

Turbulent Dissolution

Numerical Simulation of Multiphase Particle
Dissolution with Focus on Lithium Recovery from
Waste Batteries

ME55035: MSc Thesis
Amedeo Passero

Turbulent Dissolution

Numerical Simulation of Multiphase Particle Dissolution with Focus on Lithium Recovery from Waste Batteries

by

Amedeo Passero

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Wednesday September 24, 2025 at 10:00 AM.

Student number:	5850975		
Project duration:	December 9, 2024 – September 24, 2025		
Thesis committee:	Prof. dr. ir. L. Botto,	TU Delft, supervisor	
	Prof. dr. ir. S. Abrahams,	TU Delft, External Committee	
	Prof. dr. ir. W. Hogendoorn,	TU Delft, Department Committee	

Cover: DNS Turbulence velocity field from JHTDB.

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

The efficiency of hydrometallurgical leaching of cathode materials, a critical step in lithium-ion battery recycling, is often limited by the complex and poorly studied interaction between turbulent transport and multi-step reaction kinetics. To study this problem, this work uses a computationally efficient hybrid Eulerian-Lagrangian model. In this model, individual particles are simulated. Each particle evolution is described by a detailed Shrinking Core Model (SCM). The particles move through a pre-computed, high-fidelity velocity field from a Direct Numerical Simulation (DNS). This method separates the calculation of the particle reaction from the fluid dynamics simulation, which, computationally is the most expensive part. This separation makes it possible to do many parametric studies that would otherwise be too slow, which makes the model a powerful tool for process analysis.

The simulations show a clear non-monotonic performance penalty caused by turbulence. The global reaction rate is lowest at a resonant Kolmogorov-based Stokes number of $St_k \approx 0.23$. At this condition, preferential concentration is the strongest. This causes strong local clustering and reactant starvation inside the dense particle filaments. This resonant condition is the basis for a predictive engineering model for mixer design. The model defines a "clustering risk" zone for critical particle sizes as a function of mixer power and geometry. The analysis also shows internal kinetic limits, like the formation of a product layer, that if not quickly dissolved, inhibits the overall performance.

The main conclusion is that making the process faster by increasing mixing is not always as productive as expected. It is limited by a resonant clustering penalty. This challenges the common engineering idea that more mixing energy is always good for dissolution. Finally, this work gives a physics-based model to help with practical industrial problems. These problems include reactor design, process scale-up, and optimization of solids loading. This optimization must balance throughput with the performance reduction from particle clustering.

Contents

1	Introduction	1
1.1	Context and Motivation	1
1.2	Battery Recycling Processes and Hydrometallurgy	3
1.3	Problem Statement	4
1.4	Research Objectives and Questions	5
1.5	Scope and Limitations	6
1.6	Thesis Outline	7
2	Theoretical Foundations & Literature Review	9
2.1	Mass Transfer Mechanisms	9
2.1.1	Mass Transfer by Diffusion	9
2.1.2	Mass Transfer by Convection	10
2.1.3	The Mass Transfer Coefficient and Sherwood Number	10
2.1.4	Sherwood Number Correlations	11
2.1.5	Governing Equations for Particle Dissolution	12
2.1.6	Governing Equation for Bulk Concentration	13
2.2	Fundamentals of Heterogeneous Reaction Kinetics	14
2.2.1	Reaction Order	14
2.2.2	Temperature Dependence of Reaction Rates	14
2.3	The Shrinking Core Model for Fluid-Solid Reactions	14
2.3.1	Conceptual Framework and Assumptions	15
2.3.2	Governing Equations and Rate-Limiting Steps	15
2.4	The Role of Turbulence in Particle Dissolution	16
2.4.1	Turbulence Scales and the Energy Cascade	16
2.4.2	Homogeneous vs. Inhomogeneous Turbulence	17
2.4.3	Particle-Turbulence Interactions	18
2.5	Review of Numerical Modeling and Experimental Studies	21
2.5.1	Turbulence Modeling Approaches	21
2.5.2	Review of Simulating Multiphase Systems	22
2.5.3	Salt Dissolution Studies	25
2.5.4	Lithium Dissolution / Black Mass Leaching Studies	26
2.6	Motivation for the Research	26
2.6.1	Impact of Preferential Concentration on Reaction Kinetics	27
2.6.2	Deconstructing Internal Kinetic Mechanisms	27
3	A Unified Model of Dissolution	29
3.1	Problem Formulation and Objectives	29
3.2	Core Model Assumptions	29
3.3	Analysis of Serial Mass Transfer Resistances	30
3.4	The Unified Resistance Equation	31
3.5	Limiting Cases of the Unified Model	31
3.5.1	Film Resistance Limit (MassTransferOnly Model)	31
3.5.2	Coupled Diffusion-Reaction Limit (ShrinkingCoreModel)	32
3.6	Model Implementation Strategy	32
4	Numerical Framework and Verification	33
4.1	Overall Modeling Framework	33
4.1.1	A Eulerian-Lagrangian Method	33
4.1.2	Model Implementation and Choices	34
4.1.3	The Computational Workflow	34

4.1.4	Dual-Mode Simulation Capability	35
4.2	The Eulerian Concentration Field Solver	35
4.3	The Lagrangian Particle Tracking (LPT) Model	37
4.3.1	Model Simplification and Justification	37
4.3.2	Numerical Stiffness of Particle Equations	38
4.4	Particle-Fluid Coupling Strategy	39
4.4.1	One-Way Momentum Coupling	39
4.4.2	Two-Way Scalar Coupling	39
4.4.3	Mass Conservation via the Potential-Limit Algorithm	40
4.5	Particle Internal State Evolution	40
4.5.1	The Diffusion-Controlled Solver (MassTransferOnly Mode)	41
4.5.2	The Reaction-Controlled Solver (ShrinkingCoreModel Mode)	41
4.6	Model Verification	43
4.6.1	Grid Independence Study	44
4.6.2	Time-Step Independence Study	45
5	Results I: The Salt Dissolution Case	49
5.1	Influence of Turbulence on Mass Transfer	49
5.1.1	Parametric Study Setup	49
5.1.2	Fundamental Transport Regimes	50
5.1.3	Deconstructing the Ranz-Marshall Correlation	52
5.1.4	Validation Against Scaling Laws	52
5.1.5	Emergence of Polydispersity from Turbulent Heterogeneity	55
5.1.6	The Role of Particle Inertia and Stokes Number	55
5.1.7	Quantifying the Inertial Transition Threshold	57
5.1.8	Defining the Physical Limits of Process Intensification	57
6	Results II: LiCoO₂ Leaching	61
6.1	The Impact of Preferential Concentration on Reaction Rate	61
6.1.1	Decoupling Inertia and Transport	62
6.1.2	The Roles of Particle Inertia and Turbulent Transport	64
6.1.3	Evidence of Localized Reactant Starvation	68
6.1.4	Application to Reactor Design	69
6.2	Analysis of Leaching Kinetics	73
6.2.1	Parametric Analysis of Reaction Kinetics	73
6.2.2	Initial Kinetic Hierarchy and Passivation Dynamics	73
6.2.3	Kinetic Analysis of the k_1 Reaction Pathway	74
6.2.4	The Dual Role of Hydrogen Peroxide	75
6.2.5	Synthesis and Engineering Implications	76
7	Conclusions and Recommendations	79
7.1	Summary of Key Scientific Findings	79
7.1.1	Validation of Transport Physics (RQ1)	79
7.1.2	The Competing Effects of Turbulence on System Performance (RQ2)	79
7.1.3	Deconstructing the Leaching Kinetic Architecture (RQ3)	80
7.2	Implications for Industrial Process Design	80
7.2.1	Defining the Optimal Energy Input for Mixing	81
7.2.2	A Predictive Framework for Clustering Risk and Scale-Up	81
7.2.3	The Optimal Reducing Agent	81
7.3	Limitations of the Current Study	82
7.4	Recommendations for Future Work	82
7.4.1	Targeted Experimental Validation	82
7.4.2	Enhancing Fidelity with Dynamic and Higher Resolution Turbulence	83
7.4.3	Extension to Three-Dimensional Simulations	83
7.4.4	Adding Realistic Particle Properties	83
7.4.5	Development of a Non-Isothermal Model	83

A	Shrinking Core Model Rate Equations	i
A.1	Derivation of the Unified Rate Equation	i
A.2	Derivation of the Core Shrinkage Rate Equation	ii
B	Derivation of the Quiescent Concentration Profile and $Sh=2$ Limit	v
B.1	Governing Equation and Assumptions	v
B.2	Solution in Spherical Coordinates	v
B.3	Application of Boundary Conditions	vi
B.4	Derivation of the Sherwood Number Limit ($Sh = 2$)	vi
C	Numerical Solver for the <code>MassTransferOnly</code> Mode	ix
D	Derivation of the SCM Core Surface Concentration Polynomial	xi
D.1	The Physical Principle: Quasi-Steady-State Flux Balance	xi
D.2	Formulation of the Diffusive Flux	xi
D.3	Formulation of the Reactive Flux	xii
D.4	Algebraic Rearrangement to Polynomial Form	xii
E	Derivation of the Analytical Jacobian for the SCM	xiii
E.1	System Definition	xiii
E.2	Jacobian Structure	xiii
E.3	The Chain Rule and Implicit Differentiation	xiii
E.4	Derivation of Jacobian Components	xiv
E.4.1	Derivatives of Reaction Rates	xiv
E.4.2	Assembly of Final Jacobian Elements	xiv
E.5	Special Case: The No-Crust Limit	xiv

List of Figures

1.1	Battery and selected component production by country or region in the Stated Policies Scenario (STEPS) and Announced Pledges Scenario (APS), 2023-2035 [6].	2
1.2	Projected Evolution of Battery Chemistry Composition for EVs and Energy Storage Systems (ESS), 2025-2040 [7].	2
1.3	Recycling Opex Comparison for NMC 811 and LFP Batteries in Europe vs. China. . . .	3
1.4	Hydrometallurgy Process for Recycling Lithium-Ion-Batteries Showing Options for Chemical Precipitation, Solvent Extraction, and Selective Metal Adsorption Beds [11]. .	4
1.5	Overview of the numerical modeling framework and its application to different dissolution phenomena.	5
2.1	Fluid streamlines around a dissolving particle. Convection increases mass transfer by carrying solute away from the surface and bringing fresh solvent to it.	10
2.2	Plot of several Sherwood number correlations versus the particle Reynolds number. The graph shows the results for different Schmidt numbers and the valid range for each correlation.	12
2.3	A diagram that shows the connection between particle mass loss and solute release into the fluid. The dissolution rate changes the particle size. The released solute changes the bulk concentration. This new concentration then affects the dissolution rate.	13
2.4	Schematic of the Shrinking Core Model, showing the unreacted core, product layer, and fluid film. It includes concentration profiles for different rate-limiting steps: external mass transfer control, product layer diffusion control, and surface reaction control. From [13, Fig. 25.4].	15
2.5	A schematic diagram of the energy cascade at very high Reynolds number, showing the various length scales and ranges, and the processes of production, transfer, and dissipation of turbulent kinetic energy. From [28, Fig. 6.2].	17
2.6	The model spectrum normalized by the Kolmogorov scales. This log-log plot shows the energy spectrum $E(\kappa)$ versus wavenumber κ , clearly showing the inertial subrange (Kolmogorov's $-5/3$ law) and the exponential decay in the dissipation range. From [28, Fig. 6.13].	17
2.7	Comparison of typical flow fields for different turbulence categories. (a) shows the uniform statistical properties of isotropic turbulence, while (b) shows the spatial variations characteristic of impeller-driven flows.	18
2.8	Conceptual diagram showing particle response cases in turbulent flow. The cases depend on the ratio of particle response time to eddy turnover time (or Stokes number), which leads to different behaviors from passive tracers to heavy bullets.	20
2.9	Snapshots of particle number density in homogeneous isotropic turbulence for different Stokes numbers (St). The images show the effect of preferential concentration: (a) $St = 1$, (b) $St = 0.3$, (c) $St = 0.1$, and (d) $St = 0.03$. Particle clustering decreases as the Stokes number gets smaller, which leads to a more uniform distribution for particles with less inertia. From [30, Figure 3].	20
4.1	Probability density function (PDF) of the particle Reynolds number (Re_p). The distribution is mostly at very low values. These values are many orders of magnitude smaller than the limit for Stokes flow ($Re_p \approx 1$), which is shown by the dashed red line. It confirms that even the maximum observed Re_p is still in the Stokes regime. This is strong numerical proof that using the Stokes drag law is correct.	37
4.2	A schematic of bilinear interpolation. The particle at (x_p, y_p) is inside a grid cell. The value at its position is calculated from the values at the four surrounding grid points. .	40
4.3	Grid convergence study showing the solution approaching a stable value.	45

4.4	Analysis of the simulation's dependence on the time step, Δt_{sim} . Subfigure (a) shows the convergence of the final particle mass. Subfigure (b) shows the solution error and its convergence rate.	46
5.1	Deconstruction of the Sherwood Number	51
5.2	Linear decay of surface area in the diffusion-dominated regime	53
5.3	Linear decay of mass dissolution rate in the convection-dominated regime	55
5.4	The particle size distribution (PSD) over time. The particles all start at the same size. Because of turbulence, they dissolve at different rates. This makes the distribution wider, and the average diameter gets smaller.	56
5.5	Mechanistic analysis of the convective Sherwood term	57
5.6	Quantification of the Inertial Transition Radius	58
5.7	Process efficiency as a function of turbulence intensity	58
5.8	Derivatives of the process efficiency ($E = 1/t_{90}$) show two important transition points for process design. One is for operational cost, and the other is for high-throughput design.	59
6.1	Effectiveness factor $\langle \alpha \rangle$ as a function of Stokes number St for different turbulence intensity regimes.	65
6.2	Plots from Haugen et al. (2018) [30] showing mass transfer characteristics.	65
6.3	Effectiveness factor $\langle \alpha \rangle$ versus Damköhler number Da at a high turbulence intensity ($u_{rms} \approx 0.88 \text{ m s}^{-1}$) for several Stokes numbers. Performance gets worse as St is reduced over the tested range	66
6.4	Time-averaged effectiveness factor $\langle \alpha \rangle$ as a function of turbulence intensity (U_{rms}) for a fixed system (constant $n_p = 8000$, $d_0 = 50 \mu\text{m}$, $\phi_s = 1.0 \times 10^{-3}$). The plot reveals a complex, non-linear but monotonically non-decreasing trend.	67
6.5	Snapshots of the particle distribution at three different times for a case with $St \approx 1$. An initial homogeneous distribution quickly changes into thin, filament-like structures because of preferential concentration. These structures are not static and line up with the Lagrangian Coherent Structures (LCS).	69
6.6	Comparison of (a) the particle distribution and (b) the reactant (H^+) concentration field at one time in the transport-limited regime. There is a strong anti-correlation. Regions with many particles are the same as regions with low reactant concentration. This is direct visual proof of the reactant starvation mechanism.	69
6.7	The effect of solids loading (ϕ_s) on system effectiveness for different particle inertia regimes (Stokes numbers). This plot shows the practical performance penalty of clustering. Note the clear difference in robustness between the high-inertia ($St \approx 9.645$) and resonant-inertia ($St \approx 0.010$) particles.	70
6.8	Critical clustering diameter in different power zones for a fixed tank volume (V), impeller diameter (D) and Power Number (N_p).	72
6.9	Parametric Sensitivity Analysis of Mixer Design.	72
6.10	Sensitivity of Li Extraction Efficiency in the Initial Kinetically-Controlled Regime	74
6.11	The initial extraction rate as a function of the k_1 multiplier. This confirms that the process starts in a kinetically-controlled regime.	75
6.12	Average moles of passivating Co_3O_4 crust per particle. Higher k_1 values cause faster and larger crust formation.	75
6.13	The Rate Decay Index, which measures the "choking" effect of passivation. A sharp drop shows the start of a diffusion-dominated regime.	75
6.14	Relative extraction efficiency at $t = 10 \text{ s}$. The plot shows clear diminishing returns for high k_1 multipliers.	75
6.15	Time evolution of Li extraction efficiency for different k_1 multipliers. Higher k_1 causes faster initial extraction but also an earlier slowdown.	76
6.16	Final Li extraction efficiency versus the absolute k_1 rate constant (log scale). The flat part shows a transport-limited maximum.	76
6.17	Analysis of Hydrogen Peroxide Sensitivity	76

List of Tables

2.1	Summary of Key Sherwood Number Correlations for Spherical Particles.	12
2.2	Comparison of Turbulence Modeling Approaches	22
4.1	Parameters for the LiCoO ₂ SCM from Cerrillo-Gonzalez et al. [50].	42
4.2	Results of the grid refinement study. The final particle mass converges as the grid resolution N_{conc} increases. The relative change is calculated between each grid and the next finer grid.	44
4.3	Results of the time-step refinement study. Convergence, defined as a relative change below 0.3%, was achieved at $\Delta t_{sim} = 0.004\,687$ s. The error is calculated with respect to the solution at the finest time step ($\Delta t_{finest} = 0.0024$ s).	46
5.1	Key physical, numerical, and computational parameters used for the turbulence intensity parametric study. These values were kept constant across all simulation runs.	50
6.1	Statistical characteristics of turbulence for the moderate intensity simulation ($u_{rms} \approx 0.012\text{ m s}^{-1}$).	66

Nomenclature

List of Symbols and Variables		
Symbol	Variable	Unit
A_p	Particle surface area	m^2
C	Concentration	kg/m^3
C_{bulk}	Bulk fluid concentration	kg/m^3
C_s, C_i	Saturation/Interfacial concentration	kg/m^3
D	Diffusion coefficient	m^2/s
D	Impeller diameter	m
D_e	Effective diffusivity	m^2/s
Da	Damköhler number	-
d_p	Particle diameter	m
D/Dt	Material derivative	$1/\text{s}$
E_a	Activation energy	J/mol
$\mathbf{F}_i$	Force vector	N
$\mathbf{F}_a$	Added mass force vector	N
$\mathbf{F}_d$	Drag force vector	N
$\mathbf{F}_h$	History (Basset) force vector	N
$\mathbf{F}_l$	Lift force vector	N
$\mathbf{g}$	Gravitational acceleration vector	m/s^2
h	Grid spacing	m
i, j, k	Grid indices / Component index	-
J	Diffusive mass flux	$\text{kg}/(\text{m}^2 \cdot \text{s})$
k_L	Mass transfer coefficient	m/s
k_s	Surface reaction rate constant	(varies)
L	Integral length scale	m
m_f	Mass of displaced fluid	kg
m_p	Particle mass	kg
$\dot{m}$	Mass dissolution rate	kg/s
n	Reaction order	-
n_p	Number of particles	-
N	Impeller rotational Speed	-
N_p	Power Number	-
P	Power	W
Pe	Péclet number	-
r	Reaction rate	mol/s
r, r_p	Particle radius	m
r_c	Unreacted core radius	m
R	Ideal gas constant	$\text{J}/(\text{mol} \cdot \text{K})$
Re_p	Particle Reynolds number	-
Sc	Concentration source/sink term	$\text{kg}/(\text{m}^3 \cdot \text{s})$
Sc	Schmidt number	-
Sh	Sherwood number	-
St	Stokes number (e.g., St_L, St_K)	-
t	Time	s
T	Temperature	K
$\mathbf{u}$	Fluid velocity vector	m/s
u' or u_{rms}	RMS velocity fluctuation	m/s
u_{rel}	Relative velocity (particle-fluid)	m/s
V	Volume (reactor)	m^3
V_p	Particle volume	m^3
$\mathbf{v}_p$	Particle velocity vector	m/s
$\mathbf{x}, \mathbf{x}_p$	Position vector	m
x, y, z	Spatial coordinates	m
Greek Symbols		
$\langle \alpha \rangle$	Effectiveness factor	-
δ	Boundary layer thickness	m
ε	Porosity	-
ε	Turbulent kinetic energy dissipation rate	m^2/s^3
η	Kolmogorov length scale	m
λ	Taylor microscale	m
μ	Fluid dynamic viscosity	$\text{Pa} \cdot \text{s}$
ν	Fluid kinematic viscosity	m^2/s
ρ	Fluid density	kg/m^3
ρ_p	Particle density	kg/m^3
τ	Characteristic timescale (e.g., $\tau_p, \tau_L, \tau_\eta$)	s
Φ_v	Particle volume fraction	-

List of Compounds	
Compound	Name
Al	Aluminum
CaCl ₂	Calcium Chloride
CO ₂	Carbon Dioxide
Co	Cobalt
Co ₃ O ₄	Cobalt(II,III) Oxide
Cu	Copper
H ₂	Hydrogen gas
H ₂ O ₂	Hydrogen Peroxide
H ₂ SO ₄	Sulfuric Acid
HCl	Hydrochloric Acid
HNO ₃	Nitric Acid
Li	Lithium
Li ₂ CO ₃	Lithium Carbonate
LiCoO ₂	Lithium Cobalt Oxide
Mn	Manganese
Na ₂ CO ₃	Sodium Carbonate
NaCl	Sodium Chloride
Ni	Nickel
NMC 811	Nickel-Manganese-Cobalt (8:1:1 ratio)
O ₂	Oxygen gas

List of Abbreviations	
Abbreviation	Definition
APS	Announced Pledges Scenario
CFD	Computational Fluid Dynamics
DEM	Discrete Element Method
DES	Detached Eddy Simulation
DNS	Direct Numerical Simulation
EE	Eulerian–Eulerian
EL	Eulerian-Lagrangian
EMDEs	Emerging Markets and Developing Economies
EOL	End-of-Life
ERT	Electrical Resistance Tomography
ESS	Energy Storage Systems
EV	Electric Vehicle
FBRM	Focused Beam Reflectance Measurement
IEA	International Energy Agency
JHTDB	Johns Hopkins Turbulence Database
LES	Large Eddy Simulation
LFP	Lithium-Iron-Phosphate
LIB	Lithium-Ion Battery
LMFP	Lithium-Manganese-Ferro-Phosphate
LPT	Lagrangian Particle Tracking
NMC	Nickel-Manganese-Cobalt
ODE	Ordinary Differential Equation
Opex	Operational Expenditure
PBM	Population Balance Model
PEPT	Positron Emission Particle Tracking
PIV	Particle Image Velocimetry
PLIF	Planar Laser-Induced Fluorescence
PSD	Particle Size Distribution
PSSA	Pseudo-Steady-State Approximation
RANS	Reynolds-Averaged Navier-Stokes
SCM	Shrinking Core Model
SEM	Scanning Electron Microscopy
SGS	Subgrid-Scale
SPM	Smoothed Profile Method
STEPS	Stated Policies Scenario
T&E	Transport & Environment
URANS	Unsteady Reynolds-Averaged Navier-Stokes
VOF	Volume of Fluid
XRD	X-ray Diffraction

1

Introduction

1.1. Context and Motivation


Battery recycling is becoming increasingly critical for the European Union due to the rapid growth of Electric Vehicles (EVs) and widespread uses of batteries in industrial machines and household appliances. In fact, modern life is largely reliant on products that run on batteries, and so the management of waste batteries in a green way is a concern. The global market of EVs is expected to grow substantially to approximately 130 million units by 2030. Thus, approximately 11 million tonnes of lithium-ion batteries (LIBs) will reach their end-of-life (EOL) stage around that time [1]. As a result of this increasing demand for LIBs, global demand for lithium (Li) is expected to grow substantially, from approximately 0.5 million tonnes in 2022 to a possible peak of 3.2 million tonnes by 2040 [2]. The rapid increase shows the critical importance of securing secondary supplies, particularly through recycling, to have a stable and sustainable supply of lithium. Recent technical reviews emphasize the designation of lithium and cobalt as critical raw materials and highlight that spent LIBs can represent a richer source than some natural ores [3, 4].

Europe's dependence on China for battery materials is a major strategic problem. A 2022 report from the IEA, *Global Supply Chains of EV Batteries*, shows this imbalance clearly [5]. The report shows that China controls over 75% of the global production capacity for battery cells. This concentrates the supply chain in one country. China also produces over 80% of cathode materials and over 90% of anode materials. Furthermore, China refines more than half of the lithium, cobalt, and graphite for EV batteries. Although Europe produces over 25% of the world's EVs, it participates relatively minimally in the overall supply chain of batteries. An exception is cobalt processing, where Europe has about 20% of the global market share.

However, recent studies indicate a shift towards a more diversified battery production, primarily driven by strategic investments and new industrial policies. As Figure 1.1 shows, the Stated Policies Scenario (STEPS) and the Announced Pledges Scenario (APS) present different paths for battery manufacturing across regions from 2023 to 2035 [6]. China remains the main producer in both cases. However, the APS scenario predicts a large increase in production from emerging markets and developing economies (EMDEs) outside of China, such as India, Southeast Asia, and North Africa. Several factors drive this shift, including government incentives, lower manufacturing costs, and access to key raw materials.

Battery recycling presents a unique chance to achieve industrial independence for Europe, at the same time resolving material scarcity and reducing environmental impacts. According to research by T&E, by 2030 materials recovered through battery recycling (lithium, nickel, manganese, and cobalt) will be able to meet a considerable percentage of the European market's demand [7]. Recovered materials are expected to achieve 14% of the lithium demand, 16% of the nickel demand, 17% of the manganese demand, and 25% of the cobalt demand, respectively, enabling the production of up to 2.4 million electric vehicles (EVs) per year by 2030, with predictions showing a rise to 15 million EVs by 2040.

To reduce this dependence, the European Union has introduced new regulations. The EU Batteries




IEA. CC BY 4.0.

Figure 1.1: Battery and selected component production by country or region in the Stated Policies Scenario (STEPS) and Announced Pledges Scenario (APS), 2023-2035 [6].

Regulation (Regulation (EU) 2023/1542), which became effective in 2023, sets mandatory targets for recycling efficiency and for the use of recycled content in new products [8]. The goal is to increase the use of secondary raw materials and to reduce dependence on primary ones. The regulation requires companies to have a minimum recycling efficiency, calculated by mass, of 65% by the end of 2025. This target then increases to 70% by the end of 2030.

Another complication is the continuous evolution of battery chemistries. The most common chemistries used today are Nickel-Manganese-Cobalt (NMC), but it is expected a shift to Lithium-Iron-Phosphate (LFP), and Lithium-Manganese-Ferro-Phosphate (LMFP) batteries in the long-term. The long-term projection is that by 2030 these chemistries will represent a 59% market share that will grow to 63% by 2040, as seen in Figure 1.2. This shift will require recycling technology evolution to increase lithium, iron, and manganese recovery rates from a more heterogeneous feedstock of batteries [7].



T&E

Figure 1.2: Projected Evolution of Battery Chemistry Composition for EVs and Energy Storage Systems (ESS), 2025-2040 [7].

An important part of solving this problem is understanding how extraction and production processes of lithium work. Lithium is found in hard rock deposits and in saline brines within salt flats. Currently, nearly two-thirds of worldwide lithium production is gained through hard rock extraction with spodumene being the dominant lithium content mineral. Nevertheless, with increasing uses of lower-quality sources, it's predicted that in the long run, lithium content in worldwide ores will reduce. Therefore, more raw materials will have to be processed to produce a corresponding level of lithium product [7].

The 2024 T&E report [7] states that, by 2030, battery recycling in Europe can save close to 0.2 million tonnes of lithium ore, while this number can increase up to 0.8 million tonnes by 2040. In addition, by 2040, battery recycling in Europe can notably cut down the need for opening 12 new mines all over the world, including four lithium mines, three nickel mines, four cobalt mines, and one manganese mine.

