 4 simulations.



(a) Time = 10.504 s

(b) Time = 10.896 s

(c) Time = 11.680 s

(d) Time = 12.480 s

(e) Time-Average Flow

Figure 7.2: Snapshots of the time-varying flow and the time-averaged flow CFD simulations. **Left:** Velocity magnitude with LICs. **Right:** Tracer concentration.

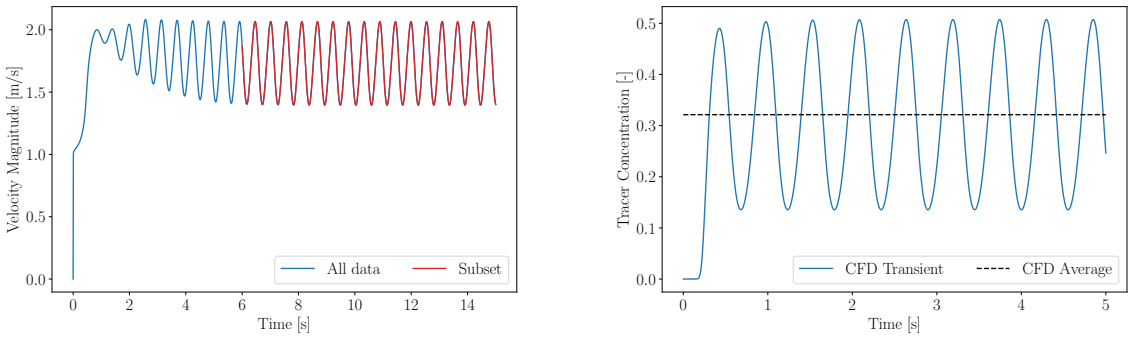
Significantly, the time-averaging removes the vertical component of the velocity field producing predominantly horizontal flow, as shown by the LICs in Figure 7.2(e). Performing a tracer simulation using the time-averaged flow leads to an inaccurate estimate of the time-averaged tracer concentration, as can be seen in the right side of Figure 7.2(e). The

lack of a vertical velocity component causes the tracer to be confined to the centerline of the domain. This system therefore presents an ideal test case, as the negative effect of the time-averaging flow topology operation on the predicted tracer distribution is evident.

The hydrodynamics required 41 minutes for the first 10 seconds, and an additional 41 minutes to calculate the time-average fields. The time-varying tracer simulation took 27 minutes to simulate 5 seconds. In comparison, calculations using the time-averaged flow field reached steady-state in under a minute. The reduced scale of this system, compared to the stirred tank reactor discussed previously, allowed for desktop computing rather than high-performance resources. This enabled the use of 4 cores, producing CFD simulation speeds that were representative of typical processing times on standard engineering workstations.

7.3 Snapshot Subsampling

Both the ensemble and the dynamic projection approaches require multiple velocity field snapshots. The snapshot selection method, detailed in Section 4.3.1, employs a frequency analysis to identify dominant temporal modes. This analysis determines both the required sampling period and guides the selection of the required number of snapshots. Figure 7.3 shows probe measurements of velocity magnitude and tracer concentration. The velocity magnitude plot shows both the complete simulation duration (blue) and the subset used for frequency analysis (red).



(a) (b)

Figure 7.3: CFD time-varying simulation measurements of oscillations in fields. **(a)** Velocity magnitude over whole simulation. **(b)** Tracer concentration over multiple oscillatory periods. Time in (b) is measured from the injection of the tracer at $t = 10$ s.

The subset of the velocity magnitude measurements outlined above are used in a frequency analysis, following Section 4.3.1. The data was mean-shifted to zero to prevent bias in the DFT calculations. The resulting power spectrum is shown in Figure 7.4.

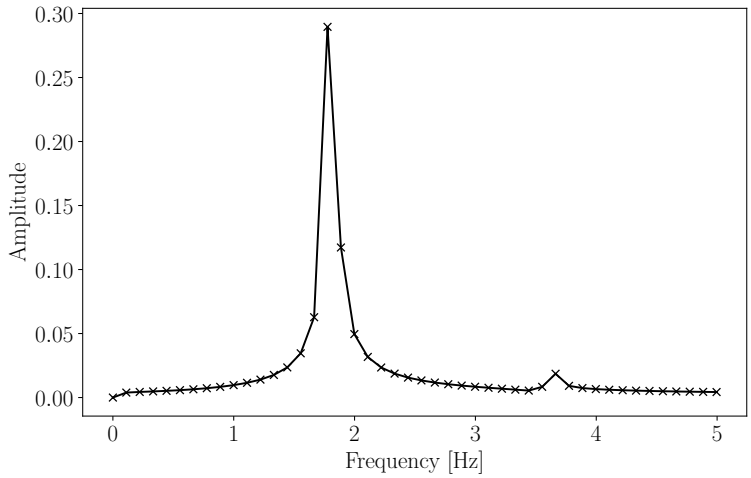


Figure 7.4: Power spectrum on the subset of velocity magnitude measurements.

The frequency analysis reveals a primary peak at 1.776 Hz, period of 0.563 s, and a secondary peak at 3.66 Hz. Given the negligible magnitude of the secondary peak, it was omitted for selection of the snapshots. Within the 0.563 s period, 70 CFD snapshots were available. Based on the discussion in Section 4.3.1, this allows for sampling either every 3^{rd} or 23^{rd} snapshot. The ensemble method does not have the restrictions on sampling, but the same snapshots were used for both methods for simplicity. The analysis proceeded with sampling every 3^{rd} snapshot. Figure 7.5 displays this sampled signal superimposed on the original data.

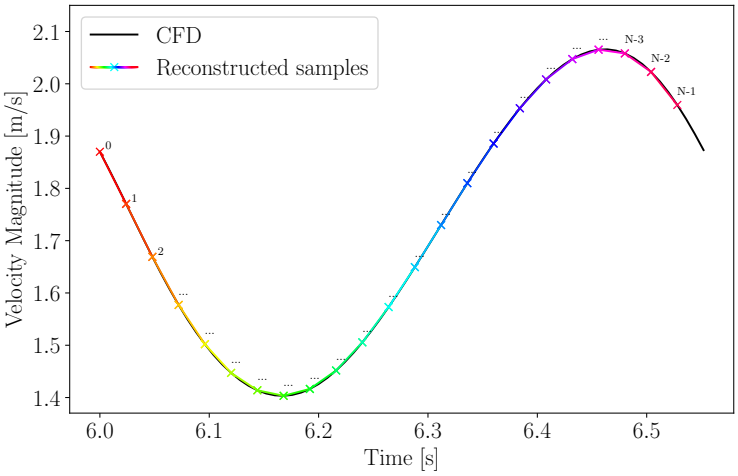


Figure 7.5: Sampled velocity magnitude with chosen snapshots labelled. Cyclic colourmap used to emphasize that the first and last compartments along this period are connected.

7.4 Compartmental Model Simulations

Three distinct approaches for creating CM are evaluated. The first uses the time-averaged hydrodynamic profile, following the method detailed in Section 4.2 and previously applied in Chapter 6. The second, described in Section 4.3.2, creates individual CMs for each snapshot and uses an ensemble approach to predict concentration ranges and mean values. The third, described in Section 4.3.3, generates compartmental models and transition matrices for each snapshot and constructs a time-varying compartmentalization which captures the time-varying nature of the velocity field.

7.4.1 Time-Averaged Compartmental Modelling Approach

While the CM using time-averaged flow was expected to predict the mean concentration values from the transient CFD simulation, without capturing the oscillatory behaviour, the results presented in Figure 7.6 demonstrate that even the predicted means is inaccurate.

The CM predictions, shown by the red line in Figure 7.6, fail to accurately represent the CFD concentration at the probe location. While the model was not expected to capture the oscillatory behaviour, it fails to approximate the mean concentration value which it was

expected to capture. This poor performance was previously indicated in Figure 7.2, where the tracer simulation using time-averaged flow fields demonstrated significant deviation from the coupled flow and tracer CFD results. The computational requirements included 20 seconds for data pre-processing, 7 seconds for CM generation, and 3 seconds for simulation.

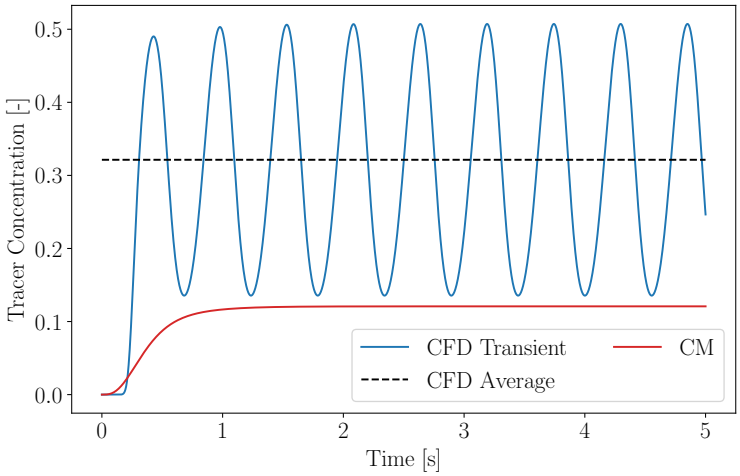


Figure 7.6: Comparison of CFD and time-averaged flow field-based CM tracer simulation results at the probe location.