From an environmental point of view, recycling of batteries offers significant advantages over extraction and processing of raw materials. The carbon footprint of recycled content is considerably lower due to considerable emissions related to both processing steps and mines. For instance, recycling of NMC 811 battery packs in Europe using the average European power grid saves CO₂ emissions by 19 % compared to lithium hydroxide extracted using raw materials. Recycling of lithium in Europe can lead to a cut of up to 20 % in CO₂ emissions compared to lithium extracted in Australia and subsequently processed in China [7].

Despite its potential, scaling up commercial battery recycling faces several economic and technical barriers. Recycling costs in Europe remain higher than in China, averaging \$14/kWh for NMC811 cell packs, compared to \$11/kWh in China. The cost gap is even wider for LFP battery recycling. Fighting these challenges requires targeted industrial support, including policy incentives, investment in hydrometallurgical recovery technologies. The application of the recycling rules in the EU Battery Regulation is essential to ensure that end-of-life batteries remain within the European recycling ecosystem.




Figure 1.3: Recycling Opex Comparison for NMC 811 and LFP Batteries in Europe vs. China.

By adopting a comprehensive strategy that prioritizes material recovery, industrial investment, and the consistent application of regulations, Europe can unlock the full potential of battery recycling, securing a sustainable and resilient battery supply chain while reducing dependence on foreign materials.

1.2. Battery Recycling Processes and Hydrometallurgy

Hydrometallurgy is the type of recycling process that will be discussed in this study. After end-of-life or production scrap lithium-ion batteries are collected, they undergo mechanical shredding in order to obtain a powder called black mass, that has a high concentration of valuable metals [9, 10].

The first step of the process is acid leaching. In this step, an acid treats the black mass to extract the target metals: cobalt, nickel, manganese, and lithium. This is commonly done using inorganic acids such as sulfuric acid (H₂SO₄), hydrochloric acid (HCl), or nitric acid (HNO₃), although greener organic acids (e.g., citric, ascorbic acid) are also used [9]. To improve the process, a reducing agent, like hydrogen

peroxide (H_2O_2), is often added. This agent helps transform less soluble higher-valence metal oxides into more soluble ionic forms [10]. A schematic illustration of the traditional process is shown in Figure 1.4. After leaching, filtration is carried out to remove insoluble byproducts, mainly graphite and other inert compounds, producing a metal-rich solution.




Figure 1.4: Hydrometallurgy Process for Recycling Lithium-Ion-Batteries Showing Options for Chemical Precipitation, Solvent Extraction, and Selective Metal Adsorption Beds [11].

After the leaching step, the metals are selectively separated from the concentrated solution. Chemical precipitation is a common technique for this purpose. The method works by careful adjustment of the solution pH or by adding specific reagents. This allows for the selective precipitation of metals based on their different solubilities [3, 4, 9]. For example, impurities like aluminum (Al) or copper (Cu) are often removed first by changing the pH. Lithium is usually recovered later. A carbonate source, such as sodium carbonate (Na_2CO_3), is added to precipitate lithium carbonate (Li_2CO_3) [10]. Other techniques, like solvent extraction and ion exchange, offer high selectivity for individual metals. However, precipitation is still a fundamental method used to recover final products or remove impurities [9]. This thesis focuses on modeling the fluid dynamics and kinetics of the dissolution step.

1.3. Problem Statement

The main goal of this thesis is to build a numerical model able to describe the dissolution of black mass powder, specifically, fine lithium-cobalt (LiCoO_2) particles, within an acid, a key step in hydrometallurgical battery recycling. This work fills an important gap in the scientific literature. Many models exist for general solid-liquid dissolution. However, there are few studies about the specific interaction between turbulent flow and the complex kinetics of hydrometallurgical leaching. This gap is especially large for lithium recovery. The literature on numerical modeling of LiCoO_2 leaching is very limited. Most of the work comes from one research group (Cerrillo-Gonzalez et al.), but their models do not include the details of the fluid dynamics. Therefore, the main challenge is to create a model that connects these two fields. The model must couple detailed transport physics with realistic reaction kinetics to predict performance at the particle scale. The research starts by developing and validating a simplified particle dissolution model for salt particles in a turbulent aqueous flow. This first step, implemented in MATLAB [12], serves as a crucial verification of the core transport physics. Then, the validated model is extended to simulate the far more complex dissolution of LiCoO_2 particles, which corresponds to the industrial leaching stage.

A critical distinction defines this two-stage approach: while simple salt dissolution is primarily governed

by external mass transfer (a diffusion-convection problem), the leaching of lithium compounds is controlled by complex multi-step chemical reactions. This fundamental difference results in highly different characteristic process timescales. The main challenge in optimizing the hydrometallurgical process is understanding which is its rate-limiting steps. The macroscopic performance of a leaching reactor, such as overall yield and residence time, is not determined by bulk-averaged properties alone. Instead, it is a result of the complex transport and kinetic phenomena at the microscale of individual particles. As established in foundational chemical engineering principles [13], the overall rate of any heterogeneous process is dictated by its slowest step, be it external mass transfer, internal diffusion, or surface reaction.

The rationale for first modeling salt dissolution is to validate the core particle tracking and mass transfer sub-models in a well defined physical system before introducing the complexities of reaction-driven kinetics. The model uses a pre-resolved, high-fidelity turbulent flow field from a Direct Numerical Simulation (DNS) database (Johns Hopkins Turbulence Database [14]). In this flow, particle motion is tracked using a point-particle, one-way momentum coupling approach. Given that black mass particles are relatively small ($d_p \approx 0.1$ mm) and the suspensions are dilute, it is a well-justified assumption that their presence does not significantly alter the turbulent structures of the surrounding fluid.

Following the validation of the transport model, the framework is adapted to include the chemistry of acid leaching via the Shrinking Core Model (SCM). The SCM models the formation of a porous solid product layer, or crust, around the unreacted particle core. This crust introduces an additional, time dependent internal resistance to mass transfer, acting in series with external transport and surface reaction. This interplay of resistances raises fundamental questions about the process's controlling mechanisms. The numerical framework, conceptually shown in Figure 1.5, is thus designed to deconstruct these distinct physical regimes, from simple diffusion-convection to multi-step, multi-resistance reaction kinetics.




Figure 1.5: Overview of the numerical modeling framework and its application to different dissolution phenomena.

Decoupling the turbulence simulation (by using a DNS database) from the particle tracking and reaction modeling provides a computationally efficient and high-fidelity alternative to more conventional methods like fully coupled CFD-DEM simulations. This strategy is effective for systems where one-way momentum coupling is a valid approximation, allowing for a focused analysis of the coupled particle dynamics and reactive mass transfer, which are central to this thesis.

1.4. Research Objectives and Questions

As previously described, the primary goal of this research is to build and test a numerical framework to describe the coupled physics of particle dissolution in turbulent, reactive environments, with a specific focus on hydrometallurgical applications. The research follows a rigorous validate then investigate methodology. This approach proceeds logically through three research questions. The framework is first benchmarked against established transport physics (RQ1) to verify its fundamental accuracy. It is then applied to see how the overall performance is affected by particle clustering (RQ2). Finally, the model is used to analyze the complex internal reaction kinetics of the chemical system (RQ3).

The key research questions that guide this study are:

RQ1: *In a diffusion-controlled dissolution process, how does the heterogeneous nature of a turbulent flow control the statistical distribution of individual particle dissolution rates? Does the numerical model accurately capture these physics when validated against semi empirical correlations such as Ranz-Marshall Equation?*

This first question is the foundation of the entire study. Before investigating complex chemical

kinetics, the model's core transport algorithms must be validated. RQ1 therefore tests the basic accuracy of the numerical framework, checking if it properly reproduce the standard physics of convective mass transfer. A successful validation against well-established correlations provides a reliable baseline, confirming that the simulator accurately solves for the external transport phenomena that govern all subsequent, more complex scenarios.

RQ2: *How does the preferential concentration of particles, a key feature of turbulent multiphase flows, create localized regions of reactant depletion? And how does this transport limitation change the system's global reaction rate compared to a uniformly distributed particle field?*

With the transport physics for individual particles verified, the research moves to a global phenomena that emerge in multi-particle flows. In turbulent suspensions, inertial particles do not remain uniformly distributed. Instead they accumulate into dense, filamentary clusters through **preferential concentration**. This second question investigates the macroscopic consequences of this microscopic non uniformity. It directly asks if the self-organization of the particle field creates local transport limitations, a phenomenon known as **reactant starvation**, and measures the resulting effect on the system's overall reaction efficiency.

RQ3: *Given the uncertainties in the experimental parameters for the multi-step Shrinking Core Model, which kinetic pathways, such as direct core dissolution versus crust formation and reaction, have the biggest effect on the predicted lithium extraction efficiency and the overall process timescale?*

The final stage of the research moves from the study of *external* transport limitations, to the *internal* chemo-physical processes inside each particle. The acidic leaching of LiCoO_2 is governed by the **Shrinking Core Model (SCM)**, which includes the formation of a porous product layer that adds a significant internal resistance to mass transfer. This final question, therefore, analyses the multi-step reaction kinetics to identify which chemical pathways are the main factors controlling lithium extraction efficiency and the overall process timescale. This analysis is critical also to understand how the uncertainty of experimental measurements of the kinetic constants affects the overall process.

1.5. Scope and Limitations

To achieve the research objectives, the numerical framework uses a set of specific modeling choices and assumptions. These choices allow a focused analysis of certain physical phenomena by excluding others. The limits of this study are:

- **Modeling Approach:** The model uses a Lagrangian Particle Tracking (LPT) method with the point particle assumption. This means that the flow and concentration boundary layers around each particle are not resolved. This simplification is valid because the model focuses on the collective behavior of many particles. The small scale physics are included via closure models (e.g., Sherwood number correlations).
- **Fluid Dynamics:** The fluid flow is a static snapshot of Direct Numerical Simulation (DNS) data. This creates a "frozen flow" environment. This choice is made mainly for computational constraints.
- **Phase Coupling:** Momentum coupling is one-way, so particles do not affect fluid velocity. Scalar coupling is *two-way*. The Eulerian concentration field affects particle reaction rates. In turn, the reacting particles act as localized sources/sinks that change the concentration field. This two way coupling is important to study effects like local reactant starvation, which happens due to the interaction between particles and scalar field.
- **Computational Parcel Model:** The model uses a Computational Parcel Model, also known as the "Super Particle" approach[15]. This is done to simulate a large number of particles with reduced computational cost. Each Lagrangian point represents a group of N_A real particles. The inertia of the parcel is the same as one real particle, however, its effect as a source or sink in the concentration field is multiplied by the number of real particles in the parcel, N_A .

- **Particle Interactions:** Particle-particle collisions and short-range hydrodynamic interactions are neglected. This assumption means the model is for dilute or semi-dilute systems. In these flows, four-way coupling effects are not the dominant factor for particle behavior.
- **Dissolution Models:** The model describes the leaching of LiCoO_2 with an advanced Shrinking Core Model (SCM). This includes a product layer (crust) where porosity and effective diffusivity change over time. The main simplification is the assumption of a *spherically symmetric* crust. This choice helps to focus the study on the competition between external mass transport, diffusion inside the crust, and surface reaction.
- **Concentration Field:** The model maps mass sources from Lagrangian particles to the Eulerian grid. It uses a conservative, area-weighted method. This method makes sure that mass is conserved, however it does not resolve the concentration gradients in the small space right next to the particle surface.
- **Dimensionality:** Simulations are performed in a two-dimensional domain. This domain is a slice from a three-dimensional DNS dataset. Using a 2D domain significantly reduces computational cost. A known limitation is that this approach does not include 3D turbulent effects like vortex stretching, which could change how particle cluster.
- **Thermal Effects:** The system is assumed to be isothermal. The effects from heat of reactions and kinetics that depend on temperature are outside the scope of this work.
- **Particle Properties:** Particles are assumed to be perfect spheres with constant density. Only their size evolves over time. This choice separates the dissolution dynamics from the complex effects, like particle shapes or internal porosity.

These assumptions define the limits of the research. They are important for the interpretation of the results.

1.6. Thesis Outline

The structure of this thesis is as follows:

- **Chapter 1: Introduction.** This chapter gives the background and reason of this research. The main one is the need for sustainable battery recycling. The chapter defines the problem, the research questions and describe the scope and limitations of this work.
- **Chapter 2: Theoretical Foundations & Literature Review.** This chapter provides the theory needed for this work. It covers the basics of particle dissolution, the physics of turbulence in multiphase flows, and common models for heterogeneous reactions like the Shrinking Core Model. The chapter also reviews numerical and experimental studies to show research gaps this thesis wants to fill.
- **Chapter 3: A Unified Model of Dissolution.** This chapter defines the main model of the thesis. It uses an analogy to electrical resistances in series to brake down the dissolution process into three steps: external film transfer, internal pore diffusion, and surface reaction. These steps are combined in a single governing equation. This equation is the base used for the numerical code.
- **Chapter 4: Numerical Framework and Verification.** This chapter explains how the conceptual model is turned into a computational tool. It describes the hybrid Eulerian-Lagrangian structure, the numerical methods used for particle tracking and the concentration field, and the way the two phases are coupled. This chapter shows also the results of verification studies, like grid and time-step independence analyses, to prove the accuracy of the simulator.
- **Chapter 5: Results I: The Salt Dissolution Case.** This chapter shows the first scientific results. It focuses on the simple case of salt dissolution. The main goal is to validate the model's transport physics, which answers Research Question RQ1. The results are compared with theoretical scaling laws to show the model is accurate.

- **Chapter 6: Results II: LiCoO₂ Leaching.** This is the main results chapter. Here the validated model is used to study the more complex reaction-controlled leaching of LiCoO₂. It answers RQ2 by calculating the loss in performance caused by particle clustering, and shows that local reactant starvation is the cause of this loss. To answer RQ3, the chapter then analyses the leaching kinetics. This shows the competing reaction pathways and the important dual role of hydrogen peroxide as both kinetic driver and passivation suppressant.
- **Chapter 7: Conclusions and Recommendations.** This final chapter summarizes the key scientific results of the study. It gives final answers to research questions base on the results obtained. The chapter explains the practical meaning of these results for industrial process design. It ends by discussing the limitations of the study and suggests specific paths for future research.
- **Appendices.** The appendices contain detailed mathematical derivations and supplementary information that support the main body of the thesis.

2

Theoretical Foundations & Literature Review

2.1. Mass Transfer Mechanisms

The dissolution of a solid particle in a fluid is the transport of solute from the particle surface to the bulk fluid. This transport happens by two main mechanisms: **molecular diffusion**, driven by a concentration gradient, and **convection**, caused by the relative motion between particle and fluid.

2.1.1. Mass Transfer by Diffusion

Mass transfer by **diffusion** is the transport of a solute driven by a concentration gradient. This process happens because of random thermal motion of molecules. The diffusive flux is described by Fick's First Law [16], which is analogous to Fourier's law for heat conduction:

$$\mathbf{J} = -D\nabla C \quad (2.1)$$

where $\mathbf{J}$ is the diffusive mass flux vector ($\text{mol m}^{-2} \text{s}^{-1}$), D is the molecular diffusion coefficient ($\text{m}^2 \text{s}^{-1}$), and ∇C is the concentration gradient (mol m^{-4}). For transport in one dimension, this simplifies to $J = -D(\partial C / \partial x)$.

During dissolution, a concentration boundary layer forms around the particle. At the particle surface ($r = R_p$), the solute concentration is at its saturation value, C_s . The concentration decreases to the bulk fluid value, C_{bulk} , far from the particle. For a spherical particle in an infinite, quiescent medium, the steady-state diffusion equation gives the exact concentration profile:

$$C(r) = C_{bulk} + (C_s - C_{bulk}) \frac{R_p}{r} \quad (2.2)$$

where R_p is the particle radius. The diffusive flux at the particle surface depends on the gradient of this profile at $r = R_p$. This analysis gives a theoretical mass transfer coefficient that corresponds to a Sherwood number of 2. This value is the minimum possible mass transfer rate for a single sphere in a still medium.

A simpler model, the **film theory**, approximates the concentration profile as a linear gradient across a thin, stagnant film of effective thickness δ . In this model, the diffusive flux is given by:

$$J \approx D \frac{C_s - C_{bulk}}{\delta} \quad (2.3)$$

This model shows a simple connection between the boundary layer thickness and the mass transfer rate. Equation (2.2) is the exact solution for the purely diffusive case. However, the film model is a useful tool to understand how fluid motion affects mass transfer. In most engineering systems, fluid motion adds convective transport. Convection is important because it changes the mass transfer rate by making the boundary layer thinner.

2.1.2. Mass Transfer by Convection

Fluid motion causes **convective mass transfer**. This process transports solute away from the particle surface. This steepens the concentration gradient and increases the dissolution rate compared to pure diffusion. Figure 2.1 shows this process. The fluid flow brings fresh solvent to the surface and carries the dissolved solute into the bulk fluid.




Figure 2.1: Fluid streamlines around a dissolving particle. Convection increases mass transfer by carrying solute away from the surface and bringing fresh solvent to it.

The complete velocity and concentration fields are complex. For this reason the effect of convection is often calculated by modeling the resistance to mass transfer inside a thin concentration boundary layer at the particle surface. This idea is based on the boundary layer theories of Prandtl and the Nernst diffusion layer model [17, 18]. The main idea is that fluid flow, particularly turbulent flow, reduces the effective thickness of this boundary layer, δ_{eff} . As shown in Equation (2.3), a smaller δ_{eff} gives a larger mass transfer flux. Turbulent eddies reach close to the particle surface, which increases mixing and makes this boundary layer thinner. This direct relationship between fluid dynamics and the boundary layer thickness is the main way that convection controls dissolution rates. The combined effects of diffusion and convection are usually described by an empirical mass transfer coefficient.

2.1.3. The Mass Transfer Coefficient and Sherwood Number

The combined effect of diffusion and convection is described using the **mass transfer coefficient**, k_L (m s^{-1}), it relates the mass flux J from the particle surface to the concentration driving force:

$$J = k_L(C_s - C_{bulk}) \quad (2.4)$$

Therefore, the mass dissolution rate, $\dot{m}$, of a single particle with surface area A_p is:

$$\dot{m} = k_L A_p (C_s - C_{bulk}) \quad (2.5)$$

The coefficient k_L is not a fixed property of the material. Its value depends on the particle diameter (d_p), the relative velocity between the particle and fluid (u_{rel}), and the fluid properties: density (ρ_f), viscosity (μ_f), and solute diffusivity (D).

These relationships are described using dimensionless numbers. The **Sherwood number** (Sh) is the dimensionless mass transfer coefficient. It is analogous to the Nusselt number (Nu) in heat transfer and

is defined as:

$$Sh = \frac{k_L d_p}{D} \quad (2.6)$$

It is the ratio of the total mass transfer rate to the rate that happens by only molecular diffusion. In practice, Sh is derived from empirical correlations. These correlations connect the Sherwood number to other dimensionless numbers that describe the flow and the fluid properties.

2.1.4. Sherwood Number Correlations

The Sherwood number is usually calculated from empirical correlations. These correlations connect it to two important dimensionless groups: the **particle Reynolds number** (Re_p) and the **Schmidt number** (Sc). The Reynolds number is the ratio of inertial forces to viscous forces in the flow around the particle. The Schmidt number is the ratio of momentum diffusivity to mass diffusivity in the fluid. They are defined as:

$$Re_p = \frac{\rho_f U_{rel} d_p}{\mu_f} = \frac{U_{rel} d_p}{\nu_f} \quad \text{and} \quad Sc = \frac{\nu_f}{D} = \frac{\mu_f}{\rho_f D} \quad (2.7)$$

where $U_{rel} = |\mathbf{u} - \mathbf{v}_p|$ is the magnitude of the relative velocity between the fluid ($\mathbf{u}$) and the particle ($\mathbf{v}_p$).

These correlations for a single sphere often have a general form that comes from boundary layer theory:

$$Sh = C_1 + C_2 Re_p^m Sc^n \quad (2.8)$$

where C_1 , C_2 , m , and n are constants from experiments. The constant C_1 is the limit for pure diffusion (typically $C_1 = 2$ for a sphere). The exponents m and n describe the physics of the transport. For example, $m = 1/2$ and $n = 1/3$ are typical for mass transfer in a laminar boundary layer.

Some common correlations for different flow regimes are:

- **Pure Diffusion Limit** ($Re_p = 0$): For a spherical particle in a quiescent medium, the theoretical limit from the steady-state concentration profile is:

$$Sh = 2 \quad (2.9)$$

This is a basic result in mass transfer theory [19].

- **Creeping Flow Regime** ($Re_p \ll 1$): For low Reynolds numbers, a correlation by Bird et al. [16] can be used for $Re_p \leq 0.1$:

$$Sh = 2 + 0.991 Re_p^{1/3} Sc^{1/3}, \quad \text{for } Re_p \leq 0.1 \quad (2.10)$$

Another form by Clift et al. [19] uses the Péclet number ($Pe = Re_p \cdot Sc$):

$$Sh = 1 + (1 + Pe)^{1/3} \quad (2.11)$$

- **Moderate Reynolds Numbers**: For flows where a laminar boundary layer forms, the Frössling equation [20] is an important correlation:

$$Sh = 2 + 0.552 Re_p^{1/2} Sc^{1/3} \quad (2.12)$$

A very common correlation for this regime is the Ranz-Marshall equation (1952) [21]:


$$Sh = 2 + 0.6 Re_p^{1/2} Sc^{1/3}, \quad \text{for } Re_p \leq 150, \quad 0.5 \leq Sc \leq 2 \quad (2.13)$$

- **Higher Reynolds Numbers / Turbulent Effects**: For flows where the boundary layer around the particle is transitional or turbulent, Clift et al. [19] give this correlation:

$$Sh = 1 + 0.724 Re_p^{0.48} Sc^{1/3}, \quad \text{for } 100 < Re_p \leq 2000, \quad Sc \geq 200 \quad (2.14)$$

Table 2.1: Summary of Key Sherwood Number Correlations for Spherical Particles.

Correlation	Source	Applicable Range
$Sh = 2$	Theoretical limit	$Re_p = 0$ (stagnant)
$Sh = 2 + 0.991Re_p^{1/3}Sc^{1/3}$	Bird et al. (2002) [16]	$Re_p \leq 0.1$
$Sh = 2 + 0.552Re_p^{1/2}Sc^{1/3}$	Frössling (1938) [20]	$Re_p \approx 50$
$Sh = 2 + 0.6Re_p^{1/2}Sc^{1/3}$	Ranz-Marshall (1952) [21]	$Re_p \leq 150$
$Sh = 1 + 0.724Re_p^{0.48}Sc^{1/3}$	Clift et al. (1978) [19]	$100 \leq Re_p \leq 2000$

**Figure 2.2:** Plot of several Sherwood number correlations versus the particle Reynolds number. The graph shows the results for different Schmidt numbers and the valid range for each correlation.

For the particle sizes (0.1 mm to 1 mm) and relative velocities (0.1 m s^{-1} to 1 m s^{-1}) used in this work, the particle Reynolds number is expected to be in the range of $0 \leq Re_p \leq 100$. In this regime, the correlations developed for creeping and laminar flow are the most useful, especially the Ranz-Marshall equation. Table 2.1 summarizes these important correlations, and Figure 2.2 shows how they change with the Reynolds number.

The standard Sherwood number correlations, which use the mean relative velocity, do not work well for very turbulent flows. In these flows, the interaction between particles and eddies control the mass transfer, especially for particle smaller than the Kolmogorov length scale (η_K). For this reason, newer correlations have been developed that include direct measures of turbulence, for example the turbulent kinetic energy dissipation rate, ε . These models describe the increased mass transfer from small-scale eddies, an effect that Re_p alone does not capture [22, 23]. Using these advanced models is not part of this thesis. However, they are important for systems with strong turbulence and small particles.

2.1.5. Governing Equations for Particle Dissolution

The rate of mass loss, $\dot{m}$, from a single spherical particle depend on the mass transfer coefficient and the concentration driving force. This relationship is described by the Noyes-Whitney equation [24]. Substituting the definitions of k_L from Equation (2.6) and the particle surface area ($A_p = \pi d_p^2$) into Equation (2.5) gives the mass loss rate in terms of the Sherwood number:

$$\dot{m} = \left(\frac{Sh \cdot D}{d_p} \right) (\pi d_p^2) (C_s - C_{bulk}) = \pi d_p Sh D (C_s - C_{bulk}) \quad (2.15)$$

The mass loss rate is related to the change in particle diameter, d_p , through the particle's mass, $m = \rho_p(\pi/6)d_p^3$. The time derivative of the mass is $\dot{m} = \rho_p(\pi/2)d_p^2(d(d_p)/dt)$. Setting this equal to Equation (2.15) and solving for the rate of change of diameter gives:

$$\frac{d(d_p)}{dt} = -\frac{2ShD(C_s - C_{bulk})}{\rho_p d_p} \quad (2.16)$$

where the negative sign shows that the size decreases. For better numerical stability in computational models, it is often better to track the square of the diameter, d_p^2 . This is because in this way the term $1/d_p$ is no more present on the right hand side. That term could in fact give numerical instabilities for dissolution processes where $d_p \rightarrow 0$. The rate of change of d_p^2 is found from Equation (2.16):

$$\frac{d(d_p^2)}{dt} = 2d_p \frac{d(d_p)}{dt} = -\frac{4ShD(C_s - C_{bulk})}{\rho_p} \quad (2.17)$$

This final equation is often called the " d^2 -law" for dissolution. It shows that the surface area of the particle decreases at a rate that does not depend on its current size, if Sh and the concentration difference are constant.

2.1.6. Governing Equation for Bulk Concentration

The change of the solute concentration field, $C(\mathbf{x}, t)$, in the bulk fluid is described by the advection-diffusion equation. This equation is a mass balance for the solution. The general form of the equation is:

$$\frac{\partial C}{\partial t} + \nabla \cdot (\mathbf{u}C) = \nabla \cdot (D\nabla C) + S_c \quad (2.18)$$

where $\mathbf{u}$ is the fluid velocity field and S_c is a source term for the mass released from the dissolving particles. For an incompressible fluid ($\nabla \cdot \mathbf{u} = 0$) and if the diffusivity is constant, the equation simplifies to:

$$\frac{\partial C}{\partial t} + \mathbf{u} \cdot \nabla C = D\nabla^2 C + S_c \quad (2.19)$$

The term $\mathbf{u} \cdot \nabla C$ is the convective transport by the fluid, and $D\nabla^2 C$ is the molecular diffusion.

In a fully-resolved simulation, this equation is solved everywhere in the fluid, with boundary conditions at each particle's surface. But, this method is computationally too expensive for systems with many particles. The Eulerian-Lagrangian framework used in this thesis avoids this cost. It models particle dissolution using the mass transfer correlations in Section 2.1.4 and coupling it to the fluid phase through the source term, S_c , in Equation (2.19).




Figure 2.3: A diagram that shows the connection between particle mass loss and solute release into the fluid. The dissolution rate changes the particle size. The released solute changes the bulk concentration. This new concentration then affects the dissolution rate.

Model Assumptions and Extensions

This analysis uses two important assumptions: that particles are perfect spheres, and the suspension is dilute, so interactions between particles can be ignored. In many industrial applications, these assumptions may not hold. For non spherical particles, the model needs to use shape factors or more complex surface models [25]. For dense suspensions, interactions between nearby particles can change the local mass transfer rates. Therefore, in such cases, the model has to include these effects, examples are the model used in CFD-DEM simulations [26]. The simple model for spherical particles in a dilute suspension is a good starting point. But these other models are important to correctly simulate systems with complex particle shapes or with many particles.

2.2. Fundamentals of Heterogeneous Reaction Kinetics

Chemical reactions can be homogeneous or heterogeneous. A homogeneous reaction happens in a single phase. A **heterogeneous reaction** happens at the interface between two or more phases, for example, a solid-fluid interface [13]. In fluid-solid systems, the total rate of the process depends on two things: the rate of the chemical reaction at the surface, and the rate of mass transfer of reactants and products to and from that surface.

The rate of a heterogeneous reaction can be defined in different ways. For example it can be per unit interfacial area, per unit mass of the solid, or per unit volume of the fluid. It is important to use the same basis to create correct kinetics models and to compare reaction rate data.

2.2.1. Reaction Order

For a reaction with a species of concentration C , the rate, r , is often described with a power-law equation:

$$r = kC^n \quad (2.20)$$

The exponent n is the **order of reaction** for that species. It is important to know that the reaction rate is found from experiments. For complex or multi-step reactions, it's important to notice that has no necessary relationship to the stoichiometric coefficient of the species in the balanced chemical equation [13, 27]. The value of n can be an integer, a fraction, or even a function of concentration. It helps to understand the mechanism of the slowest step in the overall reaction.

2.2.2. Temperature Dependence of Reaction Rates

The effect of temperature on the reaction rate constant, k , is described by the Arrhenius equation [13]:

$$k = A \exp\left(-\frac{E_a}{RT}\right) \quad (2.21)$$

where A is the pre-exponential factor, E_a is the activation energy (J mol^{-1}), R is the universal gas constant, and T is the absolute temperature. The activation energy is the energy barrier for the reaction. A high activation energy means the reaction rate is very sensitive to temperature.

This temperature dependence is a very useful tool to find the rate-limiting step in heterogeneous systems. The physical processes of mass transfer and chemical reaction have very different activation energies:

- **Mass Transfer Control:** Processes controlled by film or pore diffusion depend only a little on temperature. They have low activation energies.
- **Reaction Control:** Processes controlled by the intrinsic surface chemical reaction depend a lot on temperature. They have high activation energies.

Therefore, by measuring the overall reaction rate at different temperatures, it is possible to find the controlling process. If the rate changes a lot with temperature, the system is controlled by the reaction kinetics. If the rate changes only a little, the system is limited by mass transfer [13].

2.3. The Shrinking Core Model for Fluid-Solid Reactions

The **Shrinking Core Model (SCM)** is a common model to describe fluid-solid reactions where a product layer, often called 'ash', forms around a shrinking core of the unreacted solid [13]. This model is especially important for processes like the leaching of metal oxides from composite particles.

2.3.1. Conceptual Framework and Assumptions

The SCM models the total reaction as a series of three steps that happen in series:

1. **Film Diffusion:** Transport of the fluid reactant from the bulk fluid, through the mass transfer boundary layer, to the outer surface of the particle.
2. **Ash Layer Diffusion:** Diffusion of the fluid reactant through the porous product layer to the surface of the unreacted core.
3. **Surface Reaction:** The chemical reaction between the fluid reactant and the solid at the surface of the unreacted core.

The slowest of these three steps controls the overall rate of the process. This model is based on several important assumptions:

- The solid particles are spherical and have the same initial size.
- The reaction creates a clear, shrinking spherical core of unreacted solid with radius $r_c(t)$.
- The properties of the product layer, like its effective diffusivity (D_e), are uniform.
- A pseudo-steady-state approximation (PSSA) is assumed for the concentration profile of the fluid reactant in the film and product layers.




Figure 2.4: Schematic of the Shrinking Core Model, showing the unreacted core, product layer, and fluid film. It includes concentration profiles for different rate-limiting steps: external mass transfer control, product layer diffusion control, and surface reaction control. From [13, Fig. 25.4].

2.3.2. Governing Equations and Rate-Limiting Steps

With the pseudo-steady-state assumption, the molar rate of the fluid reactant is constant through each of the three layers in series. The classical Shrinking Core Model describes this process with an analogy to electrical resistances in series. The total rate depends on the sum of the separate resistances for mass transfer and reaction [13]. This relationship can be expressed with the following idea:

$$\text{Overall Rate} = \frac{\text{Total Concentration Driving Force}}{\sum \text{Resistances}} \quad (2.22)$$

Or, the formal equation for the molar rate of reactant consumption, $\dot{N}_A$, is:

$$\dot{N}_A = \frac{C_{Ab}}{R_{\text{film}} + R_{\text{product layer}} + R_{\text{reaction}}} \quad (2.23)$$

where C_{Ab} is the bulk reactant concentration, which is the total driving force. The terms in the denominator are the resistances of the external fluid film, the internal product layer, and the surface reaction. The full mathematical forms for each resistance term will be derived in Chapter 3.

Therefore, the slowest step, which is the one with highest resistance, determines the overall process rate. This idea leads to three distinct rate-limiting regimes:

- **Film Diffusion Control:** This happens when the resistance of the outer fluid film is dominant. In this case, the chemical reaction is fast. The overall process is limited by the rate of reactant transport from the bulk fluid to the particle's outer surface.
- **Product Layer Diffusion Control:** This happens when the resistance of the porous product layer is dominant. The process is limited by how fast the reactant can diffuse through this layer to reach the unreacted core.
- **Surface Reaction Control:** This happens when the chemical reaction is slow and its resistance is dominant. The process is limited by the rate of the chemical reaction at the surface of the unreacted core.

An important result of this model is that each controlling regime has a specific mathematical relationship. This relationship connects the time for complete particle conversion, τ , and the initial particle radius, r_p . By integrating the main differential equation and assuming one resistance is dominant, the following scaling laws can be found [13]:

- For **Film Diffusion Control**, the time for complete conversion is proportional to the initial particle radius: $\tau \propto r_p$.
- For **Product Layer Diffusion Control**, the time is proportional to the square of the initial radius: $\tau \propto r_p^2$.
- For **Surface Reaction Control**, the time is again proportional to the initial particle radius: $\tau \propto r_p$.

This model gives a good experimental method to find the controlling mechanism. By measuring the conversion time for particles of different initial sizes, the controlling regime can be found. While the first and third cases have the same dependency on radius, they can be distinguished by their differing sensitivities to fluid velocity and temperature. The standard SCM is a very important model. However, to make a predictive model for a specific process like LiCoO₂ leaching, it is necessary to include more detailed kinetics for that specific system.

2.4. The Role of Turbulence in Particle Dissolution

Turbulence is common in many industrial systems with fluids and solids. It is a chaotic, unsteady, and three-dimensional fluid motion that occurs at high Reynolds numbers [28]. Turbulence is important for this thesis because it greatly increases the transport of mass and momentum compared to laminar flow. This increased transport directly controls the rate of particle dissolution and how the solid particles are distributed in the fluid. This section describes the main characteristics of turbulence, how it interacts with suspended particles, and the computational methods used to model these processes.

2.4.1. Turbulence Scales and the Energy Cascade

A central concept in turbulence is the **energy cascade**. It describes how kinetic energy is transferred from large scales to small scales [28]. Energy enters the flow at the large, energy-containing scales, which are similar in size to the system's geometry. Through nonlinear interactions, this energy is transferred to smaller and smaller scales, or eddies. This process continues until the energy reaches the smallest scales where viscous forces become dominant. At these scales, the kinetic energy is dissipated as heat. This transfer of energy defines three main length scales:

- The **Integral Length Scale** (L) is the scale of the largest eddies. The eddies contain most of the turbulent kinetic energy and their size depends on the flow geometry.
- The **Taylor Microscale** (λ) is an intermediate scale. At this scale strong vortex stretching occurs and viscous effects start to be important.
- The **Kolmogorov Length Scale** (η) is the smallest scale of turbulence. At this scale, the energy cascade ends because viscous dissipation is stronger than inertial energy transfer. It is defined as $\eta = (\nu^3/\varepsilon)^{1/3}$, where ν is the kinematic viscosity and ε is the mean rate of turbulent kinetic energy dissipation.

This multi-scale structure is an important feature of flows at high Reynolds numbers. Kolmogorov's hypotheses states that for high enough Reynolds numbers, the statistics of the small-scale motions ($\ll L$) are universal. This means they are the same for all turbulent flows and depend only on ν and ε [28]. The Kolmogorov scale is very important in multiphase flow modeling because it is used to check the validity of the point-particle assumption. Figure 2.5 shows a diagram of the energy cascade and its associated length scales. Figure 2.6 shows the distribution of turbulent kinetic energy across different wavenumbers, which highlights the different spectral ranges.




Figure 2.5: A schematic diagram of the energy cascade at very high Reynolds number, showing the various length scales and ranges, and the processes of production, transfer, and dissipation of turbulent kinetic energy. From [28, Fig. 6.2].




Figure 2.6: The model spectrum normalized by the Kolmogorov scales. This log-log plot shows the energy spectrum $E(\kappa)$ versus wavenumber κ , clearly showing the inertial subrange (Kolmogorov's $-5/3$ law) and the exponential decay in the dissipation range. From [28, Fig. 6.13].


2.4.2. Homogeneous vs. Inhomogeneous Turbulence

In computational studies, turbulent flows are often simplified into two types. The first is homogeneous isotropic turbulence, used for fundamental analysis. The second is inhomogeneous turbulence, which represents most real engineering systems.

Homogeneous Isotropic Turbulence (HIT) HIT is a theoretical type of turbulence where the flow statistics are the same everywhere (homogeneous) and in all directions (isotropic) [28]. In simulations, it is created inside a box with periodic boundaries. This setup is very useful for basic research of how particles interact with turbulence, such as particle dispersion and clustering [29]. The research in this

this thesis uses a pre-computed **Direct Numerical Simulation (DNS)** dataset of HIT from the Johns Hopkins Turbulence Database (JHTDB) [14]. This very accurate dataset was made with a pseudo-spectral code at a Taylor-scale Reynolds number (Re_λ) of 418. It provides a fully resolved turbulent velocity field. Using this data allows a direct study of how eddies and particles interact and how this affects mass transfer. This is similar to fundamental studies like the one of Haugen et al. (2018), who used DNS to see how particle clustering affects scalar transport [30]. These studies usually assume one-way coupling or very low volume fractions. This makes it possible to focus mainly on the interaction between turbulence and mass transfer without additional complexities. A sample velocity field is shown in Figure 2.7a.

Inhomogeneous and Anisotropic Turbulence Most engineering systems, like stirred tank reactors, have turbulence that is inhomogeneous and anisotropic. The flow field is spatially non-uniform. It has important features like high-shear regions near the impeller, large-scale recirculation zones, and areas with low turbulence [31]. Because of this, the mass transfer rates and particle distributions change a lot in different parts of the vessel. Most applied dissolution simulations study these systems to optimize equipment design and scale-up. Studies by Hartmann et al. (2006), with Large Eddy Simulation (LES), and Li et al. (2024), with CFD-DEM, have shown that high turbulence near the impeller increases local dissolution. They also showed that poor mixing in "dead zones" can create strong concentration gradients and slow down the overall process [32–34]. A typical flow field from a stirred tank simulation is shown in Figure 2.7b.



(a) Homogeneous isotropic turbulence example snapshot from JHTDB.

(b) Inhomogeneous impeller-driven turbulence. From [35, Fig. 8 (a)].

Figure 2.7: Comparison of typical flow fields for different turbulence categories. (a) shows the uniform statistical properties of isotropic turbulence, while (b) shows the spatial variations characteristic of impeller-driven flows.

These two approaches work well together. HIT simulations help us understand basic physics. They are used to develop and test sub-models for particle dynamics and mass transfer in simple, controlled conditions. Then, inhomogeneous simulations use these models to predict how real engineering systems will perform, including the complex details of industrial equipment.

2.4.3. Particle-Turbulence Interactions

The motion of a particle in a turbulent flow depends on the balance between its inertia and the changing forces from the fluid. This interaction controls the particle's trajectory, its location in the fluid, and, as a result, its mass transfer rate.

Turbulence directly affects the relative velocity (u_{rel}) between particles and the fluid around them. The changing fluid velocity causes a changing drag force on the particles. This leads to continuous changes in their relative velocity. Haugen et al. (2018) showed that this is a main reason why turbulence affects mass transfer, because a higher mean relative velocity between particles and fluid increases the mass transfer rate [30]. This effect is measured using the particle Reynolds number (Re_p), which is an important parameter in mass transfer equations.

Turbulence greatly increases particle dispersion. This means particles spread out much faster than they would do by molecular diffusion alone. This happens because the random motion of turbulent eddies

carries particles across streamlines. How well the particles disperse depends on the particle inertia compared to the timescales of the turbulent eddies (Stokes number)[28].

Preferential Concentration (Clustering)

An important phenomena in turbulent flows with particles is **preferential concentration**, also called clustering [28, 30]. This happens when particles with inertia accumulate in a specific area of the turbulent flow. For example, heavy particles tend to collect in regions of low vorticity and high strain where the fluid tends to converge [28]. In contrast, very light particles (e.g., bubbles) may accumulate in high-vorticity regions. Haugen et al. (2018) show that particle clusters can become very dense. This process greatly slows down the local mass transfer rate. This is especially true for systems with large Damköhler numbers, where the chemical reaction is much faster than the turbulent mixing [30].

Turbulence Modulation

Particles can also change the turbulence of the fluid. This process is called **turbulence modulation**. Depending on their size, density, and concentration, particles can either decrease or increase the fluid turbulence [28]. For example, large particles can reduce turbulent kinetic energy through drag. Small, heavy particles can reduce the effect of small-scale eddies. Also, dissolving particles can create local turbulence from buoyancy if the fluid density changes a lot when the solid dissolves. In this thesis, a one-way coupling method is mainly used. In this method, the effect of particles on the fluid's velocity field is ignored. This is a good assumption because the particle concentration is low, and the effect of particle-particle collisions on the flow is very small [30].

The Stokes Number

How particles respond to turbulent eddies can be grouped into three cases. The case depends on the ratio between the particle response time (τ_p) and the eddy turnover time (τ_{eddy}):

- **Passive Tracers** ($\tau_{eddy} \gg \tau_p$): Particles are small and light enough to follow the fluid motion perfectly. In this case, the relative velocity between particles and the fluid is very small.
- **Resonant Eddies** ($\tau_{eddy} \approx \tau_p$): In this middle case, the particle inertia is similar to the eddy time scale. Because of their inertia, particles can be thrown out of eddies. This leads to strong particle clustering [30].
- **Heavy Bullets** ($\tau_{eddy} \ll \tau_p$): Particles have high inertia and tend to move in straight lines. They are not much affected by the fast motions of small eddies. In this case, the eddy velocities have a large effect on the relative velocity between the particle and the fluid [30].

The dimensionless **Stokes number**,

$$St = \frac{\tau_p}{\tau_L} \quad (2.24)$$

where τ_L is the integral time scale of the turbulence, is a key parameter that controls how particles respond to large-scale turbulent eddies [30]. The particle Reynolds number ($Re_p = U_{rel}d_p/\nu$), which describes the flow around the particle, is important for finding the mass transfer coefficient.

The Damköhler Number

The effect of preferential concentration on a reactive system depends on the competition between the fluid transport rate and the chemical reaction rate. This interaction is described by a dimensionless number called the **Damköhler number (Da)**. It is defined as the ratio of the fluid's characteristic time scale to the reaction's characteristic timescale:

$$Da = \frac{\tau_{flow}}{\tau_{reaction}} \quad (2.25)$$

Here, τ_{flow} is a timescale for the fluid motion, such as the turnover time of a large eddy. This time represent how long it takes for turbulence to mix reactants over a large scale. The time $\tau_{reaction}$ is the time needed for the chemical process to consume a large part of a reactant. The value of the Damköhler number separates the problem into two main cases that are important for this thesis.

High Damköhler Number ($Da \gg 1$): This case, where $\tau_{flow} \gg \tau_{reaction}$, means that the chemical reaction is much faster than the rate of turbulent mixing. The system is called **transport-limited** or




Figure 2.8: Conceptual diagram showing particle response cases in turbulent flow. The cases depend on the ratio of particle response time to eddy turnover time (or Stokes number), which leads to different behaviors from passive tracers to heavy bullets.




Figure 2.9: Snapshots of particle number density in homogeneous isotropic turbulence for different Stokes numbers (St). The images show the effect of preferential concentration: (a) $St = 1$, (b) $St = 0.3$, (c) $St = 0.1$, and (d) $St = 0.03$. Particle clustering decreases as the Stokes number gets smaller, which leads to a more uniform distribution for particles with less inertia. From [30, Figure 3].

diffusion-limited. In this situation, the effects of preferential concentration become very important. A dense cluster of particles acts as a strong, local sink for reactants. The fluid around the cluster cannot supply fresh reactants fast enough to match the fast consumption inside the cluster. This creates a situation called **local reactant starvation**. The overall reaction rate of the system is no longer controlled by the chemical kinetics. Instead, it is limited by the speed of mass transport to the particle surfaces.

Low Damköhler Number ($Da \ll 1$): This case, where $\tau_{\text{flow}} \ll \tau_{\text{reaction}}$, means that the reaction is very slow compared to the fluid mixing time. The system is called **kinetically-limited**. Turbulent mixing is very fast compared to the slow reaction. It can easily resupply any consumed reactants, and keep the concentration nearly uniform everywhere, even inside particle clusters. Because of this, the spatial distribution of particles has almost no effect on the overall process rate. The rate is controlled almost completely by the chemical kinetics at the particle surface.

This idea, used in important studies like the one of Haugen et al. [30], is needed to understand the

results of reactive multiphase simulations. The numerical studies in Chapter 5 are designed to explore how the system behaves as it changes between these two cases. They also measure the decrease in performance caused by clustering when the system becomes transport-limited.

2.5. Review of Numerical Modeling and Experimental Studies

2.5.1. Turbulence Modeling Approaches

The choice of a turbulence model is very important in computational fluid dynamics (CFD). It has a big impact on the accuracy of the flow simulation and on the computational cost. There are three main types of models: Reynolds-Averaged Navier-Stokes (RANS), Large Eddy Simulation (LES), and Direct Numerical Simulation (DNS).

Reynolds-Averaged Navier-Stokes (RANS) RANS models solve for the time-averaged flow field, not the instantaneous fluctuations. This makes the computation much cheaper. The Navier-Stokes equations are averaged, which creates new unknown terms like the Reynolds stresses. These terms must be calculated using a closure model [28]. Common RANS models, like the $k - \epsilon$ and $k - \omega$ models, use the idea of a non existing turbulent viscosity to model the Reynolds stresses [28]. RANS models are computationally cheap and are used a lot in industry. However, they average out all the transient turbulent eddies. This is a problem for processes that depend on these eddies, like particle dispersion and local mass transfer. So, RANS models often need to be adjusted for multiphase systems. For example, Hörmann et al. (2011) [36] used a RANS model to simulate a stirred vessel. The model was good for the overall behavior and gave reasonable estimates for the average dissolution rate. However, because it is an averaged model, it could not fully capture the changing concentration fields or particle clustering.

Large Eddy Simulation (LES) LES is a balance between the high accuracy of DNS and the low cost of RANS [28]. In LES, the large-scale turbulent eddies are solved directly. These eddies depend on the geometry and are anisotropic. The effects of the smaller, subgrid-scale (SGS) eddies are modeled using a subgrid-scale model, like the Smagorinsky model [28]. These small eddies are assumed to be more universal. LES gives a more detailed picture of the transient mixing and local eddy effects around particles compared to RANS. This is important for dissolution, where local turbulence has a big effect on mass transfer. However, LES is much more expensive than RANS. This means it is usually only used for smaller-scale research or for systems where high turbulence detail is necessary. Hartmann et al. (2006) [32] used LES to simulate dissolution in a stirred vessel, which gave a more detailed view of particle dispersion. But as noted by Grisafi et al. (2023) [34], the cost of LES is very high, and it can be hard to validate the model without good experimental data.

Direct Numerical Simulation (DNS) Direct Numerical Simulation (DNS) solves the unsteady Navier-Stokes equations directly. It resolves all scales of turbulent motion, from the largest energy-containing eddies down to the smallest Kolmogorov scales [28]. Because DNS does not use any turbulence models, it provides the most accurate results for turbulent flows. For this reason, DNS is very valuable for fundamental research and is used as a reference to validate other turbulence models.

The main limitation of DNS is its very high computational cost. The cost increases quickly with the Reynolds number, typically scaling with $Re^{9/4}$ or even Re^3 [28]. For example, a DNS for a moderate turbulent Reynolds number of $Re_\lambda = 100$ can require more than 10 days of CPU time, while a simulation for $Re_\lambda = 800$ would take thousands of years on the same computer [28, Table 9.1]. Because of this cost, DNS is only possible for low-to-medium Reynolds number flows in simple domains. It is not a practical tool for complex, industrial-scale problems, where methods like Large Eddy Simulation (LES) or Reynolds-Averaged Navier-Stokes (RANS) are used.

To combine the accuracy of DNS with a practical computational cost, this thesis uses a different approach. Instead of running a new, fully coupled simulation, this work uses a pre-computed DNS dataset from the Johns Hopkins Turbulence Database (JHTDB) [14]. In this method, the Lagrangian tracking of particles and their dissolution are calculated inside this static, high-resolution velocity field. The fluid turbulence is fully resolved, but the particles are treated as points, which is a common simplification in

point-particle DNS (PP-DNS) studies [30]. This approach offers several clear advantages compared to running a new, fully coupled DNS:

- **Computational Efficiency:** It removes the most expensive task, which is solving the Navier-Stokes equations for the turbulent flow. This makes it possible to make fundamental studies at high resolution.
- **High Fidelity:** It gives access to a fully resolved turbulent velocity field. This makes sure that particle-eddy interactions are captured with the accuracy of a DNS, but without the high cost.
- **Focus on Dispersed Phase Physics:** Since the fluid field is a known input, the research can focus completely on the complex physics of particle transport, mass transfer, and reaction kinetics.
- **Reproducibility:** using a public, standard database makes sure that the flow conditions are clearly defined and can be reproduced by other researchers.

This method is a good and practical way to study the basic mechanisms of particle dissolution in turbulence. It provides a level of detail that would be impossible to achieve otherwise because of the high cost.

In practical work, LES and DNS are often used for fundamental research or small studies where turbulence details are needed. RANS, or unsteady RANS (URANS), is often used for scale-up and engineering design because it is cheaper.

Table 2.2: Comparison of Turbulence Modeling Approaches

Feature	RANS (Reynolds-Averaged Navier-Stokes)	LES (Large Eddy Simulation)	DNS (Direct Numerical Simulation)
Resolved Scales	Mean flow field	Large, energy-containing eddies	All turbulent scales (integral to Kolmogorov)
Modeled Scales	All turbulent fluctuations (Reynolds stresses)	Small, subgrid-scale (SGS) eddies	None (all scales resolved)
Key Advantage	Low computational cost, good for industrial problems	Balances accuracy and cost, resolves large-scale unsteadiness	Highest accuracy, gives basic understanding, reference for models
Key Limitation	Cannot resolve transient eddies, depends on empirical models	Higher cost than RANS, needs subgrid-scale models	Extremely high cost, limited to low/moderate Reynolds numbers
Typical Use	Industrial design, steady-state flows	Unsteady flows, turbulent mixing, research	Fundamental research, model validation, simple flows
Computational Cost (relative)	Low	Medium	High

2.5.2. Review of Simulating Multiphase Systems

The dissolution of solid particles in turbulent flows has been studied for many years because it is important for many industrial applications. These include crystallization, mineral processing, and battery recycling, where turbulent mixing greatly increases mass transfer [32]. Simulating this process is hard because it must combine fluid dynamics with mass transfer at the particle surface in complex flows. Researchers have developed different computational methods to model these flows, ranging from continuum-based two-fluid models to more detailed particle tracking approaches.

This chapter reviews the main modeling strategies for simulating turbulent dissolution. The review is organized into three main types: **Eulerian-Eulerian (EE)** models, which treat both phases as continuous fluids that mix; **Eulerian-Lagrangian (EL)** models, which track single particles through the fluid; and interface-capturing methods, which resolve the fluid-solid surface directly on the grid. For the EL models, this chapter focuses on the **point-particle method** and the coupled **Computational Fluid Dynamics-Discrete Element Method (CFD-DEM)**. The chapter will also discuss how turbulence models

like Reynolds-Averaged Navier-Stokes (RANS) and Large Eddy Simulation (LES) are used with these methods.

Eulerian–Eulerian Methods Simulations of dissolving particles in turbulent flow can be grouped by how they represent the solid and fluid phases. In Eulerian–Eulerian (EE) or two-fluid models, both phases are treated as continuous fluids that mix together, and each has its own set of field equations. These models solve conservation equations for mass, momentum, and energy. The equations are connected by terms that describe the exchange of momentum (like drag and lift) and mass (dissolution) between the phases [37]. For example, Holbeach et al. (2003) [38] made an EE model for dissolving solids. In their model, the continuity equations for the liquid and solid phases have a source term, $\dot{m}_{ls}$, for the mass transfer from solid to liquid. Because these exchange terms are included, the models naturally allow two-way coupling.

A key advantage of the EE method is that it is good for high particle concentrations, where particle interactions are important for the flow. It can also be combined with population balance models (PBM) to track how the particle size distribution (PSD) changes during dissolution [39]. However, it is hard to correctly model the shrinking of particles. EE models often use simple assumptions, like all particles having the same size or a few size classes, unless a detailed PBM is also used. An example is the work by Zhan et al. (2014) [40]. They used a two-fluid computational fluid dynamics (CFD) model with a population balance to study the slow dissolution of alumina particles in a smelting bath. EE models are good for dense systems, but they become very complex if the PSD is very wide or if the dissolution is not uniform in space and cannot be captured by the averaged fields.

Eulerian–Lagrangian Methods Unlike the EE method, Eulerian–Lagrangian (EL) methods use a hybrid approach. The continuous fluid phase is described on a fixed Eulerian grid. The dispersed solid phase is described by tracking the paths of individual particles in a Lagrangian frame. This allows for a more direct and detailed description of how particles move. The EL category includes several different techniques. The most important are the **point-particle method** and the coupled **Computational Fluid Dynamics–Discrete Element Method (CFD–DEM)**. Each technique has a different balance between computational cost and accuracy.

Point-Particle Method

In the point-particle method, each solid particle is a separate object. Its position, velocity, and mass are tracked over time. The main assumption of this method is that the particles are much smaller than the smallest scales of the fluid flow. These scales can be the grid size in a CFD simulation or the Kolmogorov scale in a turbulent flow [41]. Because of this, the detailed flow and the concentration boundary layer around each particle are not solved directly. Instead, the particle acts as a point source (or sink) of momentum and mass that is smaller than the grid. The total effect of all particles on the fluid is added as source terms to the Navier-Stokes and scalar transport equations. Because the fluid-particle interactions are modeled and not solved directly, this method uses closure models. These models are often empirical equations for drag coefficients and Sherwood numbers to describe the physics of momentum and mass transfer.

The point-particle method is good for simulating systems with many particles, from tens of thousands to millions. It has been used successfully with Direct Numerical Simulation (DNS), Large Eddy Simulation (LES), and Reynolds-Averaged Navier-Stokes (RANS) models. A good example is the work of Hartmann et al. (2006). They used LES to simulate a stirred tank with seven million dissolving particles [32]. In their study, the turbulence was solved with LES. The particle motion was calculated with Newton’s second law, which included drag and collisions. The mass transfer was modeled with the Ranz-Marshall correlation. The simulation used full two-way momentum coupling and scalar coupling. It gave detailed information on particle distribution and dissolution rates, but the computational cost was very high.