The observed discrepancy likely stems from the loss of critical time-varying topological features during time-averaging, introducing significant model error. The time-averaging process eliminates the vertical velocity component without adequately accounting for its role in mass transport. This situation parallels Reynolds-averaging of the Navier-Stokes equations, where a $\langle \boldsymbol{u}'\boldsymbol{u}' \rangle$ term emerges requiring closure. However, this case presents a distinct challenge: the flow field oscillations occur over extended time-scales and exhibit large magnitudes relative to the mean flow.

One potential solution for this issue could use the fact that these oscillations are deterministic and can be directly calculated through simulation of the periodic steady-state, enabling accurate determination of both $\boldsymbol{u}'$ and $\langle \boldsymbol{u}'\boldsymbol{u}' \rangle$. These quantities can then be used to inform the calculation of an anisotropic diffusivity for mass transport, similar to the approach used in estimating a turbulent diffusivity for RANS models. Implementation of these findings within the current compartmental modelling framework would necessitate modification of the inter-compartmental mass transfer mechanisms. The existing implementation assumes convection-dominated flow, which accurately describes the system prior

to time-averaging. However, incorporating time-averaging and accounting for oscillations through anisotropic diffusivity would likely invalidate this assumption, requiring revision of the inter-compartmental mass transfer approach.

7.4.2 Ensemble-Based Compartmental Modelling Approach

The single CM approach fails to capture both the time-varying oscillations and, notably, the time-averaged tracer values. To address these limitations, an ensemble method described in Section 4.3.2 was implemented. This approach generates individual CMs for each selected snapshot, with tracer simulations executed in parallel. Results are subsequently combined during post-processing. The underlying hypothesis is that each snapshot represents a distinct flow trajectory. Hence, utilizing an ensemble of representative velocity fields rather than a single time-average field should capture the full range of time-varying flow behaviours. Rather than producing oscillatory predictions, this method aims to generate a the mean concentration as well as an envelope capturing the minimum and maximum values within which the actual values oscillate. Figure 7.7 presents the results of the ensemble CM simulation.

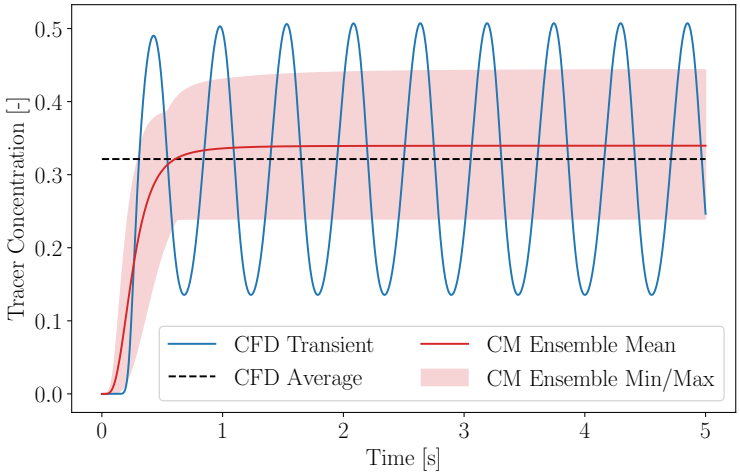


Figure 7.7: Comparison of CFD and ensemble CM tracer results at the probe location.

The ensemble approach accurately approximates the time-averaged concentration value, as shown by the close agreement between the solid red line and dashed black line. However, the red shaded region representing the predicted oscillation range fails to encompass the

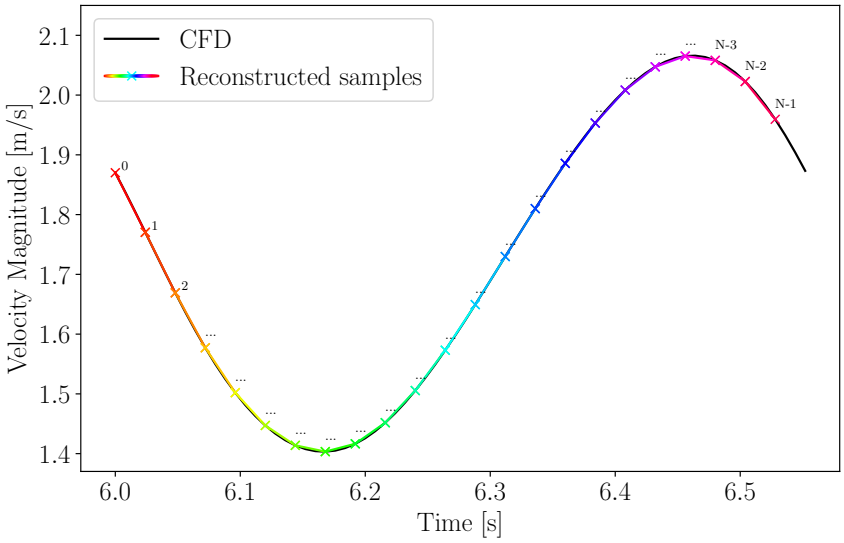
full peak-to-peak variation observed in the CFD results. While this method captures more topological information than the single CM approach, it is unable to account for convective transport between snapshots, providing bands around a mean rather than the full trajectory of the flow. The computational cost is higher than than the previous time-average approach because it solves multiple CMs. Data pre-processing required 50 seconds, CM generation 310 seconds, and the simulations 113 seconds.

It should be noted that while the same snapshots were used for both ensemble and projection methods, the ensemble approach could be run with fewer snapshots while obtaining the same results. While computational efficiency could be improved through optimized snapshot selection, the development of such a selection method remains a topic for future investigation.

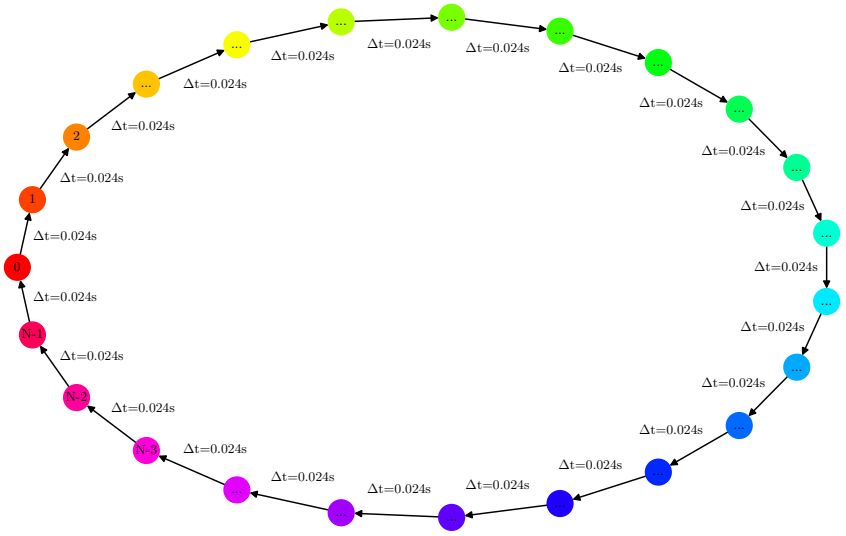
7.4.3 Dynamic-Projection-Based Compartmental Modelling Approach

The third approach in this study consists of using each snapshot to create a CM and a transition matrix. Thus, by treating the problem as a sequence of compartmental models, the intent is to address the limitations observed in both the time-average-based and ensemble-based CM approaches. Unlike the ensemble approach, this approach explicitly accounts for mass transfer between different flow states through the transition matrices projecting between CMs. A full description of this approach is provide in Section 4.3.3.

The snapshots used and the resulting closed loop of CMs are shown in Figure 7.8.



(a)



(b)

Figure 7.8: The snapshots used for the projection compartmentalization. **(a)** Concentration for full length of coupled tracer simulation. **(b)** Closed loop of CM with Δt labelled.

The preliminary results of the projection-based method are presented in Figure 7.9. The predictions exhibit several notable characteristics:

1. Correct prediction of the oscillation period.
2. The presences of a time delay.
3. An over-prediction of peak tracer concentrations.
4. The presence of spurious concentration fluctuations in trough regions.

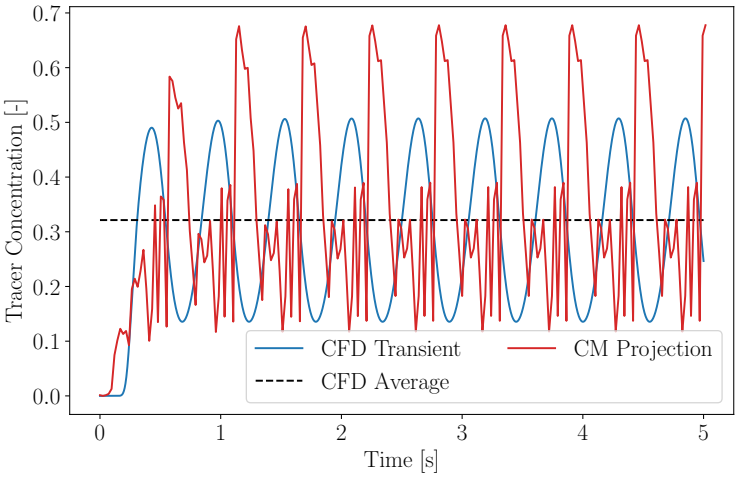


Figure 7.9: Comparison of CFD and dynamic projection CM tracer simulation results at the probe location.