The trajectory of a single particle, $\mathbf{x}_p(t)$, is found by integrating its velocity, $\mathbf{v}_p(t)$. The velocity is calculated by integrating Newton’s second law of motion:

$$m_p \frac{d\mathbf{v}_p}{dt} = \sum \mathbf{F}_i, \quad (2.26)$$

where m_p is the particle mass and $\sum \mathbf{F}_i$ is the sum of all forces on the particle. For a small, rigid spherical particle in a fluid, a complete force balance was developed by Maxey and Riley (1983) [42]. If Faxén corrections are ignored, which are for the effect of flow curvature, the equation of motion is:

$$m_p \frac{d\mathbf{v}_p}{dt} = m_p \mathbf{g} + \rho_f V_p \frac{D\mathbf{u}}{Dt} + \mathbf{F}_d + \mathbf{F}_l + \mathbf{F}_a + \mathbf{F}_h. \quad (2.27)$$

Here, the terms on the right-hand side are the gravitational force, the force from the fluid pressure gradient (where ρ_f is fluid density, V_p is particle volume, and $D\mathbf{u}/Dt$ is the material derivative of the fluid velocity $\mathbf{u}$), the drag force ($\mathbf{F}_d$), the lift force ($\mathbf{F}_l$), the added mass force ($\mathbf{F}_a$), and the Basset history force ($\mathbf{F}_h$).

For many practical problems, the full force balance equation can be simplified. A common case is a system of dense solid particles in a liquid or gas, where the particle density is much larger than the fluid density ($\rho_p/\rho_f \gg 1$). In this situation, forces that depend on the displaced fluid mass, such as the added mass and Basset history forces, are very small compared to the particle's inertia and can be ignored [43]. The pressure gradient term is also often ignored when fluid accelerations are small. Lift forces are typically important only with strong fluid shear or particle rotation. Under these conditions, the particle dynamics are dominated by gravity and drag, leading to the simplified force balance:

$$m_p \frac{d\mathbf{v}_p}{dt} = m_p \mathbf{g} + \mathbf{F}_d. \quad (2.28)$$

For small particles, the relative Reynolds number is typically low ($Re_p \ll 1$), which means the motion is in the Stokes flow regime. In this regime, the drag force is linearly proportional to the slip velocity:

$$\mathbf{F}_d = 3\pi\mu d_p(\mathbf{u} - \mathbf{v}_p), \quad (2.29)$$

where μ is the fluid's dynamic viscosity and d_p is the particle diameter. This linear form makes the drag calculation much simpler. If the particle concentration is low, the effect of the particles on the fluid can also be ignored. This approach is known as one-way coupling.

The validity of the point-particle method depends on these assumptions. This thesis studies solid particles with a diameter of about 0.1 mm and densities between 2000 – 5000 kg/m³. The particles are in water ($\rho_f \approx 1000$ kg/m³). For these conditions, the assumption that $\rho_p/\rho_f \gg 1$ is true. This means that ignoring the added mass and history forces is a good starting simplification [43].

CFD–DEM Method The coupled CFD–DEM method is a more detailed Eulerian–Lagrangian model than the point-particle method. It is more precise because it directly includes particle–particle interactions, like contact forces and collisions, using the Discrete Element Method (DEM) [44]. The fluid applies hydrodynamic forces, like drag, on each particle. The particles then apply a force back on the fluid, which is called two-way momentum coupling. The DEM part directly handles interactions between particles and between particles and walls using contact models. This complete model is called four-way coupling because it includes fluid-particle, particle-fluid, particle-particle, and particle-boundary interactions. CFD–DEM is very useful for modeling dense systems, like fluidized beds, or any case where particle collisions are important.

Because computers are now more powerful, recent studies have started to use CFD–DEM for dissolution problems. Li et al. (2024) [33] developed a CFD–DEM model to study the dissolution of solid particles in a stirred tank. They validated the model with experimental data, and the predicted particle size had an error of about 6%. Their study looked at the solute concentration, particle distribution, and local turbulent energy dissipation. They found that high concentration gradients near the particles are connected to faster dissolution rates. They also found that as particles get smaller, the momentum transfer between the fluid and solid phases becomes weaker. This helps smaller particles stay suspended in the flow. The authors also saw that a faster impeller speed greatly increases the dissolution rate at the beginning. However, this difference in dissolution rate becomes smaller after the particles are mixed evenly in the tank.

This study shows that CFD–DEM is very good at capturing both the large-scale mixing and the small-scale mass transfer. However, it has a higher computational cost than simpler point-particle

methods. Overall, EL methods, including point-particle and CFD-DEM, are used a lot for dissolution simulations because they give detailed information about particle dynamics. These methods are often combined with LES for better turbulence resolution or with RANS for cheaper simulations. The high computational cost of CFD-DEM is a motivation to find more efficient methods. This is why this work proposes a different approach, using a pre-computed DNS database with Lagrangian tracking and a point-particle method.

Interface-Capturing Methods Another type of method is the interface-capturing method, such as the Volume of Fluid (VOF) or Level-Set method. These methods are usually used for liquids that do not mix. However, they can also be used for dissolving solids by treating the solid as a separate phase with a clear interface [45]. In these methods, a large solid particle is represented directly on the computational grid. The interface moves as the solid dissolves into the fluid. These simulations solve for the full concentration and flow fields around the particle. For this reason, they are basically a direct numerical simulation (DNS) of the dissolution process.

For example, Safa et al. (2020) [46] used an interface-capturing method called the Smoothed Profile Method (SPM). They studied the dissolution of single particles in shear flows to understand how particle shape affects dissolution rates. However, these methods are very expensive because they need a lot of computer power to resolve the small boundary layers around the dissolving particles. As a result, these methods are usually only used to study one or a few particles. They are often used to develop or check empirical equations, like the relation between the Sherwood number and the Reynolds number. In most real problems, the number of particles is very high, so it is not practical to use these fully resolved methods. This makes point-particle and two-fluid models a better choice. Even with these limitations, interface-resolved studies give important basic information about the physics of dissolution at a small scale.

2.5.3. Salt Dissolution Studies

Salt dissolution (like NaCl or CaCl_2) in water is a classical example, extensively studied as a benchmark for solid-liquid mass transfer. Experimental studies dating back to the 1970s (e.g., Nienow et al. [47]) measured the dissolution rates of salts in stirred tanks, leading to widely used mass transfer correlations. Computationally, salt dissolution has been simulated in both academic and industrial cases. For instance, Hartman et al. [32] conducted LES simulations of calcium chloride dissolving in water under strong agitation in a stirring vessel. Their results indicated that the overall dissolution time was approximately ten times the particle dispersion time, demonstrating that even under strong turbulence, complete dissolution is not instantaneous but follows a finite timescale. Another example is the work by Holbeach et al. (2003) [38], where with an EE approach they modeled the dissolution of salt particles in initially quiescent water. They found that as salt dissolves, the local fluid density increases, leading to natural convection currents that influence particle dispersion. In their model it's important to notice that two-way coupling was fundamental, since the dissolving salt modified the fluid density and velocity, transforming it from quiescent to moving fluid.

This kind of buoyancy-driven flows, becomes particularly relevant in unstirred or weakly agitated conditions. Large crystals dissolving in still water, for example, generate a dense plume that sinks. This leads to an acceleration in dissolution until the plume dissipates. It has to be noticed, that most CFD simulations, assume that density variations are negligible and focus primarily on forced convection driven by turbulence. However, when concentration or temperature gradients are sufficiently large, buoyancy forces should be included in the model. In practical applications, particularly for salt dissolution in water, enhanced mixing generally leads to improved dissolution rates, since the strong recirculation keeps making unsaturated fluid reaching the dissolving surface.

Cao et al. (2016) [48] studied how small pieces of material dissolve in turbulent water using a stirred tank. They used two substances: a solid form of sodium carbonate (Na_2CO_3) plus a porous mix for detergent (PANDORA). They put 0.15 g of each substance in 800 mL of water at 20, 40, 60 °C while running the paddle at 100, 200, 300 rpm. They measured water conductivity to track the dissolution process, and then they normalised the data to yield particle release profiles. They compared these curves to a coupled CFD Noyes-Whitney model. The curves matched the model in high-speed mixing where particles remained suspended. For Na_2CO_3 , differences showed up at low speeds because mixing

was weak. For the detergent mix the model fit the data better after they reset the value for diffusion using early results. This work gives simple curves of how each substance dissolves over time with different mixing speeds, temperatures.

Carletti et al. (2018) [49] performed an investigation of mass transfer in solid-liquid systems in dilute suspensions using two baffled stirred tanks of different sizes. The work was centered on the dissolution behavior of NaCl in particle and tablet forms, and CaCl₂ particles, using different impeller types at different rotation rates and temperatures. To observe the local conductivity distribution as a function of time, Electrical Resistance Tomography (ERT) was used, providing a spatial resolution higher than that offered by standard single-point measurement methods. The mass transfer coefficients (k_L) were calculated from the ERT measurements. This work provides local temporal profiles and spatially averaged dimensionless conductivity curves, which detail the influence of particle size, agitation speed, and temperature. In addition, the results allowed the development of a dimensionless correlation for k_L , which was validated using both the authors' data and literature, thus advancing the understanding of both local and global dissolution behavior under different hydrodynamic conditions.

2.5.4. Lithium Dissolution / Black Mass Leaching Studies

The dissolution of compounds like LiCoO₂ from black mass in acid is the main process in hydrometallurgical recycling of lithium-ion batteries. This process is complex. It involves redox reactions, acid consumption, and the dissolution of several metals. For this reason, no CFD studies exist that model both the fluid dynamics and the coupled reaction kinetics in detail.

An important work, on the reaction kinetic side, is the advanced Shrinking Core Model (SCM) from Cerrillo-Gonzalez et al. (2022) developed in multiple papers[4, 50–53]. This model is the only one in literature that properly models the dissolution with competing reaction pathways and the forming crust (a passivating solid layer Co₃O₄) around each particle. This model is based on average data from bulk experiments and does not include local fluid dynamics. It cannot predict how local transport limits, for example from particle clustering, affect the chemical kinetics. This gap between advanced kinetic models and high-fidelity fluid simulations is the problem that this thesis addresses. Another important point is that the kinetic parameters used by Cerrillo-Gonzalez come from experiments. These experiments can have a high level of uncertainty. For this reason, a sensitivity analysis of these parameters is necessary.

Another study, including the fluid dynamics, by Liu et al. (2023), showed that mass transfer was the rate-limiting step for the leaching of cathode materials [54]. Their work shows that dissolution rates are controlled by the transport of fresh solvent to the particle surface. This CFD study uses RANS models with simplified kinetics and assume constant particle sizes, a big assumption and simplification of the actual phenomena. Moreover, as previously shown, RANS models do not solve the turbulence at the lower scales, so the real interaction with the particles is very simplified.

Validating models for black mass dissolution requires good experimental data. Several recent studies provide this:

Wongnaree et al. (2024) provide useful data for validation [55]. They report the dissolution efficiency over time for key metals in sulfuric acid with H₂O₂. Their study is especially valuable because it includes SEM images that show how particle shape changes over time. It also provides XRD analysis of the solid residue and kinetic equations from a shrinking core model.

Amalia et al. (2023) studied leaching with organic acids in a stirred tank [56]. They report the dissolution efficiency over time for different metals at different stirring speeds. The study also provides SEM images and rate equations from a shrinking core model.

In summary, these studies use stirred tanks, not the isotropic turbulence used in our simulations, making it hard to directly compare. However, they provide valuable time-resolved data for metal dissolution.

2.6. Motivation for the Research

The literature shows that the Eulerian-Lagrangian point-particle method is a good tool to study turbulent reactive dissolution. However, a closer look shows there are still specific gaps in the understanding of these systems. Past work shows the general problems but does not explain well how basic transport, complex reactions, and the behavior of many particles interact. This section uses these gaps to explain

the motivation for the three main research goals of this thesis. Each gap leads to a specific research question.

2.6.1. Impact of Preferential Concentration on Reaction Kinetics

Preferential concentration is a well-known process in fluid dynamics. It happens when particles with inertia accumulate in areas of low vorticity and high strain rate [28, 57]. Basic Direct Numerical Simulation (DNS) studies have shown that this clustering has a complex effect on simple mass transfer. These studies show how particle clustering can cause local reactant starvation, which reduces the total transfer rate [30]. This creates a clear transport limitation that is caused only by the structure of the turbulent flow.

The gap in the literature is that it is not known the quantitative effects of this basic fluid dynamics process on a real industrial reaction system. These systems have multiple steps, multiple species, and non-linear reactions. It is generally known that clustering can reduce performance. But it is not known how the negative effect seen in simple mass transfer appears in a system that is described by the Shrinking Core Model (SCM). The SCM has complex internal resistances, like the formation of a product layer, and involves multiple reactants (e.g., H^+ and H_2O_2) with different roles. The current literature does not measure how the transport limitation caused by preferential concentration interacts with these internal kinetic and diffusive barriers. This leads to the second research question (RQ2):

How does preferential concentration create local areas of low reactant concentration, and how does this transport limitation change the system's total reaction rate and efficiency?

This question connects basic turbulence theory with applied chemical reaction engineering. The goal is to go past the simple observation that clustering exists. The goal is to measure the negative effect of clustering on a realistic leaching process.

2.6.2. Deconstructing Internal Kinetic Mechanisms

The literature provides very good kinetic models for complex leaching processes. These models are fitted to experimental data. An example is the advanced Shrinking Core Model (SCM) for $LiCoO_2$ dissolution from Cerrillo-Gonzalez et al. [4, 50–53]. These models give the necessary reaction steps and kinetic parameters, which are a good starting point for predictive simulations.

However, these experimental studies cannot show how the internal parts of the model work in detail. The literature is missing a systematic study that finds which kinetic steps are the most important. It also does not show the feedback loops between the chemistry and the physics. It also does not measure the two different roles of reactants that can both help the reaction and create a product layer that slows it down. For example, in the SCM, a faster creation of the product layer from one reaction step can slow down the total reaction rate. This is a feedback loop that cannot be seen in simple experiments. This level of understanding can only be achieved with specific numerical simulations.

To fill this gap, a stable and strong numerical tool is needed. The tool must be able to solve the very stiff and non-linear system of ordinary differential equations (ODEs) of the SCM. This need is a good reason for the third research question (RQ3):

Because of the uncertainty in the experimental parameters, which kinetic steps and chemo-physical mechanisms have the biggest effect on the predicted lithium extraction and the total process time?

Answering this question will also give important insights on understanding how experimental errors on the measures of kinetic constants could affect the overall performance.

3

A Unified Model of Dissolution

3.1. Problem Formulation and Objectives

To model the complex dissolution of a particle in a turbulent, reactive fluid, a clear conceptual framework is needed. This framework breaks the process down into its basic physical steps. This chapter develops the mathematical model for this framework.

The main idea of this framework is based on **resistances in series**. This concept was introduced in Chapter 2 and comes from electrical circuit theory. In this model, the total dissolution rate is controlled by a total resistance made of several smaller resistances in series. This means that the slowest step, the one with the biggest resistance, controls the overall rate of the whole process. This step is called the rate-limiting step.

For a reactant molecule to travel from the main fluid to react with the solid material inside the particle, it must pass through three different barriers in order. Each of these barriers acts as a resistance to the total process:

- R1: External Film Mass Transfer.** This is the resistance for the reactant to move from the main, well-mixed fluid to the outer surface of the particle. It must cross the hydrodynamic boundary layer, or "film", around the particle.
- R2: Internal Product Layer Diffusion.** After reaching the surface, the reactant must move through the porous product layer. This layer is sometimes called the "ash" or "crust" and is formed by earlier reactions. Moving through this labyrinth of pores to reach the unreacted core is the second resistance.
- R3: Intrinsic Surface Reaction.** Finally, when the reactant reaches the surface of the unreacted core, the chemical reaction happens. The speed of this chemical reaction is the third and final resistance.

The goal of this chapter is to develop the mathematical equations for this conceptual model. First, the chapter will find the equations for each of the three resistances. Then, it will combine these parts into a single, unified equation. Finally, the chapter will show how this unified model can be simplified for specific limiting cases. This provides the basic logic for the dual-mode numerical framework that is developed in Chapter 4.

3.2. Core Model Assumptions

The unified model in this chapter is based on a set of main assumptions. These are the same assumptions used in the classical Shrinking Core Model. These assumptions separate the main physical processes. This helps to create a clear and simple theoretical model. The main assumptions are:

- **Spherical Particle.** The model assumes the particle is a perfect sphere and stays a sphere during the whole dissolution process. This simplification reduces the geometry to one dimension, the radius. This allows for a simple analytical solution that shows the main physics clearly.

- **Pseudo-Steady-State Approximation (PSSA).** The concentration profile of the reactant in the fluid is assumed to be in a *Pseudo-Steady-State*. This is true for the external boundary layer and for the pores inside the particle. This means that the concentration field changes immediately when the unreacted core shrinks. This assumption is good because the time for diffusion is much smaller than the time for the chemical reaction and for the core to shrink.
- **Isothermal Conditions.** The system is assumed to have a constant temperature (isothermal). This choice separates the transport processes from thermal effects, like the heat of reaction or kinetics that depend on temperature. This allows the study to focus only on the main mass transport and reaction processes.
- **Constant Properties.** To derive the rate equation for a single moment in time, the properties of the main fluid are assumed to be constant. This means that the reactant concentration far from the particle, C_{bulk} , and other fluid properties do not change. This assumption is valid for a single moment in time or for a system with a very large amount of fluid.

3.3. Analysis of Serial Mass Transfer Resistances

Because of the PSSA, the molar rate of reactant transport, $\dot{N}$ (in mol/s), must be the same through each layer. The governing equation for the rate through each of the three resistance layers is now given.

Resistance 1: External Film Mass Transfer

The first step is the transport of the reactant from the main fluid to the outer surface of the particle. This happens by convective mass transfer. The molar rate is the mass transfer flux multiplied by the external surface area, $4\pi r_p^2$. The rate is driven by the difference in concentration between the main fluid and the particle surface:

$$\dot{N} = k_L(4\pi r_p^2)(C_{bulk} - C_{surface}) \quad (3.1)$$

where $C_{surface}$ is the reactant concentration at the external surface (at radius r_p). The term k_L is the mass transfer coefficient. This parameter shows how well the convective transport works and depends strongly on the local fluid motion. Its value is usually calculated from equations for the dimensionless Sherwood number, Sh , as explained in Chapter 2.

Resistance 2: Internal Product Layer Diffusion

After the reactant reaches the particle surface, it must diffuse through the porous product layer to get to the unreacted core. This process is modeled as Fickian diffusion through a spherical shell. The integrated form of Fick's first law gives the molar rate (this is derived in Appendix A):

$$\dot{N} = D_{eff} \cdot 4\pi \frac{r_p r_c}{r_p - r_c} (C_{surface} - C_{core}) \quad (3.2)$$

where C_{core} is the reactant concentration at the surface of the unreacted core (at radius r_c). The term D_{eff} is the **effective diffusivity** of the reactant inside the porous layer. This parameter includes the effect of the solid material in the product layer, which makes the diffusion path more difficult.

Resistance 3: Intrinsic Surface Reaction

The last step is the chemical reaction at the surface of the unreacted core (with area $4\pi r_c^2$). For a general reaction of order n , the rate is:

$$\dot{N} = k_s(4\pi r_c^2)(C_{core})^n \quad (3.3)$$

Here, k_s is the surface reaction rate constant and n is the reaction order for the reactant.

To create a simple model that clearly shows how the resistances work together, a first-order reaction ($n = 1$) is assumed. This is the most common case and the math is simple, which makes it easy to show the main idea. It is important to know that the real kinetics used in the `ShrinkingCoreModel` are much more complex. They have multiple reactions with non-integer orders, as explained in Chapter 4. Using a first-order reaction is good enough for this explanation. This is because it keeps the rate proportional

to the concentration. This is needed to show that the resistances simply add together in the final unified model.

3.4. The Unified Resistance Equation

To combine these steps into one equation, PSSA is needed. The PSSA says that the molar rate, $\dot{N}$, is the same through each layer. The goal is to write $\dot{N}$ using only known values: the bulk concentration, C_{bulk} , and the particle geometry (r_p and r_c). To do this, the unknown concentrations at the interfaces, $C_{surface}$ and C_{core} , must be removed. This is done by rearranging the rate equation for each step to find the concentration drop. Then, these drops are added together, as shown in Appendix A. The total molar rate of reactant consumption is:

$$\dot{N} = \frac{C_{bulk}}{\underbrace{\frac{1}{k_L(4\pi r_p^2)}}_{\text{Film Resistance}} + \underbrace{\frac{r_p - r_c}{D_{eff}(4\pi r_p r_c)}}_{\text{Product Layer Resistance}} + \underbrace{\frac{1}{k_s(4\pi r_c^2)}}_{\text{Reaction Resistance}}} \quad (3.4)$$

This result shows that the total rate is equal to a total driving force (C_{bulk}) divided by the sum of the three single resistances.

To describe how the particle changes over time, this rate must be connected to the rate of core shrinkage, dr_c/dt . For a general reaction $A + bB \rightarrow \text{Products}$, where B is the solid reactant with a molar density of $\rho_{B,molar}$, a mole balance at the core surface gives the rate of core shrinkage (the derivation is in Appendix A):

$$\frac{dr_c}{dt} = -\frac{bC_{bulk}/\rho_{B,molar}}{\frac{r_c^2}{r_p^2 k_L} + \frac{r_c(r_p - r_c)}{r_p D_{eff}} + \frac{1}{k_s}} \quad (3.5)$$

This equation is the final result of the derivation. It is a single, important equation that describes the whole physical process. The process is controlled by the competition between the three basic resistances to dissolution.

3.5. Limiting Cases of the Unified Model

The strength of the unified governing equation (3.5) is its flexibility. By looking at physical limits where one resistance is much larger or smaller than the others, the general equation becomes simpler. These simpler equations describe specific physical situations. This provides the direct theoretical reason for the dual-mode design of the numerical simulator.

3.5.1. Film Resistance Limit (MassTransferOnly Model)

This mode models simple dissolution, like salt in water. It is based on two assumptions:

1. **Instantaneous Dissolution:** The surface reaction is very fast ($k_s \rightarrow \infty$). This makes the reaction resistance go to zero:

$$\text{Reaction Resistance} = \frac{1}{k_s} \rightarrow 0$$

2. **No Product Layer:** Simple dissolution does not create a product layer, so $r_c = r_p$. This makes the diffusion resistance in the product layer zero:

$$\text{Pore Diffusion Resistance} = \frac{r_c(R_p - r_c)}{R_p D_{eff}} = \frac{R_p(R_p - R_p)}{R_p D_{eff}} \rightarrow 0$$

If these limits are applied to Equation (3.5), only the film resistance is left. In this case, $dr_c/dt = dr_p/dt$, and the equation becomes:

$$\frac{dr_p}{dt} = -\frac{bC_{bulk}/\rho_{B,molar}}{r_p^2/(r_p^2 k_L)} = -(bC_{bulk}/\rho_{B,molar})k_L \quad (3.6)$$

This shows that for simple dissolution, the shrinkage rate is controlled only by the external mass transfer coefficient, k_L . This is the theoretical basis for the `MassTransferOnly` solver.

3.5.2. Coupled Diffusion-Reaction Limit (ShrinkingCoreModel)

This mode models complex leaching where turbulent mixing in the bulk fluid is highly efficient. This means there is one limit:

1. **Efficient External Transport:** The external mass transfer coefficient is very large ($k_L \rightarrow \infty$). This means the external film resistance is almost zero compared to the internal resistances:

$$\text{Film Resistance} = \frac{r_c^2}{R_p^2 k_L} \rightarrow 0$$

With this limit, the general equation for dr_c/dt becomes a simpler form. It is controlled by the sum of the two remaining resistances:

$$\frac{dr_c}{dt} = - \frac{bC_{bulk}/\rho_{B,molar}}{\underbrace{\frac{r_c(r_p - r_c)}{r_p D_{eff}}}_{\text{Product Layer Resistance}} + \underbrace{\frac{1}{k_s}}_{\text{Reaction Resistance}}} \quad (3.7)$$

This equation shows that the process is controlled by both the diffusion through the product layer and the surface reaction. This is the precise conceptual model implemented in the `ShrinkingCoreModel` solver.

3.6. Model Implementation Strategy

The conceptual model from this chapter is the physical basis for the dual-mode design of the numerical simulator. The simple, first-order model is not good enough to accurately simulate LiCoO_2 leaching. This is because the process has very non-linear kinetics and is numerically very stiff. The theoretical limiting cases from Section 3.5 map directly to the design of the implemented solvers:

- The **Film Resistance Limit** (Sec. 3.5.1) is the physical basis for the `MassTransferOnly` validation solver. The efficient numerical method used to solve its governing equation is described in Chapter 4.
- The **Coupled Diffusion-Reaction Limit** (Sec. 3.5.2) is the physical model for the main research solver, the `ShrinkingCoreModel`. The advanced numerical methods needed to solve the stiff, non-linear system of ODEs from the real multi-species reactions are shown in Chapter 4.

By treating the two simulation modes as separate limits of this single, general framework, a clear and strong basis is made for the numerical studies that follow.

4

Numerical Framework and Verification

This chapter explains how the numerical framework was built. This framework is the computer implementation of the unified dissolution model from Chapter 3. To change the theoretical model into a reliable computer tool was a difficult task that required many specific choices. Each algorithm and method was chosen to solve a specific physical or numerical problem. These problems are: the transport of chemicals, which is dominated by advection; the very stiff numerical problem of the multi-species reaction kinetics; and the important need for a fully conservative coupling between the Lagrangian particles and the Eulerian concentration fields. The next sections will describe the parts of the simulator. Each part will be explained, from its basic physical assumptions to its numerical tests. This will show that the computational tool used for the main research of this thesis is reliable.

4.1. Overall Modeling Framework

The simulation is designed in a modular way. This allows a clear progression from simple validation tests to complex scientific studies. The core of the simulator is a hybrid model that combines the strengths of the Eulerian and Lagrangian methods. This makes it possible to model the multiphase system with good accuracy and efficiency. This section gives an overview of this framework. It describes its main concepts, design choices, and what it is able to do.

4.1.1. A Eulerian-Lagrangian Method

The model makes a distinction between two components. The first is the continuous fluid, which is described in a fixed Eulerian frame. The second is the solid particles, which are tracked one by one in a Lagrangian frame.

- **The Eulerian Domain.** The fluid is represented by a set of fixed, two-dimensional Cartesian grids. These grids store the continuous fluid properties. An important part of this method is the use of a multi-resolution grid strategy:
 - A *static, high-fidelity velocity field* is defined on a main grid. This is a “frozen turbulence” snapshot from the Johns Hopkins Turbulence Database (JHTDB) [14, 58]. This velocity field controls the advective transport for everything in the domain.
 - An *evolving concentration field* for each chemical species is defined on a separate grid. This grid is usually finer and has a higher resolution. This grid tracks the location of the dissolved chemicals. The concentration changes because of the fluid motion and because of the mass exchange with the particles.
- **The Lagrangian Particles.** Solid particles are treated as separate, moving objects. Each particle has its own set of properties that are tracked over time. These properties include its position

($\mathbf{x}_p$) and velocity ($\mathbf{v}_p$), and also its physical and chemical state (like mass, core radius, and crust composition). This method allows to follow the detailed history of each particle as it changes.