Based on an investigation of the resulting compartments, it is hypothesized that the observed phase shift is caused by excessive merging of compartments in the region upstream of the cylinder. This merging reduced the superficial velocity in those compartment, which in turn increases the time delay of the system. To address the issue, the compartmentalization method would likely need to be expanded to include spatially varying merging criteria so that the upstream and downstream regions are merged differently. The over-prediction in peak concentration and the presence of the fluctuations are hypothesized to be caused by inaccuracies in the projection and/or adjustment steps. While there is a direct mapping between mesh elements and compartments, the overall approximate nature of the CM and

the 1D nature of the PFRs are also expected to introduce inaccuracies when projecting back the solutions onto the original geometric domain.

The method required 45 seconds for transition matrix generation and data preparation, with a total simulation time of 2,500 seconds. This substantial computational cost stems from two primary factors. The first is that due to practical implementation constraints necessitated the use of a fixed-step implicit Euler solver rather than the higher-order adaptive schemes employed in previous chapters. Because of their convection dominated nature, the resulting network of PFRs is a stiff system which require substantially smaller timesteps in order for the first order accurate scheme to converge. Addressing this issue requires changes to the implementation details of the method, but not the method itself; switching to a high-order time-stepping scheme would likely decrease the number of timesteps required over each Δt, and thus the computation time.

The second is that the projection operations impose a maximum timestep size of Δt. To ease this limitation, the CFD output intervals should be chosen so that they result in more favourable subsampling options, thus increasing Δt. For example, with the current 70 timesteps per cycle, subsampling is limited to the prime factors of 69; i.e. only every 3^{rd} or 23^{rd} sample can be chosen. Instead, if the CFD sampling strategy is changed to generate 71 timesteps per cycle, subsampling can be done every 2^{nd}, 5^{th}, 7^{th} timestep, the prime factors of 70, as well as every 10^{th}, 14^{th}, or 35^{th} timestep. Finally, simulations using fewer, non-uniformly spaced samples were attempted but encountered significant convergence issues and increased computational times. Switching to high-order time-stepping may address the convergence issue, possibly allowing for the use of larger and non-uniform Δt, but that investigation is left as an area for future research.

7.5 Conclusions

This chapter demonstrated that, for non-steady flow conditions, accurate compartmental modelling methods require explicit consideration of time-varying flow dynamics. Three different approaches were tested, with the results summarized in Table 7.1 below.

Traditional time-averaged CM approaches proved inadequate for both mean value predictions and flow oscillations. While the ensemble approach achieved computational speeds two orders of magnitude faster than CFD and accurately reproduced time-averaged concentrations, it could not describe the temporal oscillations. The dynamic-projection-based approach captured periodic behaviour but suffered from phase shifts and projection-related inaccuracies, with computational costs exceeding CFD for the tested scenario. However,

Table 7.1: Comparison of the three proposed approaches for CM simulations of systems with time-varying flows.

Approach	**Run time [s]**	**Positives**	**Negatives**
Time-average	3	Fast	Incorrect mean No oscillations
Ensemble	113	Correct mean	Inadequate oscillation bands
Dynamic-Projection	2,500	Correct period	Incorrect magnitudes Slow

these computational time comparisons particularly favour CFD due to the simplified domain and flow conditions studied. More complex geometries, such as the bioreactors examined in previous chapters, may yield computational advantages for CM approaches compared to the expensive CFD simulation.

Future development of these approaches should focus on addressing the time delay inaccuracies, projection errors, and computational inefficiencies identified in the projection-based approach. Additionally, research should examine methods for estimating and utilizing $\boldsymbol{u}'$ from existing CFD data, which could enable the usage of the more computationally efficient time-averaged CM approach. Resolving these limitations would enhance CM accuracy for inherently time-varying flows, extending CM practical applicability beyond steady-state scenarios.

Chapter 8

Conclusions and Recommendations

8.1 Conclusions

This thesis proposed a flow alignment-informed method for creating CMs for complex fluid flow systems. The work encompasses several key contributions related to the modelling of single-phase and multi-phase coupled hydrodynamic-mass transport problems by CMs:

1. **Developed a novel flow-alignment based compartmentalization method.** The method, outlined in Chapter 4, exploits the one-way coupling between hydrodynamics and reactions common in many engineering systems. By using flow-topology to identify unidirectional compartments, the method allows for a geometry-independent approach. The method's prescriptive nature ensures consistent application across geometries and process conditions, providing a systematic approach to compartmentalization that is independent of the underlying chemical reactions.

2. **Developed a compartment model for unidirectional compartments.** Each compartment is modelled as a series of PFRs aligned with the average flow direction, enabling accurate representation of mass transport dynamics and time delays. This approach maintains consistency with the compartmentalization criteria by utilizing the same unidirectional flow criteria. This congruence between compartment identification and modelling approach was demonstrated to improve overall CM accuracy. Further, the method preserves the spatial ordering of flows between compartments while capturing local variations in concentration. Finally, the increased DoFs, compared to traditional CSTR-based methods, allow for more accurate representation of spatial concentration gradients along the primary flow direction.

3. **Performance of the proposed method was better than previous CMs.**
   As demonstrated in Chapter 5, the proposed method achieves superior prediction accuracy for both concentration profiles and RTDs compared to traditional CSTR-based CMs. Unlike CSTR-based methods, increasing the number of compartments showed systematic convergence toward CFD simulation results. While this improved accuracy incurs additional computational cost compared to CSTR-based models, the method maintains a significant computational advantage, executing three orders of magnitude faster than the CFD simulations.

4. **Validated the model on a 3D two-phase stirred tank bioreactor.**
   Application of the method to an industrially-relevant 3D two-phase systems in Chapter 6 validated its capability to model industrially relevant geometries and operating conditions. The approach systematically captured complex flow topologies that previous methods failed to capture, achieving performance comparable to or better than geometry-specific CMs. For these more complex systems, the computational advantage became more pronounced, with the CM executing five orders of magnitude faster than corresponding CFD simulations.

5. **Incorporated spatially-varying interphase mass transport into the CM.**
   The framework readily accommodated complex reactions without requiring structural modifications, as demonstrated through the simulation of oxygen dissolution in a two-phase tank bioreactor in Chapter 6. The model successfully predicted dissolved oxygen concentration gradients driven by pressure-dependent saturation concentrations. The prediction of these gradients, their transient behaviour, and distinct recirculatory regions highlights the capabilities and benefits of the method. Operating at speeds five orders of magnitude faster than CFD and one order of magnitude faster than real-time, the method enables practical exploration of control strategies, operating conditions, and design optimization scenarios.

6. **Extended the method to time-varying velocity fields.**
   Investigation of time-varying flows revealed limitations in using time-averaged velocity profiles for compartmental modelling. An ensemble-based approach was developed to capture mean values while providing uncertainty bounds. The spatial correspondence of the proposed compartmental modelling method enabled the development of projection matrices between CMs from different velocity snapshots, successfully capturing oscillation periods despite requiring further improvements. Computational performance varied significantly: the time-averaged approach achieved three orders of magnitude speedup over CFD, the ensemble method two orders, while the dynamic-projection approach operated 50% slower than CFD. These relative performance

metrics are expected to improve with increasing CFD complexity, though the computational efficiency of the dynamic-projection method requires further optimization.

7. **Released an open-source reference implementation of the method.**
   To facilitate broader adoption and further development, the method has been implemented as an open-source Python package, `OpenCCM` [86]. Comprehensive implementation guidance, including practical considerations and detailed code documentation, has been provided to support researchers in applying and extending the method.

8.2 Recommendations for Future Work

The research presented in this thesis demonstrated the effectiveness of flow-informed approaches for creating CFD-based CMs suitable for industrial-scale multi-phase systems. Several key areas warrant further investigation to enhance the method's capabilities:

1. **Network Topology Analysis and Simplification**
   Strategic model reduction could be achieved through selective combination of PFRs based on network topology. A linear arrangement of PFRs with minimal cross-flows could be consolidated into a single PFR, while a tightly coupled recirculation zone could be simplified to a single CSTR. This topology-based simplification would maintain model accuracy while reducing computational cost.

2. **Mesh-independent Method Exploration**
   The current method creates compartments based on mesh elements and facets. However, as shown in Table 5.5, the online computational costs of the CM increases with increasing mesh density, despite the underlying system topology remaining constant. This mesh dependency arises from the algorithm's constraint of bounding surfaces to element facets, rather than allowing surfaces to intersect elements. Developing methods that either eliminate the element-by-element approach or enable bounding surfaces to cut through elements would substantially improve the model's computational efficiency and decrease it's sensitivity to mesh element shape/orientation.

3. **Compartment Model Modifications**
   The current implementation models compartments as a series of PFRs, which, while allowing for high spatial resolution within each compartment, imposes significant computational costs when implementing higher-order finite difference schemes. Because of the large number of PFRs each should be modelled with fewer DoFs, thereby

limiting the selection of available finite difference discretization schemes to low order ones. An alternative approach worth investigating would be modelling compartments using fewer PFRs but with greater number of DoFs, coupled with higher-order upwinding schemes, and compare computational efficiency and model accuracy.

4. **Complex Reaction System Integration**
   The developed method provides a foundation for modelling sophisticated reaction networks, particularly for bioreactor applications. It can accommodate complex biological systems requiring coupled hydrodynamics and mass transport equations for numerous metabolites. Integration with dynamic flux balance analysis models, which can contain hundreds of metabolic reactions, would expand the framework's capabilities for detailed bioprocess modelling.

5. **Algorithm Performance Enhancement**
   The current single-threaded, path-dependent nature of the compartmentalization algorithm limits the performance of the CM construction. Development of parallelizable algorithms would significantly reduce CM creation time. Additional optimization through implementation of performance critical sections in Cython and development of more efficient algorithms would further improve performance.

6. **Literature-based Extension**
   The CM literature contains several extensions to compartmental modelling which remain unimplemented in this work and in `OpenCCM`. These extensions include the capability to model variable working volumes in fed-batch operations as well as changes to operating conditions (e.g. impeller speeds, viscosity, or sparging rates). The integration of these features would enhance the practical utility of the method, bridging the gap between theoretical advancements and a usable tool.

7. **Conservation of Energy**
   The CM-based simulations presented in this work were limited to isothermal conditions. However, many target applications, particularly bioreactors, involve non-isothermal reactions where temperature control is crucial for process performance. While the compartmental modelling framework can account for most heat transfer phenomena, it currently lacks the capability to model heat transfer through domain boundaries, such as those occurring in jacketed stirred tank reactors. Extension of the model to incorporate wall heat transfer, potentially through conversion of boundary conditions to source terms, is necessary for accurate simulation of these processes.

8. **Control Scheme Implementation**
   Integration of industrial control algorithms, particularly PID controllers, for pro-

cess variables like dissolved oxygen, temperature, and pH, would enhance simulation realism. This capability would also enable control system tuning and strategy optimization by leveraging the computational advantage of CM over CFD simulations and accuracy advantage over simplistic CSTR models.

9. **Parameter Selection Criteria Development**
   Development of systematic guidelines for parameter selection, analogous to mesh generation criteria in CFD, would standardize model development. These criteria should address compartment size selection, discretization parameters, and tolerance settings based on flow characteristics and accuracy requirements.

10. **DoF Count Optimization at Runtime**
    The current implementation assigns an equal number of DoFs to each PFR. Since the compartments vary in size, the use of a constant number of DOFs leads to a spatially varying local error. A more sophisticated approach would allow variable DoF allocation per PFR, coupled with automatic runtime determination of required DoFs through multiple simulations and adaptive refinement. The computational efficiency of the CM framework makes this adaptive approach feasible. Implementation of such a system would simultaneously reduce the number of user-specified parameters, optimize computational resources, and enhance model accuracy.

11. **Numerical Method Improvements**
    Replacement of current FDM discretization with FVM or DG-FEM would improve numerical stability and accuracy. These methods would better handle steep concentration gradients, reduce numerical artifacts, and numerical diffusion, while maintaining computational efficiency.

12. **Seed Selection Improvements**
    Following the exhaustion of user-provided seed elements, the most upstream unassigned element from those previously encountered is used. The choice made significantly influences the resulting compartmentalization, as discussed in Section 4.1.2. For systems which possess well-defined inlets and outlets, the determination of "most upstream" element is straightforward. However, this criterion becomes ill-defined in closed systems, where the element connectivity is fully cyclical. Further research into seed selection for closed systems is important given the direct impact it has on the resulting number of compartments, and hence the computational cost.

13. **Validation Case Outlier Investigation**
    Within the 2D validation study, one intermediary mesh resolution yielded notably

different results compared to other mesh densities. Although a preliminary investigation was conducted, no conclusive explanation for this phenomenon was identified. This unexpected behaviour merits systematic examination to understand its underlying mechanisms. Specifically, given that the compartmentalization method is mesh-dependent, further investigation into the influence of mesh element location and orientation should be pursued.

14. **Time-Averaged Flow Applicability Criteria Development**
While Chapter 7 demonstrated that a CM based on time-averaged velocities was inadequate for the system under investigation, no quantitative metric existed to predict this limitation *a priori*. Development of quantitative metrics and dimensionless parameters to assess the validity of time-averaged flow approximations is essential. These criteria would guide users in selecting appropriate modelling approaches based on flow characteristics, helping to balance computational cost with required accuracy.

15. **Time-Averaging Diffusivity Closure Investigation**
The CM based on time-averaged velocities presented in Chapter 7 consistently underpredicted the time-averaged tracer concentrations obtained from CFD simulations. This underprediction stems from the time-averaging out of the vertical velocity component, which significantly affects mass transport. Because of the availability of multiple snapshots of flow field over the course of the oscillations, the Reynolds stress term could be calculated directly, rather than requiring a closure. Future work should explore the application of these terms to develop an anisotropic effective diffusivity, following methods similar to those employed in RANS turbulence modelling

16. **Ensemble Approach Extensions**
The current analysis of the ensemble results is limited to a min, max, and mean value at each point in time. The different simulation trajectories present rich data for further analysis, starting with perhaps replacing min and max with standard deviation. Further, the different trajectories can be used to develop quantities for identifying whether the ensemble method is required or if the simpler time-averaged flow will provide acceptable results.

17. **Projection Approach Improvements**
While successful in capturing oscillation frequencies, the projection-based approach requires improvements to accurately predict concentrations. Investigation of alternative projection operators and temporal discretization schemes could improve accuracy while reducing computational overhead.

References

[1] A. W. Nienow, "Stirring and Stirred-Tank Reactors," *Chemie Ingenieur Technik*, vol. 86, no. 12, pp. 2063–2074, 2014.

[2] H. Y. Farsani, J. Wutz, B. DeVincentis, J. A. Thomas, and S. P. Motevalian, "Modeling mass transfer in stirred microbioreactors," *Chemical Engineering Science*, vol. 248, p. 117146, Feb. 2022.

[3] P. Vrábel, R. G. J. M. Van der Lans, Y. Q. Cui, and K. Ch. A. M. Luyben, "Compartment Model Approach: Mixing in Large Scale Aerated Reactors with Multiple Impellers," *Chemical Engineering Research and Design*, vol. 77, no. 4, pp. 291–302, Jun. 1999.

[4] C. Haringa, R. Vandewijer, and R. F. Mudde, "Inter-compartment interaction in multi-impeller mixing: Part I. Experiments and multiple reference frame CFD," *Chemical Engineering Research and Design*, vol. 136, pp. 870–885, Aug. 2018.

[5] G. Tryggvason, R. Scardovelli, and S. Zaleski, *Direct Numerical Simulations of Gas–Liquid Multiphase Flows*. Cambridge: Cambridge University Press, 2011.

[6] C. Haringa, "An analysis of organism lifelines in an industrial bioreactor using Lattice-Boltzmann CFD," *Engineering in Life Sciences*, vol. 23, no. 1, p. e2100159, Jan. 2023.

[7] G. Larsson, M. Törnkvist, E. S. Wernersson, C. Trägårdh, H. Noorman, and S. O. Enfors, "Substrate gradients in bioreactors: Origin and consequences," *Bioprocess Engineering*, vol. 14, no. 6, pp. 281–289, May 1996.

[8] A. R. Lara, E. Galindo, O. T. Ramírez, and L. A. Palomares, "Living with heterogeneities in bioreactors," *Molecular Biotechnology*, vol. 34, no. 3, pp. 355–381, Nov. 2006.

[9] R. Spann, J. Glibstrup, K. Pellicer-Alborch, S. Junne, P. Neubauer, C. Roca, D. Kold, A. E. Lantz, G. Sin, K. V. Gernaey, and U. Krühne, "CFD predicted pH gradients in lactic acid bacteria cultivations," *Biotechnology and Bioengineering*, vol. 116, no. 4, pp. 769–780, 2019.

[10] G. Nadal-Rey, D. D. McClure, J. M. Kavanagh, S. Cornelissen, D. F. Fletcher, and K. V. Gernaey, "Understanding gradients in industrial bioreactors," *Biotechnology Advances*, vol. 46, p. 107660, Jan. 2021.

[11] C. Mayorga, J. Morchain, and A. Liné, "Reconstruction of the 3D hydrodynamics in a baffled stirred tank using Proper Orthogonal Decomposition," *Chemical Engineering Science*, vol. 248, p. 117220, Feb. 2022.

[12] I. Promma, M. G. Aucoin, N. M. Abukhdeir, and H. Budman, "A coupled metabolic flux/compartmental hydrodynamic model for large-scale aerated bioreactors," *Computers & Chemical Engineering*, p. 108806, Jul. 2024.