The connection between these two frames is made by a two-way scalar coupling. At each time step, particles get information from the local fluid. The interpolated fluid velocity is used to calculate the drag force, and the local reactant concentrations are used to calculate the dissolution rate. In return, particles act as local sources and sinks of mass. They change the Eulerian concentration field when they use reactants and release products. This continuous interaction is the main way the model captures the coupled physics of the system.

4.1.2. Model Implementation and Choices

The implementation of the numerical methods depends on the correct math and also on important design choices. These choices affect the simulation's performance and its stability. This section describes these key implementation choices.

Software The simulator is written in the MATLAB (R2024a) programming environment [12]. This choice was made to use MATLAB's main strengths on computing. MATLAB is very good for matrix operations, which makes vectorization easy. It also has many fast built-in functions (like `griddedInterpolant` and `accumarray`) and good tools for plotting.

Choices for Computational Performance The most important design choice for high computational performance in MATLAB is vectorization. Instead of using slow `for` loops for each particle, all calculations are done on all the arrays at once. This is possible because the main data structure for the particles uses the **Struct-of-Arrays (SoA)** method. In this design, a single struct, `particles`, holds fields that are complete arrays (for example, `particles.pos` is an $[N \times 2]$ matrix). This is better than the slower **Array-of-Structs (AoS)** method. The SoA design is important because it stores all values for one property together in memory.

Main Modeling Assumptions The numerical model is based on some main assumptions. It is important to know these limitations to understand the results correctly.

- **One-Way Momentum Coupling:** The pre-computed fluid velocity field is static and is not changed by the particles. This simplification is good for the dilute systems studied here ($\Phi_v \ll 1$). In these systems, the momentum lost to particle drag is very small compared to the total momentum of the turbulent fluid. Because of this, the model cannot simulate processes like turbulence modulation, which require two-way coupling.
- **The Computational Parcel Model:** It is not possible to simulate every single particle in an industrial system. To solve this problem, this work uses a Computational Parcel Model, also called a "super-particle" method [15]. Each simulated point represents a group of N_A real particles. The path and motion of the parcel are calculated for one representative particle. The effect of the parcel on the concentration field is then multiplied by N_A .
- **Two-Dimensional Domain:** The simulation is done in a two-dimensional slice of a 3D turbulent flow. This greatly reduces the computational cost. However, it does not include three-dimensional turbulence effects, like vortex stretching.

Numerical Robustness The code is written carefully to make it numerically stable. This can be seen in many details. For example, small constants (`eps`) are added to avoid division by zero. Physical limits are checked using functions like `max(0, ...)`. Special cases, like a zero crust thickness in the Jacobian, are handled directly. This method makes sure the simulation is stable and gives physically correct results.

4.1.3. The Computational Workflow

The simulation runs in a main time loop. In each step, the state of the whole system is updated from time t to $t + \Delta t_{sim}$. A single step includes a series of numerical operations. The order of operations in a single time step is as follows. First, the transport of all parts is calculated. Particle positions are updated by integrating their equations of motion, and the Eulerian concentration fields are advected. Next, the

particle-fluid interaction is calculated. The possible reaction rate for each particle is found based on its internal state and the local reactant concentrations. This rate is then limited by the amount of reactant available locally. This ensures mass conservation. Finally, the system state is updated. The limited source and sink terms are applied to the concentration fields, and the internal state of each particle is updated.

4.1.4. Dual-Mode Simulation Capability

An important design choice is the use of two different simulation modes, which are controlled by the `simulation_mode` flag. This two-mode design is not a random choice. It is a direct implementation of the limiting cases from the unified theoretical model in Chapter 3.

- **Mode 1: MassTransferOnly.** This mode is a direct numerical implementation of the **Film Resistance Limit**, from Section 3.5.1. It is designed to model simple dissolution processes that are controlled by diffusion and convection. Its main purpose is validation. By comparing the model's results to known theoretical equations, like the Ranz-Marshall correlation [21], it can be checked that the core transport algorithms are correct.
- **Mode 2: ShrinkingCoreModel.** This mode implements the **Coupled Diffusion-Reaction Limit**, from Section 3.5.2. Its purpose is to study the complex chemical and physical processes that are the main topic of this thesis. The solver models the multi-step, reaction-controlled leaching of lithium cobalt oxide (LiCoO_2). It uses the advanced numerical methods needed to solve the stiff, non-linear kinetics. This mode makes it possible to study processes like reactant starvation in particle clusters, transport limits through product layers, and how the results change with different kinetic parameters.

4.2. The Eulerian Concentration Field Solver

The concentration fields for all dissolved chemical species are simulated on a fine, two-dimensional Eulerian grid. These fields change in time and space. The design of the solver for these fields is very important, because its accuracy, stability, and conservation properties affect the whole coupled system. This is especially true for the Shrinking Core Model, because its non-linear kinetics are very sensitive to the local amount of reactants. The method used in this work uses several advanced numerical techniques.

The Challenge of Advection-Dominated Transport

The transport of a scalar C , which is the concentration of a species, in a fluid moving with velocity $\mathbf{u}$ is described by the advection-diffusion-reaction partial differential equation (PDE):

$$\frac{\partial C}{\partial t} + \mathbf{u} \cdot \nabla C = D_{mol} \nabla^2 C + S \quad (4.1)$$

where D_{mol} is the molecular diffusivity of the species and S represents all local source and sink terms from the mass exchange between particles and fluid. To choose the right numerical method, it's first needed to know which transport process is more important. The importance of advection compared to diffusion is measured by the dimensionless grid Péclet number, Pe_{grid} . This number is the ratio of the two effects at the smallest scale of the simulation, the grid cell size, h_{conc} :

$$Pe_{grid} = \frac{u' h_{conc}}{D_{mol}} \quad (4.2)$$

Here, u' is a characteristic velocity fluctuation. The root-mean-square (RMS) velocity of the JHTDB turbulent field is used, which is about 0.7 m s^{-1} . For the high-resolution concentration grid ($N_{conc} = 3136$), the cell size is $h_{conc} = 2\pi/3136 \approx 2.0 \times 10^{-3} \text{ m}$. A typical molecular diffusivity for ions in water is $D_{mol} \approx 1 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$. With these values, the Péclet number is:

$$Pe_{grid} \approx \frac{(0.7 \text{ m s}^{-1})(2.0 \times 10^{-3} \text{ m})}{1 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}} \approx 1.4 \times 10^6$$

This very high Péclet number is a clear proof that, on the grid scales, transport is dominated by advection. This physical reason allows to simplify the governing equation by removing the molecular diffusion

term. So, the model solves the pure advection-reaction equation. This is a first-order hyperbolic PDE, and it is best written using the material derivative, $D/Dt \equiv \partial/\partial t + \mathbf{u} \cdot \nabla$:

$$\frac{DC}{Dt} = S \quad (4.3)$$

This equation describes the rate of change for a small piece of fluid as it moves with the flow.

The Multi-Stage Conservative Semi-Lagrangian Scheme

The advection-reaction equation (Eq. 4.3) is a hyperbolic PDE, which is a difficult numerical problem. Standard Eulerian methods, like finite difference or finite volume schemes, are limited by the Courant-Friedrichs-Lewy (CFL) stability condition. This would force the simulation time step, Δt_{sim} , to be very small. The time step would be controlled by numerical stability and not by the time scales of the real physics.

Primary Solution: Unconditional Stability To solve this problem, a **Semi-Lagrangian (SL) method** is used. The main idea of the SL method is to find the concentration at a grid point. To do this, the position of the grid point is tracked backward in time along a trajectory. This leads to a "departure point" that is usually not on the grid. The main advantage of this method is that it is unconditionally stable. This important property separates the choice of Δt_{sim} from the CFL limit. This allows to choose a time step based on the physics studied.

Secondary Challenge and Solution: Accuracy A simple SL method has two new problems. First, the accuracy of the whole method depends a lot on the accuracy of the backward trajectory. Second, calculating the concentration at the off-grid departure point using interpolation can create unphysical errors.

To solve these problems, two improvements are done. To make the trajectory accurate, the model uses a better method than the simple first-order Euler back method. A simple Euler has large errors in the curved velocity fields of a turbulent flow. Instead, a more accurate **second-order Runge-Kutta (explicit Midpoint) method** is used to integrate the path. This makes sure the error in the departure point is small, of order $O(\Delta t_{sim}^3)$. To get physically correct results from the interpolation, the model avoids standard methods like cubic splines. These methods can create unphysical oscillations (overshoots and undershoots) near sharp gradients. The solution is to use the **Modified Akima Cubic Hermite Interpolation ('makima')** scheme. This method is designed to be *shape-preserving*. It changes its coefficients locally to stop it from creating new maxima or minima. This is very useful to correctly represent the sharp, thin concentration structures that are common in turbulent flows.

Tertiary Challenge and Solution: Mass Conservation Even with an accurate and shape-preserving SL scheme, there is one last important problem. The interpolation step does not guarantee that the total mass in the domain is conserved. This is a big problem that must be fixed.

To solve this final problem and make sure the physics are correct, the model uses a mass-fixing method from **Jin and Chen (2015)** [59]. This method is much better than a simple scaling of the whole field. It distributes the mass error based on the local shape of the field. This keeps the physical shape of the concentration gradients and makes sure mass is conserved. The method works in these steps:

1. **Calculate Global Mass Error:** The total mass before (M_{before}) and after (M_{after}) the non-conservative advection step is calculated. The mass error is $\Delta M = M_{before} - M_{after}$.
2. **Define Non-Uniform Redistribution Weights:** A weighting field W is created based on the sign of the error.
 - If $\Delta M > 0$ (Mass Deficit): This means interpolation has smoothed out the high concentration areas (peaks). To fix this, the missing mass should be added to the low concentration areas (valleys). The weights are defined as $W = C_{max} - C_{adv}$, where C_{adv} is the advected field and C_{max} is its maximum value.

- If $\Delta M < 0$ (Mass Surplus): This means interpolation has created unphysical high values. The extra mass should be removed from the highest concentration areas (peaks). The weights are defined as $W = C_{adv} - C_{min}$, where C_{min} is the minimum value.
3. **Normalize Weights:** The weighting field W is normalized to create a distribution η that sums to one: $\eta = W / \sum W$.
 4. **Apply Correction:** The final, corrected concentration field C_{corr} is calculated by adding the mass error, distributed by η , back to the advected field:

$$C_{corr} = C_{adv} + \eta \frac{\Delta M}{V_{cell}} \quad (4.4)$$

where V_{cell} is the volume (or area in 2D) of a grid cell.

This complete, conservative semi-Lagrangian scheme is a stable, accurate, and physically correct method for solving the Eulerian concentration fields when advection is dominant.

4.3. The Lagrangian Particle Tracking (LPT) Model

The solid particles are simulated using a Lagrangian method. This means the trajectory and state of each particle over time are tracked by solving its equation of motion. This section explains the model used for the particles. It shows why a simplified model can be used and it describes the numerical methods used to solve the equations.

4.3.1. Model Simplification and Justification

The complete physical model for the motion of a small, rigid, spherical particle in a changing fluid flow is the Maxey-Riley equation [42]. This equation includes many forces, such as inertia, drag, gravity, pressure gradients, added mass, and the Basset history force. However, this equation is very complex and computationally expensive. To create a model that is fast enough to be useful, this equation must be simplified. It has to be found which forces are the most important for our specific problem.

To do this, the particle Reynolds number were checked, $Re_p = \rho_f |\mathbf{u} - \mathbf{v}_p| d_p / \mu_f$. This number tells what kind of flow is around the particle and which physics are important. The distribution of Re_p was collected from all active particles in a representative MassTransferOnly simulation. The resulting probability density function (PDF) is shown in Figure 4.1.




Figure 4.1: Probability density function (PDF) of the particle Reynolds number (Re_p). The distribution is mostly at very low values. These values are many orders of magnitude smaller than the limit for Stokes flow ($Re_p \approx 1$), which is shown by the dashed red line. It confirms that even the maximum observed Re_p is still in the Stokes regime. This is strong numerical proof that using the Stokes drag law is correct.

The results clearly show that the particles are well inside the **Stokes creeping flow regime** ($Re_p \ll 1$).

The mean Reynolds number is about 3.23×10^{-6} , and the maximum value observed is only 1.85×10^{-4} . This result from the data gives a strong reason to simplify the governing equation of motion. Specifically:

1. The low Reynolds number allows to model the hydrodynamic drag force using the linear **Stokes drag law**:

$$\mathbf{F}_d = 3\pi\mu_f d_p (\mathbf{u}(\mathbf{x}_p, t) - \mathbf{v}_p(t)) \quad (4.5)$$

where μ_f is the fluid dynamic viscosity, d_p is the current particle diameter, and $\mathbf{u}(\mathbf{x}_p, t)$ is the local fluid velocity at the particle's position.

2. An analysis of the forces, supported by the data, shows that other forces can be ignored:
 - **Pressure Gradient and Added Mass Forces:** These forces depend on the volume of fluid moved by the particle. In this study, the particle density is much larger than the fluid density ($\rho_p/\rho_f \gg 1$). This means the particle's inertia is much larger than the inertia of the moved fluid, so these terms can be ignored [43].
 - **Basset History and Lift Forces:** These forces are important when the particle accelerates quickly or when there is high shear. The very low Reynolds numbers in Figure 4.1 mean the slip velocity is very small. This makes these forces not important compared to the main Stokes drag.

This simplification, reduces the governing equation of motion to a balance of inertia, gravity, and Stokes drag. This gives a model that is computationally efficient and physically accurate for this study:

$$m_p \frac{d\mathbf{v}_p}{dt} = m_p \mathbf{g} + \mathbf{F}_d \quad (4.6)$$

4.3.2. Numerical Stiffness of Particle Equations

The integration of Equation 4.6 is difficult to solve numerically because of its **stiffness**. A system is stiff when its dynamics happen on very different timescales. For the particle equation of motion, the fastest timescale is the momentum relaxation time, τ_p . This is the time a particle needs to match the velocity of the surrounding fluid:

$$\tau_p = \frac{\rho_p d_p^2}{18\mu_f} \quad (4.7)$$

For a typical $100 \mu\text{m}$ LiCoO_2 particle in an aqueous solvent, this relaxation time is very short, $\tau_p \approx 1.5 \times 10^{-5}$ s. In contrast, the simulation time step, $\Delta t_{sim} = 0.05$ s, is much larger because it is chosen to capture the large eddies of the fluid turbulence, which is the slowest timescale. The ratio of these timescales, the stiffness ratio, is very high: $\Delta t_{sim}/\tau_p \approx 3300$. This high stiffness means that any standard explicit integrator (like a Runge-Kutta method) would need to use a time step similar to τ_p to be stable. This makes the simulation too slow to be practical.

For this reason, an **A-stable implicit method** is necessary. To solve this stability problem, the **Implicit Euler method** is used. This method is numerically stable for any time step size when used for stiff linear problems. The method discretizes Equation 4.6 by evaluating the drag force at the future time step, t^{n+1} . This changes the differential equation into an algebraic equation for the unknown future velocity $\mathbf{v}_p^{n+1}$:

$$m_p \frac{\mathbf{v}_p^{n+1} - \mathbf{v}_p^n}{\Delta t_{sim}} = m_p \mathbf{g} + 3\pi\mu_f d_p^{n+1} (\mathbf{u}^{n+1} - \mathbf{v}_p^{n+1}) \quad (4.8)$$

The equation is rearranged to solve for $\mathbf{v}_p^{n+1}$:

$$\mathbf{v}_p^{n+1} = \frac{\mathbf{v}_p^n + \Delta t_{sim} \left(\mathbf{g} + \frac{1}{\tau_p^{n+1}} \mathbf{u}^{n+1} \right)}{1 + \frac{\Delta t_{sim}}{\tau_p^{n+1}}} \quad (4.9)$$

Equation 4.9 has a non-linearity because the fluid velocity $\mathbf{u}^{n+1}$ must be found at the future particle position $\mathbf{x}_p^{n+1}$. This position, however, depends on the unknown velocity $\mathbf{v}_p^{n+1}$. This problem is solved

using a **fixed-point (Picard) iteration**. For each particle, the iteration starts with a guess for the new velocity. This guess is used to predict a future position and to interpolate the fluid velocity at that position. The interpolated fluid velocity is then used in Equation 4.9 to calculate an updated velocity. This process repeats until the change between successive velocity values is smaller than a set tolerance. This method provides a stable solution and converges quickly for this problem.

4.4. Particle-Fluid Coupling Strategy

The coupling between the discrete Lagrangian particles and the continuous Eulerian fields is a central part of the simulation. This model captures the physics of multiphase reactive transport because it includes the feedback between the solid and fluid phases. The next sections describe this strategy. First, is explained the one-way treatment of momentum. Second, is detailed the two-way framework for scalar transport. Finally, is presented the custom algorithm that ensures physical conservation.

4.4.1. One-Way Momentum Coupling

In this work, the particle suspensions are dilute, with a very low particle volume fraction. For such systems, the momentum transfer from the particles to the fluid is small enough to be ignored [43]. For this reason, a **one-way momentum coupling** approach is used. This choice is made to isolate how a fixed turbulent field affects particle transport and reaction. It removes the complex effect of particles changing the turbulence itself. This simplification allows a clear analysis of the main research questions about transport and preferential concentration. The same strategy is used in other modern simulations of reactive transport [30].

4.4.2. Two-Way Scalar Coupling

While momentum coupling is one-way, the scalar coupling is **two-way**. Particles use the local reactant concentration to find their dissolution rate. In return, their dissolution acts as a source of product and a sink of reactant, which changes the concentration field on the Eulerian grid. This continuous, two-way exchange of information is how the model captures system-level effects like local reactant starvation. Numerically, this transfer is done with two operations: (1) Eulerian-to-Lagrangian Transfer, particles get local field data from the grid using interpolation. (2) Lagrangian-to-Eulerian Transfer, particle source terms are mapped back to the grid to update the fluid.

4.4.2.1. Numerical Implementation of Coupling Transfers

Eulerian-to-Lagrangian Transfer Transferring information from the grid to the particles is a trade-off between computational cost and accuracy. Particles need local fluid properties to calculate forces and reaction rates. These properties must be found at the exact particle position, which is usually located between the grid points. For this task, which is done very often in a simulation, **bilinear interpolation** is used.

Figure 4.2 shows how this method works. It uses the values from the four surrounding grid points to calculate the field value at the particle's location. This method is very fast. For this reason, it became a standard method in important simulations of particle-laden turbulent flows [57, 60]. The accuracy of this method depends on the grid resolution. If the grid is not fine enough to resolve the turbulent eddies, the interpolation can add numerical errors. This trade-off is well studied [61]. For the well-resolved fields used in this work, it gives a good balance of accuracy and speed. This use of bilinear interpolation should not be confused with the more advanced **makima** scheme, which is used for the more difficult task of scalar field advection (Section 4.2).

Lagrangian-to-Eulerian Transfer Transferring information from the particles back to the grid is a challenge for mass conservation. The source terms for mass, which are calculated at the particle locations, must be mapped to the Eulerian grid points. A simple method, like assigning the source to the nearest grid point, is not conservative and creates large errors related to the grid structure.

To solve this, the model uses a conservative **area-weighting (or Particle-in-Cell) scheme** [62]. The source term from one particle is distributed among the four grid points of the cell that contains it. The fraction given to each grid point is proportional to the area of the rectangle between the particle and the opposite corner of the cell. This method guarantees that the total source value mapped to the grid is




Figure 4.2: A schematic of bilinear interpolation. The particle at (x_p, y_p) is inside a grid cell. The value at its position is calculated from the values at the four surrounding grid points.

exactly equal to the total source value from the particles.

4.4.3. Mass Conservation via the Potential-Limit Algorithm

A key challenge in the simulation is to ensure that a particle does not consume more reactant than is available in its local Eulerian cell. If this happens, it would create unphysical negative concentrations and cause the simulation to fail. This problem is especially important for the multi-reactant Shrinking Core Model, where the reaction must stop when the limiting reactant is gone.

To solve this problem, a custom "**Potential-then-Limit**" algorithm is used. This method separates the problem into two parts: first, the calculation of the particle's reaction physics, and second, the limits from its environment.

The algorithm has two main stages:

1. **Stage 1: Calculate Potential Change.** First, the kinetic solvers calculate the *potential* change for each particle over the time step Δt_{sim} . This is done by assuming each particle reacts in an infinite amount of fluid, where the reactant concentrations are equal to the current local values. This calculation gives the unconstrained mass change. For the SCM, it also gives the unconstrained moles of each reactant that would be consumed. This value is the maximum possible reaction, based only on the particle's internal state and its own kinetics.
2. **Stage 2: Limit the Change and Update.** Second, these potential consumption values are checked against the actual amount of reactants available in the local grid cells. For each particle, a dimensionless scaling factor, $\alpha_{particle} \in [0, 1]$, is calculated. This factor represents the fraction of the potential reaction that can physically happen. Separate scaling factors are calculated for each reactant based on its local availability. The final limiting factor for the particle, α_{final} , is the minimum of all these factors. This makes sure the overall reaction is limited by the single scarcest reactant. Finally, all particle and fluid variables are updated using this single, physically correct scaling factor.

This two-stage method is key for the numerical stability of the model. It is a strong, physical mechanism that handles the complex feedback between particle kinetics and local reactant supply. It guarantees that the simulation is physically accurate, even when reactant concentrations are very low.

4.5. Particle Internal State Evolution

The main physical law of this work is the Unified Resistance Equation from Chapter 3. As shown before, this single law has two limiting cases: the Film Resistance Limit and the Coupled Diffusion-Reaction Limit. Each limit has a different mathematical structure. For this reason, an important design choice

was to implement two different solvers, controlled by the `simulation_mode` flag. This is a practical way to implement the unified theory in code. The design uses two specialized numerical solvers. Each one is optimized for the mathematical problems of its physical regime. This gives the best stability, accuracy, and speed. The next sections describe these two solvers.

4.5.1. The Diffusion-Controlled Solver (MassTransferOnly Mode)

In the `MassTransferOnly` mode, the solver integrates the equation for the change in particle radius, r :

$$\frac{dr}{dt} = -\frac{Sh(Re_p)D_{mol}}{2\rho_p r}(C_s - C_{bulk}) \quad (4.10)$$

This ordinary differential equation (ODE) has two numerical problems. The first is **stiffness**. Because of the $1/r$ term, the equation becomes very stiff as the particle gets smaller ($r \rightarrow 0$). This makes standard explicit integrators (like Forward Euler or RK4) numerically unstable and too slow. The second problem is **non-linearity**. To solve the stiffness problem, an implicit method like Implicit Euler could be used. However, this would require calculating the Sherwood number at the future radius, $Sh(r^{n+1})$. This creates a non-linear algebraic equation for r^{n+1} that needs an iterative solver (like Newton's method) for every particle at every time step. This is too slow.

To solve both problems at the same time, a **semi-implicit, linearized formulation** is used. The method starts with an Implicit Euler discretization for stability. Then, it uses a good physical approximation to remove the non-linearity: $Sh(r^{n+1}) \approx Sh(r^n)$. This approximation works well because the Sherwood number is a weak function of the radius ($Sh \propto \sqrt{r}$), and the radius changes very little in one small time step. This linearization changes the difficult non-linear problem into a simple quadratic equation for r^{n+1} :

$$(r^{n+1})^2 - r^n r^{n+1} + \left[\frac{Sh(r^n)D_{mol}(C_s - C_{bulk})\Delta t_{sim}}{2\rho_p} \right] = 0 \quad (4.11)$$

This method is a good compromise. It has the stability of an implicit method and the speed of a direct algebraic solution. The implementation solves this equation with the standard quadratic formula. A more detailed derivation is in Appendix C.

4.5.2. The Reaction-Controlled Solver (ShrinkingCoreModel Mode)

The `ShrinkingCoreModel` is used to simulate the complex leaching of LiCoO_2 . This model describes the reaction kinetics and the changes in the particle's internal structure. The physical model is a system of coupled, non-linear Ordinary Differential Equations (ODEs) for the particle state vector $\mathbf{y} = [r_c, m_{\text{crust}}]^T$. Here, r_c is the radius of the unreacted core and m_{crust} is the moles of the solid crust material (Co_3O_4). The system is described by these equations, as shown by Cerrillo-Gonzalez et al. [50]:

$$\frac{dr_c}{dt} = -\frac{M_{\text{LiCoO}_2}}{\rho_{\text{LiCoO}_2}} \frac{r_1(\mathbf{y}) + r_3(\mathbf{y})}{4\pi r_c^2} \quad (4.12)$$

$$\frac{dm_{\text{crust}}}{dt} = \frac{1}{6}r_1(\mathbf{y}) - r_2(\mathbf{y}) - r_4(\mathbf{y}) \quad (4.13)$$

These ODEs are the detailed numerical implementation of the physical ideas from the Coupled Diffusion-Reaction Limit of the Unified Model. They are a more complete version of the simple conceptual model.

- **The Molar Rate of Core Consumption:** In the simple model, the driving force for the reaction is a simple, first-order term. In the full model, as shown in Equation (4.12), this is replaced by the specific reaction rates $r_1(\mathbf{y})$ and $r_3(\mathbf{y})$. These rates include the true, multi-species, non-integer-order kinetics of the core dissolution (see Equations 4.14 and 4.16).
- **The Competing Resistances:** In the simple model, the effects of pore diffusion and surface reaction are two separate terms added together. In the full model, this competition is described more accurately. The reaction rates $r_1(\mathbf{y})$ and $r_3(\mathbf{y})$ depend on the reactant concentration at the core surface (C_{core,H^+}). This concentration is not a given value. It is calculated from a balance between the diffusive transport through the crust (Resistance 2) and the reaction at the core surface (Resistance 3). This balance is solved at every time step by finding the root of a cubic polynomial, which is explained in Appendix D.





The simple model from Chapter 3 provides the basic idea of competing resistances. Equations (4.12) and (4.13) and their sub-models translate this idea into a practical numerical tool that can handle the real complexity of the LiCoO_2 leaching system.

Solving this system of equations is difficult for standard ODE solvers. There are three main challenges. First, the system has **severe numerical stiffness**. This is because the surface reactions are very fast compared to the much slower diffusion and crust growth. Second, the system has **strong non-linear coupling** because the reaction rate functions $r_i(\mathbf{y})$ make the state variables depend on each other in a complex way. Third, the system has an **intrinsic algebraic constraint**: the reactant concentration at the core surface must be calculated by solving a non-linear equation at every time step.

To solve this difficult problem, customized function were created. Here different methods were used like the **L-stable Rosenbrock method**. This is an implicit integrator that is good for stiff systems like this one. For better accuracy and speed, the solver uses an **analytically derived Jacobian matrix**. This avoids the errors and high computational cost of numerical differentiation. The algebraic constraint is solved with a **direct analytical solver** for the internal cubic polynomial. This finds the core surface concentration without using a slow iterative method. A more detailed description of these methods is further in this chapter, the full derivations for the Jacobian and the cubic polynomial are in Appendices D and E.

The Governing Physical Model

The model uses four chemical reactions to describe the leaching process. These reactions show how the solid core (LiCoO_2) is consumed, how a solid product crust (Co_3O_4) is formed and then dissolves, and how two reactants in the liquid, H^+ and H_2O_2 , are consumed. The four main reactions are based on the work of Cerrillo-Gonzalez et al. [50]:



The rates of these reactions use non-integer order kinetic laws. The specific rate constants and other parameters are shown in Table 4.1. An important part of the model is that the properties of the Co_3O_4 crust change over time. The **crust porosity** (ϵ) and the **effective diffusivity of H^+** (D_{eff, H^+}) through the crust are calculated at each time step. These calculations depend on the current core radius and the amount of crust material. This captures how the structure of the crust affects reactant transport.