[13] O. Šrom, M. Šoóš, M. Kuschel, T. Wucherpfennig, J. Fitschen, and M. Schlüter, "Study of hydrodynamic stress in cell culture bioreactors via lattice-Boltzmann CFD simulations supported by micro-probe shear stress method," *Biochemical Engineering Journal*, vol. 208, p. 109337, Aug. 2024.

[14] F. Chinesta, A. Huerta, G. Rozza, and K. Willcox, "Model Reduction Methods," in *Encyclopedia of Computational Mechanics Second Edition*. John Wiley & Sons, Ltd, 2017, pp. 1–36.

[15] P. Benner, W. Schilders, S. Grivet-Talocia, A. Quarteroni, G. Rozza, and L. Miguel Silveira, Eds., *Model Order Reduction: Volume 3 Applications*. De Gruyter, 2020.

[16] N. Jourdan, T. Neveux, O. Potier, M. Kanniche, J. Wicks, I. Nopens, U. Rehman, and Y. Le Moullec, "Compartmental Modelling in chemical engineering: A critical review," *Chemical Engineering Science*, vol. 210, p. 115196, Dec. 2019.

[17] J. Haag, C. Gentric, C. Lemaitre, and J.-P. Leclerc, "Modelling of Chemical Reactors: From Systemic Approach to Compartmental Modelling," *International Journal of Chemical Reactor Engineering*, vol. 16, no. 8, Aug. 2018.

[18] Y. Lu, J. Dong, and J. Liu, "Zonal modelling for thermal and energy performance of large space buildings: A review," *Renewable and Sustainable Energy Reviews*, vol. 133, p. 110241, Nov. 2020.

[19] G. Nadal-Rey, "Modelling of gradients in industrial aerobic fed-batch fermentation processes," Ph.D. dissertation, Technical University of Denmark, Kgs. Lyngby, 2020.

[20] S. Rigopoulos and A. Jones, "A hybrid CFD—reaction engineering framework for multiphase reactor modelling: Basic concept and application to bubble column reactors," *Chemical Engineering Science*, vol. 58, no. 14, pp. 3077–3089, Jul. 2003.

[21] T. R. Donnelly, "A Topology-Based Method for the Compartmentalization of Multiphysics Flows," Master's thesis, University of Waterloo, Waterloo, Ontario, Canada, Sep. 2020.

[22] M. Pigou and J. Morchain, "Investigating the interactions between physical and biological heterogeneities in bioreactors using compartment, population balance and metabolic models," *Chemical Engineering Science*, vol. 126, pp. 267–282, Apr. 2015.

[23] C. Bach, "Modelling of gradients in large scale bioreactors," Ph.D. dissertation, DTU Chemical Engineering, Copenhagen, Denmark, Apr. 2018.

[24] T. Tajsoleiman, "Automating Experimentation in Miniaturized Reactors," Ph.D. dissertation, Technical University of Denmark, Kgs. Lyngby, 2018.

[25] M. Sussman, P. Smereka, and S. Osher, "A Level Set Approach for Computing Solutions to Incompressible Two-Phase Flow," *Journal of Computational Physics*, vol. 114, no. 1, pp. 146–159, Sep. 1994.

[26] M. Ishii and T. Hibiki, *Thermo-Fluid Dynamics of Two-Phase Flow*. New York, NY: Springer New York, 2011.

[27] C. E. Brennen, *Fundamentals of Multiphase Flow*. Cambridge [England] ; New York: Cambridge University Press, 2005.

[28] M. Khezrian, "Study, Analysis, and Validation of a Specific Two-fluid Model for Dispersed Two-Phase Flow," Master's thesis, University of Waterloo, Waterloo, Ontario, Canada, Jan. 2022.

[29] T. Treeratanaphitak, "Diffuse Solid-Fluid Interface Method for Dispersed Multiphase Flows," Ph.D. dissertation, University of Waterloo, Waterloo, Ontario, Canada, Aug. 2018.

[30] K. Entezari, "Simulation-based Design of Bioreactors Using Computational Multiphysics," Master's thesis, University of Waterloo, Waterloo, Ontario, Canada, 2021.

[31] F. Chen, H. Xiong, Z.-Y. Yin, and X. Chen, "Impermeable and mechanical stability of filter cake under different infiltration conditions via CFD-DEM," *Acta Geotechnica*, vol. 18, no. 8, pp. 4115–4140, Aug. 2023.

[32] J. A. Thomas, X. Liu, B. DeVincentis, H. Hua, G. Yao, M. C. Borys, K. Aron, and G. Pendse, "A mechanistic approach for predicting mass transfer in bioreactors," *Chemical Engineering Science*, vol. 237, p. 116538, Jun. 2021.

[33] B. Andersson, R. Andersson, L. Hakansson, M. Mortensen, R. Sudiyo, and B. van Wachem, *Computational Fluid Dynamics for Engineers*. Cambridge ; New York: Cambridge University Press, 2012.

[34] G. Alfonsi, "Reynolds-Averaged Navier–Stokes Equations for Turbulence Modeling," *Applied Mechanics Reviews*, vol. 62, no. 4, p. 040802, Jul. 2009.

[35] A. Behzadi, R. I. Issa, and H. Rusche, "Modelling of dispersed bubble and droplet flow at high phase fractions," *Chemical Engineering Science*, vol. 59, no. 4, pp. 759–770, Feb. 2004.

[36] R. T. Lahey, "The simulation of multidimensional multiphase flows," *Nuclear Engineering and Design*, vol. 235, no. 10, pp. 1043–1060, May 2005.

[37] I. Kataoka and A. Serizawa, "Basic equations of turbulence in gas-liquid two-phase flow," *International Journal of Multiphase Flow*, vol. 15, no. 5, pp. 843–855, 1989.

[38] A. D. Gosman, C. Lekakou, S. Politis, R. I. Issa, and M. K. Looney, "Multidimensional modeling of turbulent two-phase flows in stirred vessels," *AIChE Journal*, vol. 38, no. 12, pp. 1946–1956, Dec. 1992.

[39] C. Yang and Z.-S. Mao, "Multiphase stirred reactors," in *Numerical Simulation of Multiphase Reactors with Continuous Liquid Phase*. Elsevier, 2014, pp. 75–151.

[40] C. Haringa, R. Vandewijer, and R. F. Mudde, "Inter-compartment interaction in multi-impeller mixing. Part II. Experiments, sliding mesh and large Eddy simulations," *Chemical Engineering Research and Design*, vol. 136, pp. 886–899, Aug. 2018.

[41] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, Apr. 2016.

[42] K. Willcox, "Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition," *Computers & Fluids*, vol. 35, no. 2, pp. 208–226, Feb. 2006.

[43] K. Lee and K. T. Carlberg, "Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders," *Journal of Computational Physics*, vol. 404, p. 108973, Mar. 2020.

[44] Y. Du, H. Zhao, A. Ma, and C. Yang, "Equivalent Reactor Network Model for the Modeling of Fluid Catalytic Cracking Riser Reactor," *Industrial & Engineering Chemistry Research*, vol. 54, no. 35, pp. 8732–8742, Sep. 2015.

[45] U. Rehman, "Next generation bioreactor models for wastewater treatment systems by means of detailed combined modelling of mixing and biokinetics," Ph.D. dissertation, Ghent University, 2016.

[46] S. Yang, S. Kiang, P. Farzan, and M. Ierapetritou, "Optimization of Reaction Selectivity Using CFD-Based Compartmental Modeling and Surrogate-Based Optimization," *Processes*, vol. 7, no. 1, p. 9, Jan. 2019.

[47] A. Delafosse, F. Delvigne, M.-L. Collignon, M. Crine, P. Thonart, and D. Toye, "Development of a compartment model based on CFD simulations for description of mixing in bioreactors," *BASE*, Jan. 2010.

[48] E. K. Nauha, Z. Kálal, J. M. Ali, and V. Alopaeus, "Compartmental modeling of large stirred tank bioreactors with high gas volume fractions," *Chemical Engineering Journal*, vol. 334, pp. 2319–2334, Feb. 2018.

[49] E. K. Nauha, "Modeling of bioreactors," Ph.D. dissertation, Aalto Univerty, Espoo, Finland, 2019.

[50] S. Trespi, H. Nicolai, P. Debiagi, J. Janicka, A. Dreizler, C. Hasse, and T. Faravelli, "Development and Application of an Efficient Chemical Reactor Network Model for Oxy-fuel Combustion," *Energy & Fuels*, vol. 35, no. 9, pp. 7121–7132, May 2021.

[51] A. Alvarado, S. Vedantam, P. Goethals, and I. Nopens, "A compartmental model to describe hydraulics in a full-scale waste stabilization pond," *Water Research*, vol. 46, no. 2, pp. 521–530, Feb. 2012.

[52] T. H. Nguyen, "Improved Chemical Reactor Network Application for Predicting the Emission of Nitrogen Oxides in a Lean Premixed Gas Turbine Combustor," *Combustion, Explosion, and Shock Waves*, vol. 55, no. 3, pp. 267–273, May 2019.

[53] C. Haringa, W. Tang, and H. J. Noorman, "Stochastic parcel tracking in an Euler–Lagrange compartment model for fast simulation of fermentation processes," *Biotechnology and Bioengineering*, vol. 119, no. 7, pp. 1849–1860, 2022.

[54] G. Liu and B. A. Wilhite, "A generalized zoning framework for development of CFD-based reduced-order compartment models," *AIChE Journal*, Feb. 2023.

[55] Y. Bai, C. P. Fonte, A. Kowalski, W. Rosales-Trujillo, and T. L. Rodgers, "A general methodology for extracting and categorising flow patterns in turbulent stirred tank mixers based on 2-D network-of-zones (NoZ) model," *Chemical Engineering Research and Design*, vol. 194, pp. 829–841, Jun. 2023.

[56] H. Hartmann, J. J. Derksen, and H. E. A. Van Den Akker, "Mixing times in a turbulent stirred tank by means of LES," *AIChE Journal*, vol. 52, no. 11, pp. 3696–3706, Nov. 2006.

[57] D. W. Pepper, S. Pirbastami, and D. B. Carrington, "A Comparison Between FEM and FVM via the Method of Weighted Residuals," in *50 Years of CFD in Engineering Sciences: A Commemorative Volume in Memory of D. Brian Spalding*, A. Runchal, Ed. Singapore: Springer, 2020, pp. 753–777.

[58] B. Cockburn, G. E. Karniadakis, and C.-W. Shu, Eds., *Discontinuous Galerkin Methods: Theory, Computation and Applications*, ser. Lecture Notes in Computational Science and Engineering. Berlin: Springer, 2000, no. 11.

[59] V. Dolejší and M. Feistauer, *Discontinuous Galerkin Method*, ser. Springer Series in Computational Mathematics. Cham: Springer International Publishing, 2015, vol. 48.

[60] O. Bokhove and J. J. van der Vegt, "Introduction to (dis) continuous Galerkin finite element methods," *Department of Applied Mathematics, University of Twente*, 2005.

[61] E. Monte, "OpenCMP: An Open-Source Computational Multiphysics Package," Master's thesis, University of Waterloo, Aug. 2021.

[62] J. Andrejevic and C. Ding, "Solving differential-algebraic systems of equations (DAEs)," Nov. 2020.

[63] M. Yousefi, K. van Heusden, G. A. Dumont, I. M. Mitchell, and J. M. Ansermino, "Model-invariant safety-preserving control," in *2016 American Control Conference (ACC)*, Jul. 2016, pp. 6689–6694.

[64] M. F. Villa-Tamayo, M. A. Caicedo, and P. S. Rivadeneira, "Offset-free MPC strategy for nonzero regulation of linear impulsive systems," *ISA Transactions*, vol. 101, pp. 91–101, Jun. 2020.

[65] T. Marzullo, M. M. Keane, M. Geron, and R. F. D. Monaghan, "A computational toolchain for the automatic generation of multiple Reduced-Order Models from CFD simulations," *Energy*, vol. 180, pp. 511–519, Aug. 2019.

[66] J. Le Nepvou De Carfort, T. Pinto, and U. Krühne, "An Automatic Method for Generation of CFD-Based 3D Compartment Models: Towards Real-Time Mixing Simulations," *Bioengineering*, vol. 11, no. 2, p. 169, Feb. 2024.

[67] G. D. Wehinger, B. Kreitz, and C. F. Goldsmith, "Non-Idealities in Lab-Scale Kinetic Testing: A Theoretical Study of a Modular Temkin Reactor," *Catalysts*, vol. 12, no. 3, p. 349, Mar. 2022.

[68] M. A. Agizza, G. Bagheri, S. Bürkle, T. Faravelli, S. Wagner, and A. Dreizler, "Investigation of Oxy-Fuel Combustion through Reactor Network and Residence Time Data," *Energies*, vol. 15, no. 1, p. 252, Jan. 2022.

[69] T. Tajsoleiman, R. Spann, C. Bach, K. V. Gernaey, J. K. Huusom, and U. Krühne, "A CFD based automatic method for compartment model development," *Computers & Chemical Engineering*, vol. 123, pp. 236–245, Apr. 2019.

[70] J. Lindfors, "Performance of cracked ammonia combustion in a gas turbine engine - Evaluation through CFD and chemical reactor network modeling," Master's thesis, Chalmers University of Technology, Gothenburg, Sweden, 2022.

[71] M. Savarese, A. Cuoci, W. De Paepe, and A. Parente, "Machine learning clustering algorithms for the automatic generation of chemical reactor networks from CFD simulations," *Fuel*, vol. 343, p. 127945, Jul. 2023.

[72] Y. Du, S. Li, X. Chen, Y. Li, B. Wang, and C. Yang, "A novel methodology to construct compartment models for a circulating fluidized bed riser," *Chemical Engineering Science*, p. 118470, Jan. 2023.

[73] A. Delafosse, M.-L. Collignon, S. Calvo, F. Delvigne, M. Crine, P. Thonart, and D. Toye, "CFD-based compartment model for description of mixing in bioreactors," *Chemical Engineering Science*, vol. 106, pp. 76–85, Mar. 2014.

[74] A. Delafosse, S. Calvo, M.-L. Collignon, F. Delvigne, M. Crine, and D. Toye, "Euler–Lagrange approach to model heterogeneities in stirred tank bioreactors – Comparison to experimental flow characterization and particle tracking," *Chemical Engineering Science*, vol. 134, pp. 457–466, Sep. 2015.

[75] A. Innocenti, A. Andreini, D. Bertini, B. Facchini, and M. Motta, "Turbulent flow-field effects in a hybrid CFD-CRN model for the prediction of NOx and CO emissions in aero-engine combustors," *Fuel*, vol. 215, pp. 853–864, Mar. 2018.

[76] N. Jourdan, M. Kanniche, T. Neveux, and O. Potier, "Compartmental Modeling of Biofouling in Industrial Cooling Circuits," *Chemical Engineering & Technology*, vol. 46, no. 6, pp. 1140–1148, 2023.

[77] G. Nadal-Rey, D. McClure, J. Kavanagh, B. Cassells, S. Cornelissen, D. Fletcher, and K. Gernaey, "Development of dynamic compartment models for industrial aerobic fed-batch fermentation processes," *Chemical Engineering Journal*, vol. 420, 2021.

[78] J. Morchain, C. Mayorga, P. Villedieu, and A. Liné, "A dynamic compartment model for spatially heterogeneous reactors: Scalar and Monte-Carlo particle mixing," *Chemical Engineering Research and Design*, vol. 205, no. Complete, pp. 628–639, 2024.

[79] V. Fichet, M. Kanniche, P. Plion, and O. Gicquel, "A reactor network model for predicting NOx emissions in gas turbines," *Fuel*, vol. 89, no. 9, pp. 2202–2210, Sep. 2010.

[80] T. Karches, "Towards a Dynamic Compartmental Model of a Lamellar Settler," *Symmetry*, vol. 15, no. 4, p. 864, Apr. 2023.

[81] A. A. Vasile, M. G. Aucoin, H. Budman, and N. M. Abukhdeir, "A flow alignment-informed method for compartmental modelling," *Computers & Chemical Engineering*, vol. 185, p. 108650, Jun. 2024.

[82] E. J. Monte, A. A. Vasile, J. Lowman, and N. M. Abukhdeir, "OpenCMP: An Open-Source Computational Multiphysics Package," *Journal of Open Source Software*, vol. 7, no. 73, p. 3742, May 2022.

[83] H. S. Fogler, *Elements of Chemical Reaction Engineering*, 5th ed. Boston: Prentice Hall, 2016.

[84] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, Mar. 1951.

[85] R. Bird, W. Stewart, and E. Lightfoot, *Transport Phenomena*, 2nd ed. New York: John Wiley and Sons, 2002.

[86] A. A. Vasile, M. P. Tino, Y. Aseri, and N. M. Abukhdeir, "OpenCCM: An Open-Source Continuous Compartmental Modelling Package," *Journal of Open Source Software*, vol. 9, no. 102, p. 6963, Oct. 2024.

[87] S. Stocks, "Industrial enzyme production for the food and beverage industries: Process scale up and scale down," in *Microbial Production of Food Ingredients, Enzymes and Nutraceuticals*. Elsevier, 2013, pp. 144–172.

[88] M. Ishii and N. Zuber, "Drag coefficient and relative velocity in bubbly, droplet or particulate flows," *AIChE Journal*, vol. 25, no. 5, pp. 843–855, Sep. 1979.

[89] A. Tomiyama, H. Tamai, I. Zun, and S. Hosokawa, "Transverse migration of single bubbles in simple shear flows," *Chemical Engineering Science*, vol. 57, no. 11, pp. 1849–1858, Jun. 2002.

[90] P. Vrábel, R. G. J. M. van der Lans, F. N. van der Schot, K. C. A. M. Luyben, B. Xu, and S.-O. Enfors, "CMA: Integration of fluid dynamics and microbial kinetics in modelling of large-scale fermentations," *Chemical Engineering Journal*, vol. 84, no. 3, pp. 463–474, Dec. 2001.