Table 4.1: Parameters for the LiCoO_2 SCM from Cerrillo-Gonzalez et al. [50].

Parameter	Value	Units	Description
k_1	2.17×10^{-8}	$\text{m} \cdot \text{s}^{-1}$	Rate constant for core dissolution (Eq. 4.14)
k_2	8.33×10^{-9}	$\text{m}^2 \cdot \text{mol}^{-2/3} \cdot \text{s}^{-1}$	Rate constant for crust dissolution (Eq. 4.15)
k_3	6.67×10^{-11}	$\text{m}^7 \cdot \text{mol}^{-7/3} \cdot \text{s}^{-1}$	Rate constant for core dissolution (Eq. 4.16)
k_4	1.67×10^{-12}	$\text{m}^8 \cdot \text{mol}^{-5/3} \cdot \text{s}^{-1}$	Rate constant for crust dissolution (Eq. 4.17)
k_D	-3.2×10^7	m^{-1}	Empirical constant for effective diffusivity

Physical Model of the Porous Crust

An important part of the SCM is the model for the porous product layer (the crust) that forms around the unreacted core. The properties of this crust control how fast reactants diffuse to the core surface. Here the main properties are two:

- **Porosity (ϵ):** The porosity of the crust is calculated at each step. It is based on the volume of the spherical shell and the volume of the solid crust material (Co_3O_4).

- **Effective Diffusivity ($D_{eff,H}$):** The diffusion of reactants depends on the effective diffusivity. This value is much lower than the normal molecular diffusivity because the path through the porous material is complex. It is very hard to model it. For this reason, this work uses an empirical model for the effective diffusivity of H^+ from Cerrillo-Gonzalez et al. [50]. This model connects $D_{eff,H}$ to the changing crust porosity and thickness. It is used because it is made to capture the complex effects of pore structure in leaching systems.

The Rosenbrock Method for Stiff Systems

To solve the stiff ODE system, a **First-Order Rosenbrock method (ROS1)** is used. Rosenbrock methods are L-stable. This means they are very stable for stiff problems because they correctly handle the fast-changing parts of the solution that can cause instability [63]. The ROS1 method updates the solution by first calculating an intermediate vector $\mathbf{k}$. This is done by solving a single linear system:

$$\left(\frac{\mathbf{I}}{\gamma \Delta t_{sim}} - \mathbf{J}(\mathbf{y}^n) \right) \mathbf{k} = \mathbf{f}(\mathbf{y}^n) \quad ; \quad \mathbf{y}^{n+1} = \mathbf{y}^n + \mathbf{k} \quad (4.18)$$

Here, $\mathbf{I}$ is the identity matrix, $\gamma = 1$ is a constant for the method, and $\mathbf{J}$ is the Jacobian matrix $\partial \mathbf{f} / \partial \mathbf{y}$. It is stable because it solves a linear system with the Jacobian, but it only needs one Jacobian calculation and one matrix factorization per time step.

Analytical Jacobian Formulation

An important feature that improves the accuracy and speed of the Rosenbrock solver is the use of an **analytical Jacobian**. The other option is to use a finite-difference approximation, but this method has errors and is sensitive to the chosen step size. In contrast, the analytical Jacobian is exact and is faster to calculate.

The derivation of the Jacobian elements, $J_{ij} = \partial f_i / \partial y_j$, is complex because of the structure of the SCM equations. For this reason, the full analytical derivation is shown in Appendix E.

An important part of this implementation is that it works correctly at physical limits. For example, the formulation correctly handles the case when there is no crust. In this case, the crust thickness is almost zero, and normal diffusion equations would fail. This ensures the Jacobian is always well-defined, which prevents numerical errors and makes the solver more stable.

Adaptive Sub-stepping for Efficiency and Accuracy

An internal while loop for adaptive sub-stepping is used. The main simulation step Δt_{sim} is divided into smaller, variable sub-steps, Δt_{sub} . The size of Δt_{sub} is controlled by a simple adaptive method that checks how much the state vector $\mathbf{y}$ changes. If the change is large, it means the system is changing quickly, so Δt_{sub} is made smaller. If the change is small, Δt_{sub} is made larger to save time. This makes the solver take small steps when the reaction is fast and large steps when the reaction is slow.

The Analytical Cubic Solver

In each evaluation of the SCM physics function $\mathbf{f}(\mathbf{y})$, an important sub-problem must be solved: finding the reactant concentration at the core-crust interface, C_{core,H^+} . This value is found from a balance between the diffusion of reactants to the core and the consumption of reactants by the reaction at the core surface.

As detailed in Appendix D, this physical balance creates a non-linear algebraic equation. For the specific kinetic laws used here, this equation can be changed into a special cubic polynomial of the form $Px^3 + Qx + R = 0$, where $x = C_{core,H^+}^{1/3}$. A function solves this equation directly and analytically using Cardano's method. This choice is much better than an iterative numerical solver. It is very fast, it is guaranteed to find all real roots without needing a good initial guess, and it includes logic to choose the single root that is physically correct ($0 \leq C_{core} \leq C_{bulk}$).

4.6. Model Verification

Before using the numerical model from Chapter 4 for the main research, it is important to check its accuracy and robustness. This chapter shows the results from key verification studies. These studies

make sure that the numerical tool works correctly. Verification confirms that the code correctly solves the mathematical equations. These checks are necessary before doing the scientific studies in Chapters 5 and 6.

4.6.1. Grid Independence Study

First, it is important to check the individual parts of the model. This section describes a grid convergence study. The study checks the spatial accuracy of the Eulerian field solver. It also helps to choose a grid resolution that gives a good balance between computational cost and accuracy for the next simulations.

Methodology

A standard test case was used to check only the spatial discretization errors. The test case simulates the dissolution of a single, stationary particle at the center of the domain. The simulation was run in the `MassTransferOnly` mode, and the fluid velocity was set to zero. This removes the effects of fluid motion and lets the test focus only on how the solver handles concentration gradients. The simulation time step was kept constant and small ($\Delta t_{sim} = 0.00625$ s) to make sure that errors from the time integration did not affect the results.

The same simulation was run on six different grids. Each grid was finer than the one before, with resolutions (N_{conc}) from 640 to 4672. The final total particle mass, $M(h)$, was used to compare the results. Here, h is the grid spacing, $h = L_{domain}/N_{conc}$. The solution error for each grid, $E(h)$, was calculated by comparing its result to the result from the finest grid (h_{finest}):

$$E(h) = |M(h) - M(h_{finest})| \quad (4.19)$$

Analysis of Solution Convergence

The main goal of the study was to show that as the grid gets finer (as $h \rightarrow 0$), the numerical solution becomes stable and does not change anymore. The results are shown in Table 4.2 and Figure 4.3.

Table 4.2: Results of the grid refinement study. The final particle mass converges as the grid resolution N_{conc} increases. The relative change is calculated between each grid and the next finer grid.


Grid Resolution, N_{conc}	Grid Spacing, h (m)	Final Particle Mass (kg)	Relative Change (%)
640	9.82×10^{-3}	6.09×10^{-10}	-
960	6.55×10^{-3}	5.13×10^{-9}	88.12
1408	4.46×10^{-3}	1.22×10^{-8}	58.05
2112	2.98×10^{-3}	2.26×10^{-8}	45.89
3136	2.00×10^{-3}	4.43×10^{-8}	48.99
4672	1.34×10^{-3}	4.57×10^{-8}	3.11

Figure 4.3a clearly shows the solution convergence. For coarse grids ($N_{conc} < 2500$), the final mass changes a lot with the resolution. This shows that the grid is not fine enough to correctly capture the concentration gradients. When the grid resolution is higher than $N_{conc} \approx 3000$, the curve becomes flat. This means the solution changes very little if the grid is made even finer. Table 4.2 shows this numerically. The relative change between the two finest grids is only 3.1%. This demonstrates that the solution is practically grid-independent.


Convergence Rate Verification

To check that the numerical scheme is implemented correctly, its observed order of accuracy was calculated. Figure 4.3b shows a log-log plot of the solution error versus the grid spacing. This is a good verification tool because the slope of this line for fine grids should be equal to the theoretical order of accuracy of the numerical method.

The plot shows two different regions. For coarse grids (large h), the error does not decrease in a predictable way. This is typical when the grid is too coarse to show the true convergence rate. However, for fine grids (small h , specifically for $N_{conc} \geq 2112$), the slope of the error curve becomes parallel to the reference line for second-order ($p = 2$) convergence.



(a) Final particle mass for different grid resolutions (N_{conc}). The value becomes constant for fine grids.



(b) Log-log plot of the solution error versus grid spacing (h_{conc}). The slope shows second-order convergence for fine grids.

Figure 4.3: Grid convergence study showing the solution approaching a stable value.

This result is very important. It provides strong numerical evidence that the main spatial discretization schemes used in the model, like interpolation and gradient calculations, are second-order accurate. The fact that the observed convergence rate matches the theoretical rate confirms that the code is implemented correctly.

Discussion and Final Selection of Grid Resolution

The grid convergence study shows that the numerical model produces a solution that converges to a stable, grid-independent value. The convergence is predictable and follows the theory. The observed second-order convergence rate gives high confidence that the spatial discretization schemes are implemented correctly.

These results help to choose a good grid resolution for the main scientific simulations. The finest grid, $N_{conc} = 4672$, gives the most accurate solution, but it has a high computational cost. The $N_{conc} = 3136$ grid is about 3.5 times faster, and its result is different by only 3.1%. For studies with many simulations, the $N_{conc} = 3136$ grid is a very good compromise between accuracy and computational cost.

For this reason, a concentration grid resolution of $N_{conc} = 3136$ is selected for the main simulations in this thesis. This choice is supported by the convergence data. It ensures that the simulation results have small spatial discretization errors. This allows for a reliable analysis of the physical phenomena.

4.6.2. Time-Step Independence Study

This section shows a study of the simulation time step. This is done to check the accuracy of the time integration and to make sure the code is implemented correctly. The study has two goals. First, to confirm that the numerical method converges at the correct theoretical rate, which shows the code is working as expected. Second, to find a time step, Δt_{sim} , that is small enough so that the results do not change if the time step is made even smaller.

Methodology

To study only the errors from the time step, a standard test case was used. The test case simulates a single particle dissolving in a "frozen turbulence" velocity field. The simulation was run in MassTransferOnly

mode. To reduce errors from the spatial grid, a fine grid with $N_{conc} = 1280$ was used for all tests. An automated study was done. It started with a large time step of $\Delta t_{sim} = 0.1$ s. The simulation was then run again and again with smaller time steps. The study stopped when the relative change in the final particle mass between two runs was less than 0.3%.

For each run, the final particle mass, $M(\Delta t)$, and the computational time were recorded. To analyze the convergence rate, the solution error, $E(\Delta t)$, was calculated. The error is the difference between the result from a given time step and the result from the smallest time step in the study ($\Delta t_{finest} = 0.0024$ s). The error is defined as:


$$E(\Delta t) = |M(\Delta t) - M(\Delta t_{finest})| \quad (4.20)$$

Analysis of Solution Convergence


The main goal of the study was to find a time step that gives a converged solution. The results are shown in Table 4.3 and Figure 4.4.

Table 4.3: Results of the time-step refinement study. Convergence, defined as a relative change below 0.3%, was achieved at $\Delta t_{sim} = 0.004687$ s. The error is calculated with respect to the solution at the finest time step ($\Delta t_{finest} = 0.0024$ s).

Time Step, Δt_{sim} (s)	Final Particle Mass (kg)	Relative Change (%)	Error vs. Finest (kg)
0.100000	9.080×10^{-9}	-	2.818×10^{-9}
0.050000	9.987×10^{-9}	8.38	1.911×10^{-9}
0.025000	1.095×10^{-8}	5.58	9.450×10^{-10}
0.012500	1.148×10^{-8}	3.45	4.220×10^{-10}
0.006250	1.196×10^{-8}	1.72	6.500×10^{-11}
0.004687	1.193×10^{-8}	0.28	3.200×10^{-11}
0.003125	1.200×10^{-8}	0.54	9.800×10^{-11}
0.002400	1.190×10^{-8}	0.82	0



(a) Final particle mass as a function of the simulation time step, Δt_{sim} , showing the convergence to a stable value.



(b) Log-log plot of the solution error versus the time step. The reference lines show first-order ($p = 1$) and second-order ($p = 2$) convergence.

Figure 4.4: Analysis of the simulation's dependence on the time step, Δt_{sim} . Subfigure (a) shows the convergence of the final particle mass. Subfigure (b) shows the solution error and its convergence rate.

The data clearly shows that the solution converges. Figure 4.4a shows that for large time steps ($\Delta t_{sim} > 0.01$ s), the final particle mass depends on the step size. As the time step is made smaller, the curve becomes flat and approaches a stable value. Table 4.3 shows this with numbers. The change from $\Delta t = 0.006250$ s to $\Delta t = 0.004687$ s resulted in a relative change of only 0.28%. This met the 0.3% convergence goal. For even smaller time steps, the relative change shows a small oscillation. This can happen when the solution error is very small, close to the precision limit of the computer.

Convergence Rate Verification

It is also important to check that the numerical method behaves as expected from theory. The log-log plot of solution error versus time step in Figure 4.4b shows this verification. The convergence process has two stages.

Stage 1: First-Order Convergence (Large Time Steps: $\Delta t > 0.01$ s)

For large time steps, the observed error (blue curve) follows the line for first-order ($p = 1$) convergence. This is the expected theoretical behavior. The main part of the numerical method which solves for the different physical processes (like fluid motion and chemical reactions) in separate, sequential stages, is first-order accurate. This match shows that the main error source in the code is implemented correctly.

Stage 2: Transitional Behavior (Small Time Steps: $\Delta t < 0.01$ s)

When the time step becomes smaller than about $\Delta t = 0.01$ s, the error curve changes. The slope becomes steeper, which looks like the convergence rate is better than first-order. This is a good sign. It means that the main first-order error is now very small. It is smaller than other, higher-order errors in the system, like the second-order advection. As the solution gets very close to the true value, the error decreases very fast. This is another confirmation that the solution has converged.

Discussion and Final Selection of Time Step

The study was successful for two reasons. First, the observed first-order convergence for large time steps confirms that the main part of the numerical method is implemented correctly. Second, the study found a time step for which the solution is converged. The results show that a time step of $\Delta t_{sim} = 0.004687$ s meets the convergence criterion of less than 0.3% relative change. The oscillation seen for smaller time steps shows that using an even smaller time step costs more computation time but does not make the solution much more accurate.

For these reasons, a time step of $\Delta t_{sim} = 0.004687$ s is chosen for all scientific simulations in this thesis. This choice is based on the convergence data. It ensures that the results are not affected by large errors from the time step. This allows for an accurate analysis of the physics.

5

Results I: The Salt Dissolution Case

This chapter starts the investigation with a simple physical system: the diffusion-controlled dissolution of solid particles in a turbulent fluid. This system is like salt dissolving in water. This simple case is used to answer the first research question (RQ1), which is about the effect of fluid motion. By studying the system without chemical reactions, the simulator's core transport algorithms can be carefully validated. This process creates a reliable baseline for the transport physics, which is needed for the more complex reactive studies in Chapter 6.

The reason for this approach is to study interacting physical processes in isolation. To do this, the simulations in this chapter use a special `MassTransferOnly` mode. This mode turns off all chemical reaction models. This simplification allows for a direct study of the interaction between fluid turbulence and mass transfer at the particle scale, without the complication of surface chemistry. The result is a clear picture of the hydrodynamic transport mechanisms. This provides an important benchmark for the simulator's accuracy and a reference for understanding the more complex systems that follow.

The investigation is structured around one main question that gives a baseline answer to the primary research questions.

1. **RQ1 (Turbulence Influence):** How does the intensity of turbulence affect the rate and efficiency of mass transfer? This analysis validates the numerical model by comparing simulation results against the well-known Ranz-Marshall correlation. This confirms that the simulator correctly captures standard transport physics. This analysis also shows an important system behavior called **particle clustering**.

This chapter is organized as follows. It starts with model validation, then explains the physical mechanisms, and ends with a practical guide for process optimization.

5.1. Influence of Turbulence on Mass Transfer

This investigation first analyzes how turbulence intensity controls the rate of solid-liquid mass transfer. The analysis shows that the overall process is a result of two competing physical mechanisms. These are a baseline molecular diffusion, which is always present, and a convective transport driven by turbulence, which enhances the rate under specific flow conditions.

This section demonstrates this effect by separating the simulated Sherwood number into its diffusive and convective parts. It then uses the Ranz-Marshall correlation as a theoretical framework that correctly models this two-part structure. Finally, this framework is used to derive distinct scaling laws for both the diffusion-dominated and the convection-dominated regimes. The good agreement between the simulation data and these theoretical laws validates the physical accuracy of the model.

5.1.1. Parametric Study Setup

The results in this section come from a parametric study where the intensity of the background turbulent velocity field was changed. All other physical and numerical parameters were kept constant. This makes

sure that any differences in the results are only because of the change in hydrodynamic conditions. The main variable was the `turbulence_scaling_factor`, which was changed across 20 logarithmically spaced points from 0.1 to 1000. For each scaling factor, 5 separate simulations, called ensembles, were run.

Ensemble averaging is necessary to get reliable statistics. The simulation uses a static, or "frozen," turbulent velocity field. Because of this, the result of a single simulation depends a lot on the random starting positions of the particles. For example, if particles start in a good spot, they might dissolve quickly. If they start in a bad spot, like inside a vortex, they might get trapped and dissolve slowly. A single simulation run is not enough to show the average system behavior.

To solve this problem, an ensemble-averaging method was used. For each turbulence level, the simulation was run 5 times (`num_ensembles` = 5). The initial particle positions were re-randomized for each run. The number 5 was chosen as a good balance between getting a stable average result and the high computational cost of the study. In total, the study includes 100 independent simulations (20 turbulence levels $\times$ 5 ensembles).

The main parameters were kept constant across all simulations. They are listed in Table 5.1.

Table 5.1: Key physical, numerical, and computational parameters used for the turbulence intensity parametric study. These values were kept constant across all simulation runs.

Parameter	Symbol	Value	Unit / Description
Fluid Properties			
Fluid Density	ρ_f	1000	kg/m ³
Fluid Kinematic Viscosity	ν_f	1.85×10^{-4}	m ² /s
Solid & Mass Transfer Properties (MassTransferOnly Mode)			
Particle Material	-	NaCl (equiv.)	-
Particle Density	ρ_p	2160	kg/m ³
Initial Particle Diameter	d_0	1×10^{-4}	m
Solute Molecular Diffusivity	D_{mol}	1.6×10^{-9}	m ² /s
Saturation Concentration	C_s	360	kg/m ³
Molar Mass of Solid	M_s	0.05844	kg/mol
Computational & Domain Parameters			
Domain Length	L	2π	m
Number of Particles Tracked	n_p	1000	-
Molar Aggregation Factor	N_A	1×10^5	particles/parcel
Velocity Grid Resolution	N_{vel}	64	-
Concentration Grid Resolution	N_{conc}	3136	-
Temporal and Solver Parameters			
Time Step	Δt	0.005	s
Final Time	t_{final}	6.0	s
Implicit Solver Tolerance	$tol_{implicit}$	1×10^{-6}	-

5.1.2. Fundamental Transport Regimes

To understand the effect of hydrodynamics, which is part of the first research question (RQ1), it is important to look at the two parts of the mass transfer process. The total mass transfer is not one single process. It is a competition between two mechanisms. The first is molecular diffusion, which is always present. The second is convective transport from turbulence, which adds to the diffusion. Which of these two is stronger controls the physical regime of the system. This section shows this competition by looking at the parts of the Sherwood number (Sh) that are used in the numerical model. This shows the main transport physics of the system.

Figure 5.1 shows this competition. The figure separates the total Sherwood number into its two parts during the simulation. The vertical axis is the value of each Sherwood number component on a logarithmic scale. The horizontal axis is time. The black dashed line is the constant **diffusive component**, Sh_{diff} . This is the baseline for mass transfer. The colored solid lines are the variable

convective component, Sh_{conv} . The color of each line shows the intensity of the background turbulence, as shown on the color bar.




Figure 5.1: Comparison of the convective (colored lines) and diffusive (dashed black line) components of the Sherwood number across a range of turbulence intensities. The analysis shows a clear transition where the convective part becomes larger than the diffusive baseline.

The black dashed line has a constant value of $Sh_{diff} = 2$. This value is not random. It is the theoretical limit for mass transfer from a sphere in a fluid that is not moving. This idea is explained in Appendix B. It is the minimum possible rate of mass transfer, controlled only by molecular diffusion. It is the baseline used to measure all convective effects.

The colored lines, which show Sh_{conv} , behave in a more complex way. As expected, their value depends directly on the turbulence intensity. Higher intensities (warmer colors) cause a stronger convective mass transfer. There is also another effect. For any intensity, the convective term gets smaller over time. This happens because of the physics in the model. Convective transport depends on the particle Reynolds number (Re_p), which depends on the particle radius, r . As the particle dissolves, its radius gets smaller. This makes the Reynolds number smaller, and the convective effect gets weaker.

Putting these two parts together shows the system's operational regime at any time. The regime is found by comparing the size of the convective term (a colored line) to the diffusive term (the black dashed line).

- **The Diffusion-Dominated Regime:** When a colored line is below the black dashed line ($Sh_{conv} < Sh_{diff}$), molecular diffusion is the main transport mechanism. In this state, turbulence has only a small effect on the total mass transfer rate. Figure 5.1 shows this is true for a wide range of turbulence intensities, for all scaling factors up to about 700. It is important to note that even a small convective enhancement, while weaker than diffusion, still improves the total dissolution over time.
- **The Convection-Dominated Regime:** When a colored line is above the black dashed line ($Sh_{conv} > Sh_{diff}$), turbulent convection is the main transport mode and controls the overall rate. This only happens for the highest simulated turbulence intensities.

The data show a clear transition around a turbulence scaling factor of $I \approx 700$. At this point, the system's physics changes from being diffusion-dominated to convection-dominated.

The next important question is: How does this change in the main transport mechanism affect the observed dissolution over time, like the change in particle surface area and mass? The next sections will answer this question by looking at the simulation results and comparing them to the theoretical scaling laws for each of these two regimes.

5.1.3. Deconstructing the Ranz-Marshall Correlation

The analysis in the previous section showed that the system has two different mass transfer regimes. To study these regimes, a quantitative physical model is needed to describe mass transfer from a sphere with different levels of convection. The **Ranz-Marshall correlation** is a good model for this purpose. It is a well-known semi-empirical model that works well for the physical problem in this work: convective mass transfer from spheres at low or medium particle Reynolds numbers. The correlation is:

$$\text{Sh} = 2.0 + 0.6 \cdot \text{Re}_p^{1/2} \cdot \text{Sc}^{1/3} \quad (5.1)$$

where Sh is the Sherwood number, Re_p is the particle Reynolds number, and Sc is the Schmidt number.

The Ranz-Marshall correlation can be separated into two parts that have a physical meaning. The first term, 2.0, is the **diffusive Sherwood number**, Sh_{diff} . Its value, $\text{Sh}_{\text{diff}} = 2.0$, is the theoretical minimum for mass transfer. It describes pure molecular diffusion from a sphere into a fluid with no motion. This term comes from first principles (Laplace's equation), as shown in Appendix B. This part represents the basic resistance to mass transfer that does not change with the turbulent flow. This is the constant baseline that was shown in Figure 5.1.

The second term in Equation (5.1) is the **convective Sherwood number**, Sh_{conv} , which is $\text{Sh}_{\text{conv}} = 0.6 \cdot \text{Re}_p^{1/2} \cdot \text{Sc}^{1/3}$. This term describes the increase in mass transfer that is caused only by the relative motion between the particle and the fluid. Its value depends on two dimensionless groups. The Schmidt number, $\text{Sc} = \nu / D_{\text{mol}}$, is a ratio of fluid properties and is constant for a simulation. The important dynamic variable is the particle Reynolds number, $\text{Re}_p = u_{\text{rel}} d_p / \nu$. Here, u_{rel} is the relative particle-fluid velocity and d_p is the particle diameter.

Because it depends on the particle Reynolds number, Sh_{conv} changes in space and time during the simulation. This change is caused by two things. First, the particle diameter, d_p , gets smaller as the particle dissolves. This explains why the convective term in Figure 5.1 decreases over time. Second, the relative velocity, u_{rel} , is not a fixed value. It is a result of the complex interaction between a particle and the turbulent eddies. The average of this velocity is controlled by the **Turbulence Scaling Factor**, I . For these reasons, the convective enhancement term can be seen as a function of the two main variables of the system: turbulence intensity and particle radius. This relationship can be written as $\text{Sh}_{\text{conv}} = f(I, r)$. This provides a clear theoretical explanation for the variable part of mass transfer seen in the simulation.

The Ranz-Marshall correlation, therefore, gives a good model that separates the total mass transfer into a constant diffusive baseline ($\text{Sh}_{\text{diff}} = 2$) and a dynamic convective part ($\text{Sh}_{\text{conv}} = f(I, r)$). This model is not just for description; it can also be used to make predictions. It allows us to derive different mathematical scaling laws that describe how a particle should dissolve in each of the two limiting regimes. The next section will use this model to derive these scaling laws for the change in particle surface area in both the diffusion-dominated limit ($\text{Sh} \approx 2$) and the convection-dominated limit ($\text{Sh} \approx \text{Sh}_{\text{conv}}$). These theoretical predictions will then be compared to the simulation data. This comparison is a good test to check if the numerical model is physically correct.

5.1.4. Validation Against Scaling Laws

Section 5.1.3 established the Ranz-Marshall correlation as the theoretical model for this system. From this model, it is possible to derive testable predictions about the system's behavior. This section validates the numerical model by deriving the expected mathematical scaling laws for the two limiting regimes, diffusion-dominated and convection-dominated. These theoretical predictions are then compared directly to the simulation results.

This comparison is an important test of the model's physical accuracy. If the simulation data match the derived scaling laws, it gives high confidence that the model is implemented correctly. The analysis examines each of the two limiting regimes, first deriving the theoretical prediction and then comparing it to the simulation data.

The Diffusion-Dominated Limit: A Constant Rate of Surface Area Decay

The theoretical behavior in the diffusion-dominated limit comes from the main Ordinary Differential Equation (ODE) for the change in a particle's radius, r :

$$\frac{dr}{dt} = -\frac{\text{Sh} \cdot D_{\text{mol}}(C_s - C_{\text{bulk}})}{2\rho_p r} \quad (5.2)$$

As shown in Section 5.1.2, the system is diffusion-dominated for low turbulence (scaling factor $I < 700$). In this limit, the convective effect is negligible, so the Sherwood number can be approximated by its minimum value, $\text{Sh} \approx 2$. Putting this value into Equation (5.2) gives a simpler equation:

$$\frac{dr}{dt} \approx -\frac{2 \cdot D_{\text{mol}}(C_s - C_{\text{bulk}})}{2\rho_p r} = -\frac{K_1}{r} \quad (5.3)$$

where $K_1 = D_{\text{mol}}(C_s - C_{\text{bulk}})/\rho_p$ is a group of constant physical properties. To compare this theory with an observable result from the simulation, the particle's surface area, $A_p = 4\pi r^2$, is used. The rate of change of the area is found with the chain rule:

$$\frac{d(A_p)}{dt} = \frac{d(4\pi r^2)}{dt} = 8\pi r \frac{dr}{dt} \quad (5.4)$$

Substituting the simplified equation for dr/dt (Equation (5.3)) into Equation (5.4) shows that the radius term r cancels out:

$$\frac{d(A_p)}{dt} \approx 8\pi r \left(-\frac{K_1}{r}\right) = -8\pi K_1 = \text{constant} \quad (5.5)$$

This derivation gives a clear prediction that can be tested: for a system in the diffusion-dominated regime, the particle surface area should decrease linearly with time.