[91] J. C. Lamont and D. S. Scott, "An eddy cell model of mass transfer into the surface of a turbulent liquid," *AIChE Journal*, vol. 16, no. 4, pp. 513–519, 1970.

[92] B. Xu, M. Jahic, G. Blomsten, and S. O. Enfors, "Glucose overflow metabolism and mixed-acid fermentation in aerobic large-scale fed-batch processes with Escherichia coli," *Applied Microbiology and Biotechnology*, vol. 51, no. 5, pp. 564–571, May 1999.

[93] M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher, "Benchmark Computations of Laminar Flow Around a Cylinder," in *Flow Simulation with High-Performance Computers II: DFG Priority Research Programme Results 1993–1995*, E. H. Hirschel, Ed. Wiesbaden: Vieweg+Teubner Verlag, 1996, pp. 547–566.

[94] C. Greenshields, *OpenFOAM V11 User Guide*. London, UK: The OpenFOAM Foundation, 2023.

[95] U. Ayachit, *The ParaView Guide: Updated for ParaView Version 4.3*, L. Avila, Ed. Clifton Park, NY: Kitware Inc, 2015.

[96] Quintessa, "Compartment modelling software," Online.

[97] D. G. Goodwin, H. K. Moffat, I. Schoegl, R. L. Speth, and B. W. Weber, "Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes," 2023.

APPENDICES

Appendix A

Implementation Details of Method in OpenCCM

This chapter provides an overview[1] of the implementation of the method described in Section 4.1.

A.1 Summary

`OpenCCM` is a compartmental modelling [16] software package which is based on recently-developed fully automated flow alignment compartmentalization methods [81]. It is primarily intended for large-scale flow-based processes where there is weak coupling between composition changes, e.g. through (bio)chemical reactions, and convective mass transport in the system. Compartmental modelling is an important approach used to developed reduced-order models [14, 15] using *a priori* knowledge of process hydrodynamics [16]. Compartmental modelling methods, such as those implemented in `OpenCCM`, enable simulations of these processes to be performed with far less computational complexity while still capturing the key aspects of their performance.

`OpenCCM` integrates with two multiphysics simulation software packages, `OpenCMP` [82] and `OpenFOAM` [94], allowing for ease of transferring simulation data for compartmentalization. Additionally, it provides users with built-in functionality for calculating residence times, exporting to transfer data to simulation software, and for exporting results for

[1] Adapted from [86].

visualization using `ParaView` [95]. Post-processing methods are included for mapping simulation results from compartment domains to the original simulation domain, useful for both visualization purposes and for further simulations in other software (e.g. multi-scale modelling).

A.2 Statement of Need

Simulation-based design and analysis continues to be widely applied in research and development of physicochemical processes. Processes with large differences in characteristic time- and length-scale result in the infeasibility of direct multiphysics simulations due to computational limitations. This imposes a significant computation costs which severely reduces the utility of these simulations for entire classes of processes. Compartmental modelling is well-suited for such applications because it produces reduced-order models which are orders-of-magnitude less computationally demanding than direct simulation by taking advantage of the weak coupling between the short and long time-scale phenomena.

However, there are several barriers to the more widespread use of compartmental models. The largest of these is the lack of software for automatically generating compartmental models. There exists closed-source software, specifically `AMBER` [96], for manually creating and solving well-mixed compartment networks. However, this software is cost-prohibitive for much of the research community and lacks automated compartmentalization. There also exists open-source software, `Cantera` [97], for solving compartment networks. However, it does not incorporate flow information, when available, either from direct observation or multiphysics simulation. Furthermore, neither of these software allow for usage and direct transfer of flow information from multiphysics simulations, such as computational fluid dynamics (CFD) simulations, which are typically feasible over short times-scales.

The overall aim of `OpenCCM` is to fill this need for an open-source compartmental modelling package which is user-friendly, compatible with a variety of simulation package backends (e.g. `OpenFOAM` and `OpenCMP`), and which fits into the users existing simulation and post-processing software toolchain, i.e. `ParaView`.

A.3 Features

Feature	Description
Model support	Accepts `OpenCMP` and `OpenFOAM` results
Compartmentalization	Single-phase flow-based compartment identification
Compartmental Modelling	Plug Flow Reactors (PFRs)-in-series-based model Previous state-of-the-art CSTR-based models
CM Simulations	Linear, non-linear, and reversible arbitrary reactions 1st Order upwinding finite-difference-based Adaptive time-stepping
Post-Processing	Residence time distribution
Output	Intermediary mesh format Labelled compartments in `Paraview` format Results in both `Paraview` and simulation package format
Performance	Multi-threading Caching of intermediary results to speed-up subsequent runs

A.4 User Interface

The `OpenCCM` Python package can be used via text-based configuration files centered around the CLI (command line interface) where each simulation run/project is in a self-contained directory. In addition to the `OpenCCM` configuration files, the required contents include flow information from one of two open source simulation packages: `OpenCMP` and `OpenFOAM`. For `OpenCMP` three files are required:

1. the `OpenCMP` config file,
2. the mesh on which the simulation was run, and
3. the `.sol` solution file that contains the velocity profile to use for creating the compartmental model.

For `OpenFOAM`, two directories are required:

1. the `constant` directory which contains the mesh information in ASCII format,
2. a directory containing the simulation results to be used for creating the compartmental model saved in ASCII format.

The path to the solution directory is specified in the `OpenCCM` config, and the `constant` directory is assumed to be in the same parent folder. The `OpenCCM` software will create several output directories: `log/` which contains detailed debugging information (if enabled), `cache/` which contains intermediary files, and `output_ccm/` both simulation results of the compartmental model (in various user-specified formats) and `ParaView` files for visualization.

A sample config file, `CONFIG`, which outlines the available parameters is included in the main directory. An excerpt of it, showing the compartmental modelling parameters, is shown below. Square brackets indicates values which have a default with the default provided inside the brackets.

```
[COMPARTMENT MODELLING]
# Whether to use PFRs-in-series or a CSTR to model each compartment.
model                    = PFR
# Volumetric flow through a single facet
# below which the flow is considered 0.
flow_threshold_facet     = [1e-15]
# Volumetric flow threshold through a surface
# below which the flow is considered 0.
flow_threshold           = [1e-15]
# Maximum allowable difference (in % of compartment volume)
# between connections for merging them to one location.
dist_threshold           = [5 / 100]
# Absolute tolerances for checking that conservation of mass
# after the flow optimization is performed.
atol_opt                 = [1e-2]
```

A.5 Reaction Config File

The chemical reaction parser in `OpenCCM` reads and parses the reactions configuration files and can handle general reaction equations of the form,

```
aA + bB + [...]  -> cC + dD + [...]
```

with associated numeric rate constants. It intentionally does not support the standard `<->` symbol for reversible chemical reactions, so that each independent reaction has its rate constant explicit defined. Therefore, a reversible reaction must be written as two

independent forward reactions, each with its own rate constants. Each species name can contain letter and numbers, but cannot contain brackets "(" or ")" or special characters, e.g. "+", "-", "^", etc. Kinetic rate constants must be expressed as positive real numbers in standard or scientific notation. Additionally, each reaction/rate pair must have a unique *identifier* (i.e. R1, R2). For example, take the reversible reaction,

$$2\mathrm{NaCl} + \mathrm{CaCO_3} \Leftrightarrow \mathrm{Na_2CO_3} + \mathrm{CaCl_2}$$

with $k_f = 5 \times 10^{-2}$ and $k_r = 2$, dimensionless numbers picked for demonstration purposes. A configuration file for this reversible reaction may then be:

```
[REACTIONS]
R1: 2NaCl   + CaCO3 ->  Na2CO3 + CaCl2
R2:  Na2CO3 + CaCl2 -> 2NaCl   + CaCO3

[RATE CONSTANTS]
R1: 5e-2
R2: 2
```

where `R1` and `R2` are the reaction *identifiers* for the forward and reverse reactions, respectively.

An example reaction config file, `CONFIG_REACTIONS` is provided in the main directory.

A.6 Examples of Usage

Several examples are provided in the `OpenCCM` documentation which demonstrate the usage of both `OpenCMP` and `OpenFOAM` simulation flow information for compartmentalization. One example, inside `examples/OpenCMP/pipe_with_recird_2d`, uses the geometry from Vasile et al. [81] and shows how to execute the needed CFD simulation for flow information (both using `OpenCMP` and `OpenFOAM`), create/visualize the compartmental model results, and compare the predicted RTD to the reference result directly from CFD simulation.

For this illustrative example, the steady-state hydrodynamic flow-profile is obtained by running the `OpenCMP` simulation through the `run_OpenCMP.py` script in the folder. The resulting flow profile was opened in `ParaView` and the line integral convolution of the velocity field is shown below, colored by velocity magnitude.

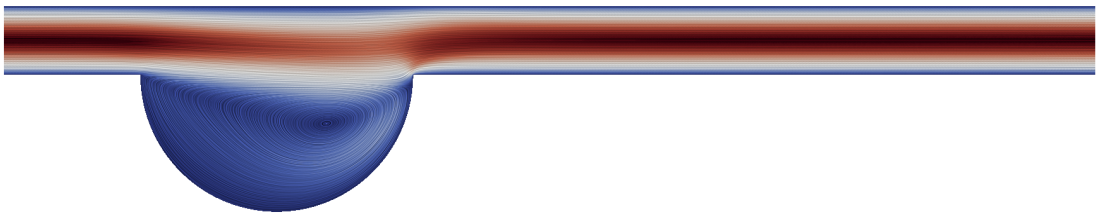


Figure A.1: Visualization of hydrodynamics from CFD simulation with line integral convolutions indicating local flow direction and color corresponding to velocity magnitude.

The underlying velocity field data is then processed using `OpenCCM` to produce a network of compartments by executing the `run_compartment.py` script. The figure below shows each element of the original mesh colored according to the compartment it belongs to.

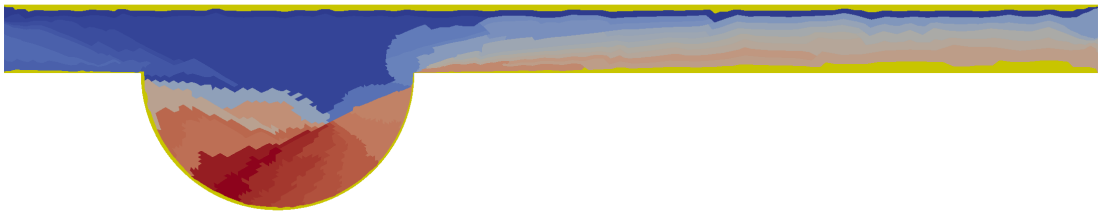


Figure A.2: Visualization of flow-informed compartmentalization with coloring corresponding to compartment number.

That network of compartments is further processed as each compartment is represented by a series of plug flow reactors (PFRs). The resulting network (graph) of PFRs is shown in the figure below; nodes are the centers of the PFRs and edges are connections (flows) between PFRs.

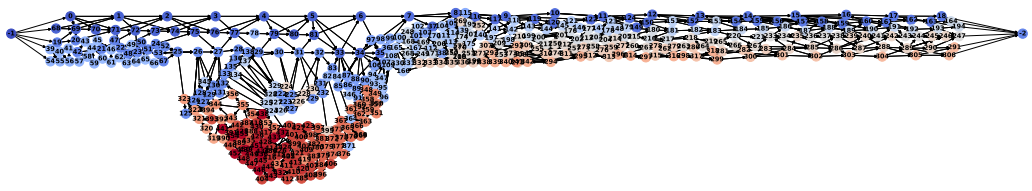


Figure A.3: Undirected graph of the compartment network resulting from both (i) flow-information compartmentalization and (ii) the use of spatially-varying compartment approximations (PFRs).

A.6.1 RTD Curves

The Residence Time Distribution (RTD) curve for both the CFD and Compartmental Model (CM) are calculated using the script in the supplementary material of [@Vasile2024].

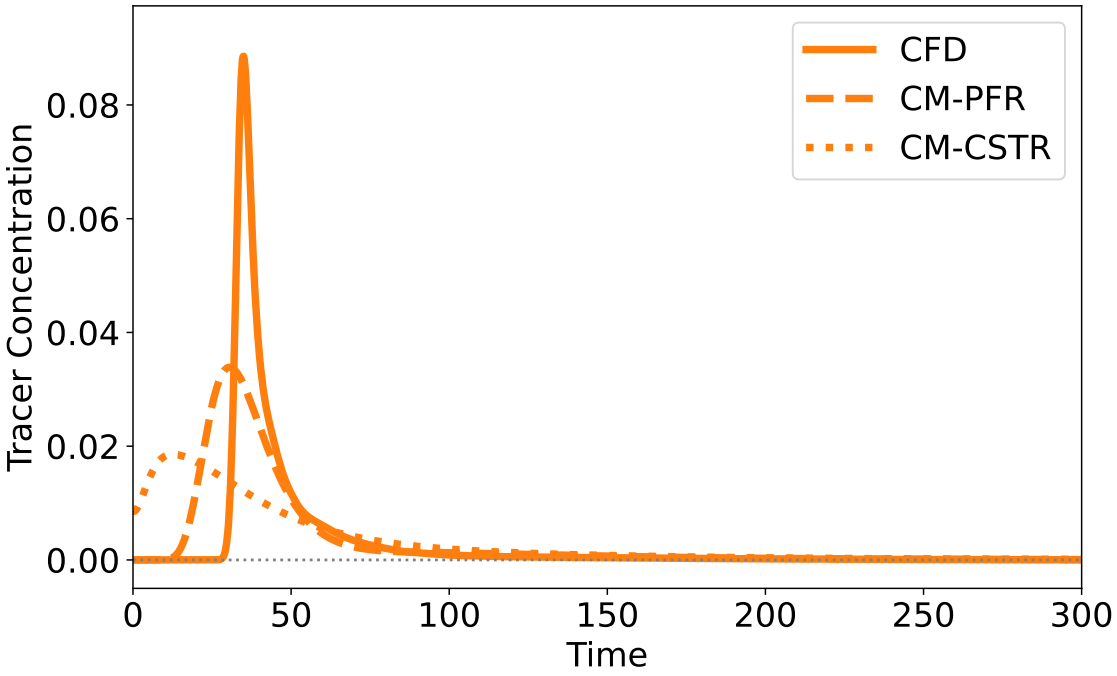


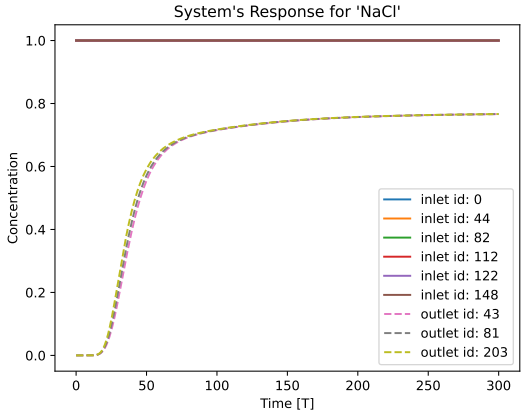
Figure A.4: Residence time distribution curves for CFD and CM simulations.

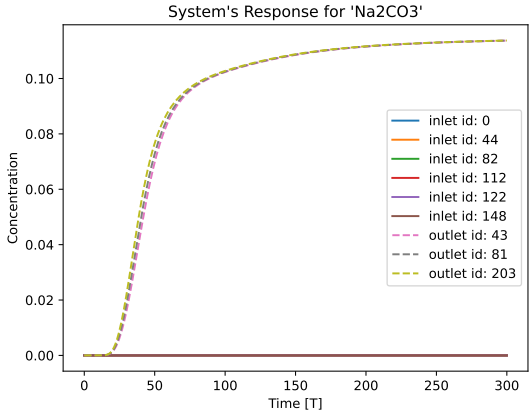
A.6.2 Reactions

Finally, to demonstrate how to use the reaction system we will implement the reversible reaction system mentioned above:

$$2\text{NaCl} + \text{CaCO}_3 \Leftrightarrow \text{Na}_2\text{CO}_3 + \text{CaCl}_2$$

with $k_f = 5 \times 10^{-2}$ and $k_r = 2$ as the dimensionless forward and backwards rate constants whose values were picked for demonstration purposes. All species start at an initial





concentration of 0, dimensionless concentration, and have inlet boundary conditions of $[\text{NaCl}] = [\text{CaCO}_3] = 1$ and $[\text{Na}_2\text{CO}_3] = [\text{CaCl}_2] = 0$, again both dimensionless concentration. The equations and conditions have already been specified, and the simulation can be run by using the `run_compartment_w_rxn.py` script. Note that when you re-run this script it will finish much faster than the first time since the compartmental model is loaded from disk rather than having to be re-created each time.

To analyze the results, the equilibrium values for this reversible system are calculated as follows:

$$k_f[\text{NaCl}]^2[\text{CaCO}_3] = k_r[\text{Na}_2\text{CO}_3][\text{CaCl}_2]$$
$$\frac{k_f}{k_r} = \frac{[\text{Na}_2\text{CO}_3][\text{CaCl}_2]}{[\text{NaCl}]^2[\text{CaCO}_3]}$$
$$\frac{5 \times 10^{-2}}{2} = \frac{(x)(x)}{(1-2x)^2(1-x)}$$
$$x \approx 0.1147$$

where `x` is the change in CaCO_3, in dimensionless units.

The expected equilibrium concentrations for the four species are: $[\text{NaCl}] = 0.7706$, $[\text{CaCO}_3] = 0.8853$, $[\text{Na}_2\text{CO}_3] = 0.1147$, and $[\text{CaCl}_2] = 0.1147$. Based on the figures below, and from opening up the results, it can be seen that these steady state values are obtained at the outlet of the reactor.

To output the `ParaView` visualizations, change `output_VTK` to `True` in the `CONFIG` file and re-run the simulation.