Figure 5.2: Change in the total particle surface area over time for different turbulence intensities. The curves for low intensities ($I < 700$) show a clear linear decay, which confirms the theoretical scaling law from Equation (5.5).

This prediction is compared to the simulation data in Figure 5.2. To check the theory, the curves for a turbulence scaling factor $I < 700$ are examined. These are the blue and green curves. These curves show a nearly perfect linear decay during the entire simulation, and the constant negative slope is easy to see. The observed linear decay of the surface area for low turbulence matches the theoretical scaling law exactly. This exact match is a strong validation. It confirms that the numerical model correctly captures the basic physics of diffusion-limited mass transfer.

The Convection-Dominated Limit: A Linear Decay of Mass Dissolution Rate

In the convection-dominated limit (for $I \gg 700$), the constant diffusive term $\text{Sh}_{\text{diff}} = 2$ is small compared to the convective term, so it can be ignored. The Sherwood number can then be approximated by the

convective part of the Ranz-Marshall correlation:

$$\text{Sh} \approx \text{Sh}_{\text{conv}} = 0.6 \cdot \text{Re}_p^{1/2} \cdot \text{Sc}^{1/3} \quad (5.6)$$

If the definition of the particle Reynolds number is used, the scaling becomes $\text{Sh} \propto (u_{\text{rel}} \cdot r)^{1/2}$. For a high turbulence intensity I , the relative velocity u_{rel} depends mostly on I . Its dependence on the radius is weak because the radius changes slowly. Therefore is a good approximation that u_{rel} is almost constant. With this, the Sherwood number scaling becomes simpler: $\text{Sh} \propto r^{1/2}$.

The instantaneous mass dissolution rate, $\dot{m}$, is given by $\dot{m} = k_L A_p (C_s - C_{\text{bulk}})$. The mass transfer coefficient is $k_L = \text{Sh} \cdot D_{\text{mol}} / d_p \propto \text{Sh} / r$. So, the mass rate scaling is:

$$\dot{m} \propto \left(\frac{\text{Sh}}{r} \right) \cdot A_p \propto \left(\frac{\text{Sh}}{r} \right) \cdot r^2 = \text{Sh} \cdot r \quad (5.7)$$

Finally, substituting the derived scaling for the Sherwood number gives the theoretical prediction for the mass dissolution rate's dependence on the radius:

$$\dot{m} \propto (r^{1/2}) \cdot r = r^{3/2} \quad (5.8)$$

To determine the temporal evolution of $\dot{m}$, it has to be found the temporal evolution of the radius, $r(t)$, by integrating the governing ODE in this limit. The radius ODE, $\frac{dr}{dt} \propto -\frac{\text{Sh}}{r}$, with the convective scaling $\text{Sh} \propto r^{1/2}$, becomes:

$$\frac{dr}{dt} \propto -r^{-1/2} \implies \frac{dr}{dt} = -K_2 r^{-1/2} \quad (5.9)$$

where K_2 is a constant incorporating the fluid properties and turbulence intensity. This ODE can be solved by separation of variables and integration from an initial state (r_0 at $t = 0$) to a state at time t :

$$r^{1/2} dr = -K_2 dt \quad (5.10)$$

$$\int_{r_0}^{r(t)} r^{1/2} dr = - \int_0^t K_2 dt' \quad (5.11)$$

$$\frac{2}{3} \left[r(t)^{3/2} - r_0^{3/2} \right] = -K_2 t \quad (5.12)$$

$$r(t)^{3/2} = r_0^{3/2} - \frac{3}{2} K_2 t \quad (5.13)$$

This result shows that the quantity $r^{3/2}$ should decrease linearly with time. Since it was previously established that $\dot{m} \propto r^{3/2}$ from Equation (5.8), this means that the mass dissolution rate itself should also decrease linearly with time. This leads to the final temporal scaling law:

$$\dot{m}(t) \propto r(t)^{3/2} \implies \dot{m}(t) \propto t \quad (5.14)$$

The theory predicts that in the convection-dominated regime, the instantaneous mass dissolution rate $\dot{m}$ should decrease linearly with time.

This prediction is checked with the data in Figure 5.3. The curves for high turbulence intensities ($I > 700$), which are the yellow and orange lines, look very different from the curves for low turbulence. These curves show a clear linear decay over time. This behavior is different from the non-linear, concave-up decay seen in the diffusion-dominated regime. This observed linear decrease matches the predicted linear decay from Equation (5.14) very well.

The model's behavior has been successfully validated against the theoretical scaling laws in both of its physical limits. This gives high confidence in its implementation and physical accuracy. The model is now verified and can be used to study more complex effects.




Figure 5.3: Change in the mass dissolution rate over time. The curves for high intensities ($I > 700$) show a clear linear decay, which matches the theoretical scaling from Equation (5.14) for the convection-dominated limit.

5.1.5. Emergence of Polydispersity from Turbulent Heterogeneity

The average properties of all particles show the main behavior of the system. However, this average is made from many different particle histories. An important result of this is the change in the particle size distribution (PSD). A group of particles that all start with the same size (monodisperse) will have different sizes after moving in a turbulent flow.

This creation of different sizes (polydispersity) happens because the turbulence is not the same everywhere. Each particle follows a different path, $\mathbf{x}_i(t)$. Because of this, it sees a different history of local fluid velocities, $\mathbf{u}(\mathbf{x}_i(t), t)$. This leads to a different history of relative velocities for each particle, $u_{\text{rel},i}(t)$, which controls its Sherwood number, $Sh_i(t)$. As a result, the total mass lost over time is different for each particle. Turbulence is the mechanism that changes the initial group of same-sized particles into a wide distribution of sizes. The same convection that increases the average mass transfer rate also creates the differences between individual particles.

The simulation results in Figure 5.4 show this. The figure shows that the particles start at the same size, but the distribution becomes wider over time, like a Gaussian distribution. The average diameter of the distribution gets smaller as the particles dissolve. At the same time, its variance increases, which shows that the dissolution rates are not the same for all particles. This confirms that the model correctly captures how each particle's path affects its dissolution.

This result has two important consequences. First, it shows why it is important to use ensemble averaging. The final state of each particle is random. For this reason, many simulations are needed to find the main system behavior from this random process. Second, it shows a mechanism for creating polydispersity that is only based on transport. This result gives a good baseline for studying more complex systems in the next chapters. In those systems, the velocity and reactant concentration are not uniform, which should create even wider particle distributions.

5.1.6. The Role of Particle Inertia and Stokes Number

The previous sections showed that the numerical model correctly simulates the mass transfer for both diffusion and convection. However, this validation does not explain the physical mechanism of the convective term itself. The important question is: why does the convective enhancement, Sh_{conv} , depend on both turbulence intensity (I) and particle radius (r)? The answer is **particle inertia**. Convective transport is effective when a particle does not follow the fluid streamlines. This effect is measured by the dimensionless **Stokes number** (St).

The Stokes number measures particle inertia relative to the fluid. It is the ratio of the particle's




Figure 5.4: The particle size distribution (PSD) over time. The particles all start at the same size. Because of turbulence, they dissolve at different rates. This makes the distribution wider, and the average diameter gets smaller.

momentum response time, τ_p , to a typical timescale of the fluid flow, τ_f :

$$St = \frac{\tau_p}{\tau_f} \quad (5.15)$$

The **particle response time**, $\tau_p = \rho_p d_p^2 / (18\mu_f)$, is the time a particle needs to adapt to a change in the fluid velocity. It depends strongly on the particle size (d_p^2). This means a particle's inertia changes a lot as it dissolves. The **fluid timescale**, τ_f , is the large-eddy turnover time. It is inversely proportional to the turbulence intensity (I).

The value of the Stokes number shows the different physical regimes of particle-fluid interaction:

- **$St \ll 1$ (Tracer Regime):** The particle responds very quickly to changes in the fluid velocity. It follows the fluid streamlines, so the relative velocity (u_{rel}) is very small.
- **$St \approx 1$ (Inertial/Centrifuge Regime):** The particle response time is similar to the eddy turnover time. This causes the largest separation between the particle and the fluid paths. Particles are thrown out of vortices into regions of high strain. This regime has higher relative velocity and higher convective enhancement.
- **$St \gg 1$ (Ballistic Regime):** The particle has too much inertia and is not affected much by the turbulent eddies. This regime is not important for the particle sizes considered in this work.

To study the effect of inertia, an analysis was done using the data from the parametric study. The results are shown in Figure 5.5. This plot is made from data taken at specific times when the average particle radius, r_{avg} , was a certain fraction of the initial radius (r_0). This method makes it possible to study Sh_{conv} as a function of turbulence intensity for fixed particle sizes. The dotted lines show a theoretical $k \cdot \sqrt{I}$ fit. This fit is a model for a perfect inertial regime ($St \approx 1$), where the particle's relative velocity is directly related to the turbulence. When the simulation data are different from this fit, it shows the true physical regime of the system.

The large particles (yellow curve, $r_{avg} = 0.95r_0$) show a clear transition. For low turbulence, the fluid




Figure 5.5: The convective Sherwood term (Sh_{conv}) versus the turbulence scaling factor (I) for different fixed average particle radii. The dotted lines show a theoretical $k \cdot \sqrt{I}$ fit. The agreement or difference from this fit can be explained by the Stokes number of the particles.

timescale τ_f is long, so the Stokes number is low ($St \ll 1$). The particles act like tracers. Their relative velocity u_{rel} is low and does not follow the turbulence intensity. For this reason, Sh_{conv} is weak and is below the $\sqrt{I}$ fit. As the turbulence intensity increases, τ_f gets shorter, and the Stokes number for these large particles gets closer to one ($St \approx 1$). The system enters the inertial regime. Here, the separation from the fluid is large, u_{rel} is high, and it follows the turbulence intensity. The simulation data match the $k \cdot \sqrt{I}$ model very well.

The smallest particles (dark blue curve, $r_{avg} = 0.20r_0$) behave differently. Because the particle response time depends on d_p^2 , their τ_p is very short. As a result, even at high turbulence, their Stokes number is always small ($St \ll 1$). These particles are always in the tracer regime. Their convective enhancement is always weak and does not follow a clear trend. For this reason, the data deviates from the ideal inertial model.

5.1.7. Quantifying the Inertial Transition Threshold

The previous analysis linked the system behavior to the Stokes number. Now, a quantitative study is done to find a threshold radius, r_{trans} . This radius is the size that separates particles with predictable inertial behavior from particles that act like simple tracers.

To find this radius, a quantitative analysis of the data from Figure 5.5 was done. For each curve of fixed radius, a linear regression was done on the log-log plot of $\log(Sh_{conv})$ versus $\log(I)$. Only data from the high-intensity range ($I > 100$) were used, because convection is dominant there. The coefficient of determination (R^2) from this fit was used as a measure of how well the data fit the inertial model. An R^2 value close to 1.0 means a strong, predictable inertial response. The results of this analysis are shown in Figure 5.6, which plots the calculated R^2 value against the normalized particle radius.

Figure 5.6 clearly shows the transition. For large particles ($r_{avg}/r_0 > 0.25$), the R^2 value is high and stable. This confirms that these particles follow the inertial model well at high turbulence. But, as the particle radius gets smaller, the plot shows a sharp transition. The quality of the fit gets much worse, which shows a change in the physics. A quality threshold of $R^2 = 0.97$ is used to define the inertial regime.

From this analysis, the transition radius is found to be about $r_{trans} \approx 0.22r_0$. This is marked with a star in the figure. Particles larger than this size are controlled by predictable inertial physics in high turbulence. Particles smaller than this size act like tracers, even with high mixing energy.

5.1.8. Defining the Physical Limits of Process Intensification

To connect the validated micro-scale physics to practical process design, the analysis now uses a key performance metric: the **ensemble-averaged** t_{90} . This is the time needed to dissolve 90% of the initial




Figure 5.6: The quality of the log-log fit (R^2) is used to find the transition between tracer-like and inertial particle behavior. The plot of R^2 versus the normalized particle radius (r_{avg}/r_0) shows that the fit quality gets worse for smaller particles. A threshold of $R^2 = 0.97$ is used to find the transition radius.

solid mass. This metric is a good measure of the batch time and operational cost. Figure 5.7 plots this metric against the turbulence intensity (I). The plot shows a process behavior that is controlled by diminishing returns.





Figure 5.7: The ensemble-averaged dissolution time (t_{90}) as a function of turbulence intensity (I). The curve shows a transition from a diffusion-limited plateau to a regime of strong diminishing returns. This shows that simply increasing energy input is not a good optimization strategy.

5.1.8.1. Two Critical Thresholds for Optimization

The "best" operating point depends on the engineering goal. The analysis of the derivatives of the process efficiency ($E = 1/t_{90}$) shows two different and important thresholds. These are shown in Figure 5.8.

The Efficiency Optimum (Operational Cost). The first threshold is found at the peak of the linear marginal efficiency (dE/dI), which is at $I = 0.21$. This point gives the maximum return on energy. This **Efficiency Optimum** is the best operating point to minimize operational cost (OpEx). At this point, each unit of energy gives the largest increase in performance.



(a) The linear marginal efficiency has a peak at $I = 0.21$. This point defines the **Efficiency Optimum** for minimizing operational cost.

(b) The logarithmic marginal efficiency becomes flat for $I \geq 10^4$. This defines the **Performance Plateau** for maximizing throughput.

Figure 5.8: Derivatives of the process efficiency ($E = 1/t_{90}$) show two important transition points for process design. One is for operational cost, and the other is for high-throughput design.

The Performance Plateau (Capital Design). The second threshold is important for designing a high-throughput process. The logarithmic marginal efficiency ($dE/d(\log_{10} I)$) shows that for intensities $I \geq 10^4$, the system reaches a stable **Performance Plateau**. At this point, the relative return on energy becomes constant. This plateau is the upper limit for making the process faster. Investing more money to create turbulence above this level is not efficient, because it does not give a proportional increase in speed.

The Performance Plateau exists because of two competing physical effects. Higher turbulence increases mass transfer, which makes t_{90} smaller. But, it also causes the inertial clustering of particles. This **particle clustering** creates regions where particle concentration is high, which makes the dissolution slower and increases t_{90} . The plateau is a state of balance, where the benefit from more convection is canceled by the negative effect of particle clustering.

Results II: LiCoO₂ Leaching

Chapter 5 validated the transport model and gave a baseline analysis for hydrodynamic influence (RQ1) and introduced the concept of clustering. This chapter now studies the main scientific problem of this thesis. It investigates the complex, coupled process of acid leaching by using the `ShrinkingCoreModel` simulation mode. The analysis goes beyond the simple diffusion-controlled case to revisit the main research questions in a system where hydrodynamics and multi-step surface chemistry are strongly connected (RQ2). A new question about the chemical kinetics (RQ3) is also introduced.

This chapter has two main goals. These goals build on and extend the results from the previous chapter.

1. **Addressing RQ2 (The Consequence of Clustering):** The previous chapter introduced particle clustering. This analysis studies the effects of this non-uniform particle distribution in a reactive system where transport is limited. The main hypothesis is that clustering reduces the performance of the process. The goal is to find the exact physical mechanism that causes this.
2. **Introducing RQ3 (Deconstructing the Kinetics):** What are the controlling mechanisms in the kinetics of the leaching process? This study analyzes the Shrinking Core Model to find the main rate-determining steps, to find important feedback loops like surface **self-passivation**, and to quantify the role of key reactants like hydrogen peroxide (H₂O₂).

To meet these goals, specific computational experiments are used. To analyze the kinetics (RQ3), a series of parametric sensitivity studies is done. These studies show the two roles of hydrogen peroxide: it is the main driver of the reaction and it also stops unwanted side reactions that cause passivation. At the same time, the effect of particle clustering (RQ2) is studied with a comparative numerical experiment. This experiment compares the system performance with a hypothetical uniform particle distribution to the performance with a realistic, clustered distribution. This comparison shows clearly that **local reactant starvation** is the main mechanism that reduces reactor performance because of clustering.

6.1. The Impact of Preferential Concentration on Reaction Rate

In real industrial systems, an important complication emerges: the particles are not evenly dispersed. Instead, they tend to accumulate in certain regions and avoid others. This is called **preferential concentration**. This effect happens because inertial forces push particles out of eddies (regions of high vorticity) and into zones of high strain. The result is a pattern of dense, clusters of particles separated by large, empty regions.

This creates a "clustering paradox". On one hand, turbulence increases the rate of mass transfer for each individual particle. On the other hand, the same turbulence organizes particles into clusters, which can act as local barriers to transport. The process that helps mixing can also hurt it.

To understand this, it is necessary to quantify the competition between particle inertia and fluid motion. The particle's inertia is described by its momentum response time, τ_p . This time depends on the particle's density and the square of its diameter ($\tau_p \propto \rho_p d_p^2$).

Turbulence has a spectrum of timescales. The two most important scales are: the **integral timescale**, T_L (called τ_f in this work), which is the turnover time of the largest eddies, and the **Kolmogorov timescale**, τ_η , which is the characteristic time of the smallest eddies, where velocity gradients are the strongest.

This leads to two important, dimensionless Stokes numbers:

- **The Large-Eddy Stokes Number (St_L):** This is the ratio of the particle response time to the integral timescale. It controls how the particle responds to the large-scale fluid motions.

$$St_L = \frac{\tau_p}{\tau_f} \quad (6.1)$$

- **The Kolmogorov Stokes Number (St_K):** This is the ratio of the particle response time to the Kolmogorov timescale, $St_K = \tau_p/\tau_\eta$. This parameter controls preferential concentration. The strongest clustering happens when particle inertia is in resonance with the small, dissipative eddies, classically when $St_K \approx 1$ [64].

The value of St_L defines different regimes for the particle behavior. When $St_L \ll 1$, particles act like passive tracers and follow the flow. When $St_L \gg 1$, particles are ballistic and are not affected much by the fluid.

This study tests a hypothesis about how clustering affects the total system performance. It is proposed that the performance reduction from clustering can be measured by a dimensionless *effectiveness factor*, $\langle \alpha \rangle$. This is the ratio of the measured global reaction rate to the ideal rate that would happen if particles were perfectly uniform.

It is also hypothesized that $\langle \alpha \rangle$ is not a constant. It depends on two key dimensionless groups:

1. **The Damköhler Number (Da).** When $Da \ll 1$, the reaction is slow and is the limiting step, so local transport is not important. In this regime, clustering has little effect and $\langle \alpha \rangle \approx 1$. When the system is limited by transport ($Da \gg 1$), the effect of clustering should be larger, and $\langle \alpha \rangle$ should decrease.
2. **The Stokes Number (St).** For a given Damköhler number, the size of the performance drop is controlled by St. Simple theories suggest the penalty is largest when the particle response time (τ_p) matches a fluid timescale (e.g., $St_K \approx 1$). However, the real mechanism may be more complex. It is hypothesized that the performance has a non-monotonic dependence on St, with a minimum at a specific resonant condition. An important goal of this work is to empirically find this dependence across a wide range of Stokes numbers and to find the true resonant condition.

The next sections will test this hypothesis directly with a series of controlled numerical experiments. The final goal is to develop a predictive model for how turbulence-induced clustering affects dissolution performance.

6.1.1. Decoupling Inertia and Transport

To test how system performance depends on both transport and inertia, as described in Section 6.1, a set of targeted numerical experiments was done using the validated framework. The main idea of the design was to separate the effects of particle inertia and turbulent transport.

Two sets of parametric sweeps were done to separate these competing effects:

1. **Mapping the Inertial Performance Penalty.** To study the non-monotonic effect of particle inertia on system performance $\langle \alpha \rangle$, the turbulent flow was kept fixed, and the Stokes number (St) was changed over a wide range. This was done for two different turbulence intensities: a moderate intensity case ($u_{rms} \approx 0.012 \text{ m s}^{-1}$) and a high intensity case ($u_{rms} \approx 0.88 \text{ m s}^{-1}$). This method allowed us to find the resonant Stokes number that gives the worst performance and to see how this resonance changes with turbulence intensity.
2. **Isolating the Role of Turbulent Transport.** To isolate the effect of turbulent transport, a second study was done. In this study, the particle properties were kept constant. With the other parameters fixed, the turbulence intensity, measured by the root mean square velocity u_{rms} , was changed over a wide range. This was done to measure how only the transport intensity affects system performance for a given particle type.

In both studies, the Damköhler number, Da , was also changed over a wide range. This was done by changing the number of particles, n_p , which changes the overall chemical reaction timescale.

All simulations were averaged over multiple runs with different initial particle positions. This was done to get reliable statistics and to reduce errors from the random starting positions.

The Damköhler Number

To measure the balance between reaction and transport, a system-level Damköhler number (Da) is defined. Da is the ratio of the transport timescale, τ_f , to the chemical reaction timescale, τ_{chem} :

$$Da = \frac{\tau_f}{\tau_{chem}} \quad (6.2)$$

For the homogeneous, isotropic turbulence considered here, the relevant transport timescale is the large-eddy turnover time, the time taken by the largest, most energetic eddies to move a scalar across the domain:

$$\tau_f = \frac{L}{u_{rms}} \quad (6.3)$$

where L is the integral length scale (the domain size) and u_{rms} is the root-mean-square velocity of the turbulence.

The chemical reaction timescale, τ_{chem} , is the time the particles need to consume a large part of the reactant. For a multi-particle system, it is defined as:

$$\tau_{chem} = \frac{C_{H^+,bulk}}{R_0} \quad (6.4)$$

where $C_{H^+,bulk}$ is the initial bulk concentration of the reactant, and R_0 is the initial global volumetric reaction rate:

$$R_0 = \dot{n}_{p,0} \cdot n_p \quad (6.5)$$

Here, $\dot{n}_{p,0}$ is the initial molar consumption rate of H^+ for a single isolated particle, and n_p is the particle number density.

It is crucial to distinguish this system-level definition of the chemical timescale from those employed in homogeneous, single-phase reacting flows. In a continuous medium where reactants are uniformly available, τ_{chem} is typically defined based on local, intrinsic kinetic rate constants [65]. However, for a dispersed reactive phase such as this, the reaction is fundamentally heterogeneous, occurring only at the discrete particle-fluid interfaces. Consequently, the overall process rate is governed not by an intrinsic property of the fluid volume, but by the collective chemical demand of all particles within that volume. Our definition, which incorporates the particle number density n_p into the global reaction rate R_0 , correctly captures this multiphase nature. This formulation, consistent with the approach used by Haugen et al. (2018) [30] for discrete reactive particles, ensures that the Damköhler number accurately represents the competition between the rate of turbulent reactant supply and the total rate of reactant consumption by the particle ensemble.

Combining these definitions, the Damköhler number for this study becomes:

$$Da = \frac{\tau_f \cdot (\dot{n}_{p,0} \cdot n_p)}{C_{H^+,bulk}} \quad (6.6)$$

Justification for Modulating Da via Particle Number Density (n_p): The final expression for Da also explains our choice of experimental control. In our studies, we changed the Damköhler number by changing the particle number density, n_p . This method is simple and makes physical sense. The equation shows that Da is directly proportional to n_p . More particles increase the total chemical demand, which makes the chemical timescale, τ_{chem} , shorter and Da larger.

Changing n_p is also a good way to isolate the effect of the total reaction rate. In each study, all other terms in the Da formula— τ_f , $\dot{n}_{p,0}$, and $C_{H^+,bulk}$ —are kept constant. This makes sure that the change from the kinetically-limited regime ($Da \ll 1$) to the transport-limited regime ($Da \gg 1$) is only caused by the change in the total reaction rate from all particles. It is not caused by changes in turbulence or the kinetics of a single particle.

The System Effectiveness Factor

To quantify the system-level impact of clustering, it has been used the time-averaged **effectiveness factor**, $\langle \alpha \rangle$. This metric compares the actual dissolution performance of the system to an idealized baseline, averaged over the main period of the simulation:

$$\langle \alpha \rangle = \left\langle \frac{\dot{m}_{\text{actual}}(t)}{\dot{m}_{\text{ideal}}(t)} \right\rangle_t \quad (6.7)$$

Here, $\dot{m}_{\text{actual}}(t)$ is the actual mass dissolution rate at a given time, taken directly from the multiphase simulation. At each time step t , it is calculated as the sum of the dissolution rates of all particles, p :

$$\dot{m}_{\text{actual}}(t) = \sum_{p=1}^{n_p} \dot{m}_p(C_{\text{local},p}(t), y_p(t)) \quad (6.8)$$

where $\dot{m}_p$ is determined by its internal state vector $y_p(t)$ and the local reactant concentration $C_{\text{local},p}(t)$. This concentration is found at the exact position of the particle inside the non-uniform concentration field.

To create a good baseline for comparison, $\dot{m}_{\text{ideal}}(t)$ is defined as the ideal dissolution rate. This is a hypothetical value. It is calculated at the same time t using the same particles from the simulation. For each particle p , the ideal rate is found by replacing the local concentration with the domain-averaged reactant concentration, $\bar{C}(t)$:

$$\bar{C}(t) = \frac{1}{V} \int_V C(\mathbf{x}, t) dV \quad (6.9)$$

The total ideal rate for the system is then the sum of these individual ideal rates:

$$\dot{m}_{\text{ideal}}(t) = \sum_{p=1}^{n_p} \dot{m}_p(\bar{C}(t), y_p(t)) \quad (6.10)$$

This dynamic baseline is very important. It changes over time to match the true state of all particles, including their sizes and surface properties, and the total reactant depletion. By using the average concentration $\bar{C}(t)$ for every particle, this definition removes the spatial differences caused by clustering, but it keeps all other parts of the system's condition the same. Therefore, the effectiveness factor $\langle \alpha \rangle$ isolates and measures the performance loss that happens only because clustered particles are in regions with lower reactant concentrations than the domain average.


6.1.2. The Roles of Particle Inertia and Turbulent Transport

This section presents the main results from the parametric studies. The analysis first shows the non-monotonic performance penalty caused by particle inertia at a moderate turbulence intensity. This result is then compared to other works to show the competing physical mechanisms. Finally, the analysis shows how this behavior changes with higher turbulence.


The Resonant Penalty

The analysis starts by studying the effect of particle inertia under moderate turbulence ($u_{rms} \approx 0.012 \text{ m s}^{-1}$). Figure 6.1a plots the system effectiveness factor, $\langle \alpha \rangle$, as a function of the Stokes number for two different, transport-limited Damköhler numbers ($\text{Da} = 1$ and $\text{Da} = 2$).

The results clearly show a non-monotonic performance penalty. For both Damköhler numbers, the system performance is initially high in the tracer-like regime ($\text{St} < 10^{-3}$), here the particle doesn't have enough inertia to leave the streamlines, avoiding the formation of clustering. A further increase in Stokes, reach a minimum in the range of $\text{St} \approx 5 \times 10^{-3}$. This point represents the resonant condition where preferential concentration has its most negative impact on the dissolution performance. As St increases more, performance recovers as the system enters the ballistic regime where clustering weakens. In Figure 6.1b is possible to notice the exact same trend, but for a different turbulence intensity. Here



(a) Effectiveness factor $\langle \alpha \rangle$ versus Stokes number St for moderate turbulence ($u_{rms} \approx 0.012 \text{ m s}^{-1}$). The plot shows a non-monotonic performance penalty. The minimum performance is at $St \approx 5 \times 10^{-3}$.




(b) Effectiveness factor $\langle \alpha \rangle$ versus Stokes number St for high turbulence ($u_{rms} \approx 0.88 \text{ m s}^{-1}$). The plot also shows a non-monotonic penalty, with a minimum performance at $St \approx 2 \times 10^{-3}$.


Figure 6.1: Effectiveness factor $\langle \alpha \rangle$ as a function of Stokes number St for different turbulence intensity regimes.

the resonant Stokes number is slightly higher, suggesting a dependency between the resonant Stokes and turbulence intensity. These plots clearly show the effect of the clustering penalty and are the basis for the next part of the analysis.

To check if these results are physically correct, they are compared to the benchmark DNS results from Haugen et al. (2018) [30], shown in Figure 6.2b. The qualitative agreement is excellent. The benchmark data, conducted at a higher Reynolds number of $Re = 2200$, similarly reveals a non-monotonic performance penalty with respect to the Stokes number. Specifically, performance is lowest for intermediate Stokes numbers ($St = 0.3, 1.0$) compared to the near-tracer case ($St = 0.03$). This strong qualitative agreement confirms that our simulation framework accurately captures the essential physics of the resonant clustering penalty. There is a quantitative difference in the location of the minimum performance. The minimum is at $St \approx 0.005$ in this work, but at $St_k \approx 1$ in the work by Haugen et al. This difference is caused by the different definitions of the Stokes number used. This simulation defined the Stokes number based on the large-eddy timescale, while Haugen et al. use the Kolmogorov timescale. The next analysis will show that the results are very similar when this difference is accounted for.



(a) Single-particle Sherwood number (Sh) as a function of Stokes number (St) for $Re = 400$, from Haugen et al. (2018) [30]. The plot demonstrates a monotonic increase in individual particle mass transfer efficiency with inertia.



(b) Benchmark DNS results from Haugen et al. (2018) [30] showing the effect of St at a fixed Reynolds number ($Re = 2200$).

Figure 6.2: Plots from Haugen et al. (2018) [30] showing mass transfer characteristics.

Quantitative Analysis of the Resonant Mechanism

For the moderate intensity case ($u_{rms} \approx 0.012 \text{ m s}^{-1}$), the resonant large-eddy Stokes number was found to be $St_{L,res} \approx 0.005$. This observed value seems to be different from the common hypothesis that maximum clustering happens when the "Stokes number is one" [64].

This difference can be explained by using the dual-Stokes framework from Section 6.1. The hypothesis of resonance at unity is for the **Kolmogorov Stokes number** (St_K), not the large-eddy Stokes number. We can now test this hypothesis directly by calculating the value of St_K that corresponds to our observed resonant condition.

First, we calculate the particle response time at the point of the worst performance. The fluid properties are from Table 6.1:

$$\tau_{p,res} = St_{L,res} \times \tau_f \approx 0.005 \times 1.99 \text{ s} \approx 0.00995 \text{ s} \quad (6.11)$$

Next, it's calculated the corresponding Kolmogorov Stokes number using the fluid's Kolmogorov timescale, $\tau_\eta = 0.0424 \text{ s}$:

$$St_{K,res} = \frac{\tau_{p,res}}{\tau_\eta} = \frac{0.00995 \text{ s}}{0.0424 \text{ s}} \approx 0.23 \quad (6.12)$$

This result is an important finding of this analysis. It proves that the observed resonant penalty, which occurs at a small large-eddy Stokes number ($St_L \approx 0.005$), is driven by a near-resonance between the particle response time and the dissipative, small-scale motions of the turbulence. This gives a Kolmogorov Stokes number of order one ($St_K \approx 0.23$).

Table 6.1: Statistical characteristics of turbulence for the moderate intensity simulation ($u_{rms} \approx 0.012 \text{ m s}^{-1}$).

Parameter	Symbol	Value
Large-eddy turnover time	T_L	1.99 s
Kolmogorov timescale	τ_η	0.0424 s
Integral length scale	L	0.1 m
Kolmogorov length scale	η	0.002 m

The Influence of Turbulence

This section examines how these dynamics change with a significant increase in turbulence intensity. Figure 6.3 plots the effectiveness factor as a function of Damköhler number for several Stokes numbers at high turbulence intensity ($u_{rms} \approx 0.88 \text{ m s}^{-1}$).




Figure 6.3: Effectiveness factor $\langle \alpha \rangle$ versus Damköhler number Da at a high turbulence intensity ($u_{rms} \approx 0.88 \text{ m s}^{-1}$) for several Stokes numbers. Performance gets worse as St is reduced over the tested range

The plot shows that for the tested range of $St = [0.48, 1.45]$, system performance degrades monotonically as the Stokes number is reduced. Based on the resonant behavior seen before at moderate turbulence, this result can be understood as the 'right-hand side' of a similar, but shifted, non-monotonic curve. The data show that for this higher turbulence intensity, the resonant point with the worst performance has shifted to a different value of $St \leq 0.48$. This dependence of the resonant St on the turbulence intensity is an important result for modeling and scaling up these systems.

6.1.2.1. The Competing Effects of Turbulence on System Performance

The previous analysis showed that inertial clustering can reduce the performance in transport-limited systems. This leads to an important engineering question: for a fixed amount of solids in a reactor, what is the global effect of increasing the mixing energy? It might seem that more turbulence is always better because it improves mass transfer. This section shows that the relationship is more complex. It is controlled by the competing effects of turbulence at different scales.

To study only the effect of turbulence, the physical properties of the system were kept constant. This includes the number of particles (n_p), their initial diameter (d_0), and the solids volume fraction (ϕ_s). The turbulence intensity, measured by its root-mean-square velocity (U_{rms}), was the only variable that was changed. This is like a real situation where an engineer changes the mixing power for a given batch. The results in Figure 6.4 show the system's total efficiency as a function of the energy input.




Figure 6.4: Time-averaged effectiveness factor $\langle \alpha \rangle$ as a function of turbulence intensity (U_{rms}) for a fixed system (constant $n_p = 8000$, $d_0 = 50 \mu m$, $\phi_s = 1.0 \times 10^{-3}$). The plot reveals a complex, non-linear but monotonically non-decreasing trend.

The data in Figure 6.4 demonstrate a complex relationship defined by the interplay of several physical mechanisms, all driven by turbulence. The overall trend is the result of a competition between phenomena that enhance and inhibit mass transfer:

1. **Particle-Scale Convective Enhancement:** Turbulence, as deeply discussed, creates relative motion between particles and the fluid. This sharpens concentration gradients at the particle surface and makes the boundary layer thinner. This increases the mass transfer coefficient (or Sherwood number). This effect helps to supply more reactant to each particle and improves the effectiveness factor $\langle \alpha \rangle$.
2. **Inter-Particle Competition (The "Clustering Tax"):** The same turbulence organizes particles into dense clusters. Inside these clusters, particles must compete for a limited local supply of reactants. This competition can create local zones with low reactant concentration. This slows down the reaction and reduces the performance, which lowers $\langle \alpha \rangle$.
3. **Cluster-Scale Convective Transport:** Turbulence acts on individual particles, but it also acts on the clusters themselves. Turbulence can bring fresh reactants to the outside of a cluster or even inside it. This brings more reactant to the starved regions. This mechanism is more important at high turbulence intensities.

The shape of the curve in Figure 6.4 can be understood by looking at three different regimes. In each regime, a different mechanism is the most important:

- **Regime I: Convection-Dominated Enhancement** ($U_{rms} \lesssim 0.05 \text{ m s}^{-1}$): At this low intensity, the main limit is the slow mass transfer to each particle. The curve has a positive slope because increasing U_{rms} improves the particle-scale convection. At these energy levels, clustering is weak, so the "clustering tax" is small. The total result is a fast improvement in performance.
- **Regime II: Clustering and Diminishing Returns** ($0.05 \lesssim U_{rms} \lesssim 0.2 \text{ m s}^{-1}$): In this regime, the slope of the curve becomes much smaller, almost flat. This plateau means the two main competing effects are in balance. The particle-scale convection continues to get better with U_{rms} , but this benefit is now canceled by the negative effect of clustering. The "clustering tax" becomes important, and adding more mixing energy gives almost no benefit.
- **Regime III: Cluster-Scale Transport Dominance** ($U_{rms} > 0.2 \text{ m s}^{-1}$): After the plateau, the effectiveness factor starts to increase quickly again. This means the main transport mechanism has changed. The turbulence is now strong enough to bring reactants to the clusters. This overcomes the mass transfer limits inside the clusters that created the plateau. The system is no longer limited by transport within clusters. Instead, it is driven by the rate of transport to clusters, and this rate increases a lot with U_{rms} .

This analysis gives an important insight for process design. It changes the simple idea that more mixing is always better. The relationship between mixing energy and performance is not a simple curve with one optimum, but rather is a series of bottlenecks. The data show that for this system, increasing the mixing power can still be a good strategy, but only if the energy is high enough to reach the third regime where cluster-scale transport is the most important. This changes the idea of preferential concentration. It is not always negative. It is a limit that depends on the regime and must be understood to make the process better.

6.1.3. Evidence of Localized Reactant Starvation

The previous analysis showed that preferential concentration causes a large performance penalty when transport is limited ($Da > 1$). However, this result does not show the exact mechanism. This section studies the system at the micro-scale to test the main hypothesis: that the bad performance is caused by **localized reactant starvation** inside the dense particle clusters.

First, the change in the particle distribution is shown. Figure 6.5 shows how a uniform particle distribution quickly changes into a network of thin filaments. This shows that preferential concentration is a dynamic process. The particles are moved and line up along the fluid's Lagrangian Coherent Structures (LCS) by turbulence.

Visualizing the Clustered State and Its Impact on the Concentration Field

It starts by visually characterizing the evolution of the particle field. Figure 6.5 illustrates how an initially uniform particle distribution rapidly transforms into a highly anisotropic, filamentary network. This dynamic clustering confirms that preferential concentration is an active, time-dependent process: particles are continuously advected, stretched, and folded by turbulent structures, aligning preferentially along the fluid's Lagrangian Coherent Structures (LCS).

This changing structure is important for the local reactant concentration. As the dense particle filaments form, their combined chemical reaction changes the concentration field around them. This creates local areas with low reactant concentration. The connection between clustering and mass transfer limitation is what causes the bad performance of the total system. It links the micro-scale effects to the system-level behavior.

To test the connection between clustering and reactant availability, Figure 6.6 compares the particle distribution and the reactant concentration field at the same time. The anti-correlation is very clear: regions with a high number of particles are the same as regions with low reactant concentration. This is direct visual evidence for reactant starvation. The dense particle filaments create channels of low reactant concentration in the fluid. This shows that the transport limitation is a local effect, not a global depletion of reactants.




Figure 6.5: Snapshots of the particle distribution at three different times for a case with $St \approx 1$. An initial homogeneous distribution quickly changes into thin, filament-like structures because of preferential concentration. These structures are not static and line up with the Lagrangian Coherent Structures (LCS).




Figure 6.6: Comparison of (a) the particle distribution and (b) the reactant (H^+) concentration field at one time in the transport-limited regime. There is a strong anti-correlation. Regions with many particles are the same as regions with low reactant concentration. This is direct visual proof of the reactant starvation mechanism.

Synthesis of Microscopic Evidence

The microscopic study gives a clear and consistent result. Preferential concentration arranges particles into changing, filament-like structures (Figure 6.5). Inside these dense filaments, particles compete for the reactant. This creates small areas with very low concentration (Figure 6.6). This confirms the main hypothesis: the bad performance of the whole system, which is measured by the effectiveness factor $\langle \alpha \rangle$, is caused by local reactant starvation.

6.1.4. Application to Reactor Design

The previous analysis showed the basic non-dimensional conditions where inertial clustering causes a large performance penalty. The main finding is that this penalty is largest at a resonant Stokes number, St_{crit} , which also depends on the flow Reynolds number. This dimensionless result is important for science. But for practical use in industrial reactor design and optimization, it must be translated into a dimensional model.

This section connects the Damköhler number to the practical parameter of solids loading. Then, it creates a quantitative engineering model. This model changes the resonant Stokes number into a design chart. This chart connects the mixer power (P/V) to the critical particle diameter ($d_{p,crit}$) that is most affected by clustering. Finally, a parametric sensitivity analysis is shown. It uses this model to understand how design choices, like impeller type, agitator diameter, and reactor volume, change the risk of clustering in a system.

6.1.4.1. From Damköhler Number to Solids Loading

The previous analysis showed that the Damköhler number is the key parameter that decides if a system's performance is limited by transport. But, the Damköhler number is a complex dimensionless group. It is not a direct input for an engineer. To make these results useful, the Damköhler number must be connected to a practical parameter: the solids loading, or **solid volume fraction** (ϕ_s).

For the monodisperse system in this study, the solid volume fraction is the particle number density multiplied by the volume of one particle:

$$\phi_s = n_p \cdot \frac{\pi}{6} d_p^3 \quad (6.13)$$

The Damköhler number is also directly proportional to the particle number density:

$$Da = n_p \cdot K, \quad \text{where} \quad K = \frac{\tau_f \cdot \dot{n}_{p,0}}{C_{H^+, \text{bulk}}} \quad (6.14)$$

Because both Da and ϕ_s are directly proportional to the particle number density, n_p , a change in one causes a direct change in the other for a given particle size and flow condition. This linear relationship makes it possible to change the results from the Damköhler number analysis into a new type that is useful for industrial process design.

Figure 6.7 shows this practical view. It plots the time-averaged effectiveness factor, $\langle \alpha \rangle$, as a function of the solid volume fraction, ϕ_s , for several Stokes numbers.




Figure 6.7: The effect of solids loading (ϕ_s) on system effectiveness for different particle inertia regimes (Stokes numbers). This plot shows the practical performance penalty of clustering. Note the clear difference in robustness between the high-inertia ($St \approx 9.645$) and resonant-inertia ($St \approx 0.010$) particles.

The plot shows that the system's ability to handle more solids depends a lot on particle inertia. For highly inertial particles ($St \approx 9.645$, purple line), the system is very robust. The effectiveness factor stays high even when the solid volume fraction is close to 10%. This shows that for these 'ballistic' particles, clustering is weak. The system can have a high throughput without a large transport penalty.

In contrast, particles near the resonant clustering condition ($St \approx 0.001$, light blue line) are very sensitive to solids loading. The effectiveness factor drops quickly, by more than 50% at a solid fraction of only 0.1%. This shows a large performance penalty. A small increase in particle concentration causes strong local reactant starvation, which reduces the total process efficiency.

This analysis gives an important design principle: **the maximum allowed solids loading in a reactor is not constant. It is a strong function of the particle Stokes number.** For this reason, a scale-up strategy that only tries to maximize solids loading to get higher throughput is not a good idea. This approach

can cause very bad performance if the system operates near a Stokes number that causes strong particle clustering.

For good process optimization, an integrated method is needed. The fluid mixing conditions, which set the fluid timescale (τ_f), and the particle properties, which set the particle relaxation time (τ_p), must be considered at the same time to find the operational Stokes number. Only with this analysis is it possible to choose an optimal solids loading. This choice must balance the process throughput with the performance penalties from turbulence-induced clustering.

6.1.4.2. Theoretical Basis for the Engineering Framework

The engineering model is based on a series of known physical equations. These equations connect the large-scale mixer parameters to the small-scale turbulence properties that control the particle behavior. The key connection is the critical clustering condition, which is based on the resonant Stokes number:

$$St_{crit} = \frac{\tau_{p,crit}}{\tau_f} \quad (6.15)$$

where $\tau_{p,crit}$ is the momentum response time of the critical particle size, and τ_f is the characteristic turnover time of the large eddies. The value of St_{crit} is determined from the fundamental numerical study as a function of the agitator Reynolds number, $St_{crit} = f(Re)$.

To use this condition, standard mixer design equations are used to calculate the turbulence timescales from engineering parameters [31]. First, the power (P) from an impeller is given by:

$$P = N_p \rho_f N^3 D^5 \quad (6.16)$$

where N_p is the dimensionless **Power Number** characteristic of the impeller geometry, ρ_f is the fluid density, N is the agitator rotational speed, and D is the agitator diameter. In a well-mixed, turbulent system at steady state, this power input is balanced by the turbulent energy dissipation rate, ε :

$$\varepsilon = \frac{P}{\rho_f V} \quad (6.17)$$

where V is the reactor volume. The timescale of the large eddies, τ_f , which are about the size of the impeller, depends on this dissipation rate:

$$\tau_f = \left(\frac{D^2}{\varepsilon} \right)^{1/3} \quad (6.18)$$

The momentum response time of a spherical particle is given by:

$$\tau_p = \frac{\rho_p d_p^2}{18\mu_f} \quad (6.19)$$

where ρ_p and d_p are the particle density and diameter, and μ_f is the fluid dynamic viscosity. By substituting Equations 6.18 and 6.19 into the critical condition (6.15) and solving for the particle diameter, we obtain the final equation:

$$d_{p,crit} = \sqrt{\frac{18\mu_f}{\rho_p} St_{crit} \tau_f} \quad (6.20)$$

This equation is the main part of the design model. It is used to directly calculate the critical particle diameter that is most likely to cluster for any mixer design and operating conditions.

Figure 6.8 shows the critical diameter zone that should be avoided when designing a mixer. The plot is for a fixed tank volume (V), impeller diameter (D), and Power Number (N_p).

6.1.4.3. Parametric Sensitivity Analysis of Mixer Design

The engineering model from Equation (6.20) is now used to study how the critical clustering zone is affected by different design variables. This analysis shows the trade-offs in reactor design and gives quantitative help for process optimization.




Figure 6.8: Critical clustering diameter in different power zones for a fixed tank volume (V), impeller diameter (D) and Power Number (N_p).




Figure 6.9: Parametric Sensitivity Analysis of Mixer Design.

Effect of Impeller Type (Power Number, N_p)

The choice of impeller, which is described by the Power Number N_p , has a large effect on the hydrodynamics of the system. Figure 6.9a shows how changing N_p affects the critical clustering diameter.

The analysis shows a clear trend: for a constant specific power (P/V), impellers with a higher Power Number move the critical particle diameter to larger values. The physical reason for this is the relationship between power and rotational speed. To have the same power input, a less efficient impeller (higher N_p) must turn at a lower speed N , as shown in Equation (6.16). This lower speed creates larger and slower turbulent eddies. This increases the large-eddy turnover time τ_f . To meet the resonant condition from Equation (6.15), a longer particle response time τ_p is needed. This means the critical particle diameter $d_{p,crit}$ is larger.

Engineering Implication: Choosing an impeller type is a trade-off. Impellers with high efficiency (low N_p) use less energy. But, they must turn faster to give the same power. This creates stronger turbulence, which creates a risk of clustering for smaller particles. For processes with fine powders where clustering is a problem, a less efficient, high- N_p impeller that turns at a lower speed can be a better choice. It will use more energy, but it can reduce the performance loss from clustering.

Effect of Agitator Diameter (D)

The agitator diameter is an important design parameter. It controls the size of the large eddies. Figure 6.9b shows how it affects the critical clustering zone.

The results show that for the same specific power (apart from the unclear region in the middle of the plot), a larger agitator diameter makes the critical particle diameter larger. This is because of turbulence scaling. A larger agitator diameter D creates larger eddies. As shown in Equation (6.18), for a constant energy dissipation rate ε , the eddy turnover time τ_f scales with $D^{2/3}$. This longer fluid timescale needs a longer particle response time τ_p for resonance, which means a larger particle size.

Engineering Implication: For processes that dissolve large or dense particles, a large agitator that turns slowly is better than a small agitator that turns fast, if they use the same power. The larger impeller creates a flow with longer timescales, so it is less likely to cause resonant clustering for large particles.

Effect of Reactor Volume (V)

Reactor volume is the main parameter for process scale-up. The analysis in Figure 6.9c shows that scaling a reactor by keeping the specific power constant can change how the system clusters. The data show that for a constant specific power (P/V), a larger reactor volume makes the critical clustering zone move to smaller particle diameters. The physical reason is that a constant P/V in a larger volume V needs a higher total power input P . If the impeller is the same, this extra power needs a higher agitator speed N . The higher speed creates stronger turbulence with a shorter eddy turnover time τ_f . This moves the resonant condition to smaller particles.

Engineering Implication: Scaling up a process by keeping the specific power constant is not enough and can be a bad strategy if clustering reduces performance. This analysis shows that this scale-up method does not keep the particle-turbulence interactions the same. A larger reactor with the same P/V will have a different clustering risk. This is an important point that must be considered when designing large industrial systems. This model gives a tool to measure this change during the scale-up process.

6.2. Analysis of Leaching Kinetics

The previous sections showed how external hydrodynamics and particle clustering affect the dissolution process. This section now studies the chemical kinetics that controls this process. The Shrinking Core Model (SCM), described in Chapter 4, uses kinetic constants from experimental work [50]. These constants have some uncertainty. For this reason, it is important to analyze how sensitive the system is to these parameters to understand the kinetics of the LiCoO_2 leaching process. This analysis finds the rate-determining steps and the feedback mechanisms that control the efficiency of lithium extraction.

6.2.1. Parametric Analysis of Reaction Kinetics

The goal of this analysis is to understand the connection between the reaction rates and the amount of available reactant. Three different numerical studies were done using the `ShrinkingCoreModel` mode of the simulator. To study only the chemical effects and not the transport limits, all simulations started with a uniform particle distribution. This makes sure that any differences in the results are only because of changes in the kinetic parameters, and not because of the non-uniform hydrodynamics that were studied before.

The study is done in three steps:

1. **Analysis I: Kinetic Hierarchy and Passivation.** First, a basic study finds which kinetic constants are the most important. A baseline simulation is compared to other simulations where each of the four main kinetic constants (k_1, k_2, k_3, k_4) is changed by $\pm 50\%$. This analysis finds the first rate-determining steps.
2. **Analysis II: The Self-Limiting Nature of the Acid Attack Pathway.** Second, a detailed study looks at the direct acid attack pathway. The kinetic constant k_1 is changed over a wide range. This is done to measure its non-linear effect and to understand the self-limiting feedback mechanism.
3. **Analysis III: The Dual Role of Hydrogen Peroxide.** Third, this analysis studies the effect of a reducing agent (H_2O_2) concentration to the system performance. The initial concentration of hydrogen peroxide is changed in each simulation. This helps to understand its two roles: as a driver for the reaction and as a way to stop the passivation mechanism.

6.2.2. Initial Kinetic Hierarchy and Passivation Dynamics

The initial stage of the leaching process ($t \rightarrow 0$) is purely kinetically-controlled. In this regime, reactant concentrations are at their maximums, and the inhibiting product layer has not yet formed. Figure 6.10 shows how the lithium extraction efficiency changes with time. It compares a baseline case to other cases where the four SCM kinetic constants are changed by $\pm 50\%$.

Figure 6.10 shows that the extraction efficiency for all cases increases linearly with time. This linear




Figure 6.10: Sensitivity of Li extraction efficiency over the first 10 seconds of the reaction to $\pm 50\%$ changes in the SCM kinetic constants. The slope of each line is a direct measure of the initial reaction rate.

behavior is because of the initial conditions. The reactant concentrations, the total surface area, and the product layer thickness (which is zero) are all constant at the start. For this reason, the total reaction rate is also constant. The slope of each line is a direct measure of the initial reaction rate for that set of kinetic parameters.

An analysis of these slopes shows which reaction pathways are most important. The effect of each kinetic constant on the initial rate gives this order of sensitivity:

$$\text{Sensitivity Hierarchy: } k_3 > k_1 \gg k_2, k_4 \quad (6.21)$$

The peroxide-assisted pathway (k_3) is dominant, as its variation produces the largest change in reaction rates. The direct acid attack pathway (k_1) is the second most important.

An important result is that changing the passivation kinetics has no effect at the start. The curves for the baseline case and the cases with reduced k_2 and k_4 are exactly the same. This result shows that the model is physically correct. The constants k_2 and k_4 control the dissolution of the Co_3O_4 product crust. Because this crust is a product, it cannot affect the process at $t = 0$ when there is no crust. The numerical model correctly simulates this sequence: first the core dissolves, then the passivating layer forms and dissolves.

6.2.3. Kinetic Analysis of the k_1 Reaction Pathway

Analysis I showed that the direct acid attack pathway (k_1) is an important part of the total reaction rate. This section studies this pathway in detail. A study was done where k_1 was changed over a wide range. This shows a complex, self-limiting mechanism with diminishing returns.

This is the clear kinetic benefit of a faster direct acid attack. But, this benefit has a cost. The same reaction creates the passivating Co_3O_4 product. Figure 6.11 and 6.12 shows that cases with higher k_1 values create this crust much faster. This shows the main problem: a faster reaction also speeds up the mechanism that stops it.




Figure 6.11: The initial extraction rate as a function of the k_1 multiplier. This confirms that the process starts in a kinetically-controlled regime.




Figure 6.12: Average moles of passivating Co₃O₄ crust per particle. Higher k_1 values cause faster and larger crust formation.

The "Rate Decay Index" in Figure 6.13, defined as the ratio of the final to the initial reaction rate, measures the self-passivation of the system. For k_1 multipliers up to about 20, the rate is stable. After this point, the index drops sharply, and at a 200x multiplier, the final rate is only 7% of its initial value. This sharp decay happens because the system shifts from an initial, kinetically-controlled regime, dominated by the fast acid attack (k_1), to a later regime that is limited by reactant transport through a newly formed passivating layer. As a result, the dissolution of this crust itself (governed by k_2 and k_4) becomes the new process bottleneck. This bottleneck explains the strong diminishing returns seen in the final extraction efficiency (Figure 6.14), where increasing the k_1 multiplier from 100 to 200 gives only a very small improvement in the total extraction.




Figure 6.13: The Rate Decay Index, which measures the "choking" effect of passivation. A sharp drop shows the start of a diffusion-dominated regime.




Figure 6.14: Relative extraction efficiency at $t = 10$ s. The plot shows clear diminishing returns for high k_1 multipliers.

The change in extraction over time is shown in Figure 6.15. It shows this dynamic conflict. Higher k_1 values have a steeper initial slope, but their curves become flatter as the passivating crust forms. When plotted against the absolute value of the kinetic constant (Figure 6.16), the final extraction efficiency does not decrease. Instead, it reaches a regime of strong diminishing returns, getting closer to a transport-limited maximum. This study shows that the direct acid attack pathway is controlled by a strong, self-regulating feedback loop. Making this pathway faster gives large initial gains, and at the same time, future bottlenecks.

6.2.4. The Dual Role of Hydrogen Peroxide

This analysis examines the effect of an important external parameter: the initial concentration of hydrogen peroxide, [H₂O₂]. The results show a complex, non-linear response that reveals the dual function of H₂O₂. It acts not only as a primary reactant but also as a process stabilizer that prevents the system from failing due to self-passivation.

Figure 6.17a shows the initial reaction rate at different [H₂O₂] concentrations. At low concentrations ([H₂O₂] < 2000 mol/m³), the rate is limited by peroxide and increases quickly. At high concentrations ([H₂O₂] > 4000 mol/m³), the rate reaches a plateau. This plateau is clear evidence that the rate is no




Figure 6.15: Time evolution of Li extraction efficiency for different k_1 multipliers. Higher k_1 causes faster initial extraction but also an earlier slowdown.






Figure 6.16: Final Li extraction efficiency versus the absolute k_1 rate constant (log scale). The flat part shows a transport-limited maximum.



(a) The initial reaction rate ($t \rightarrow 0$) as a function of initial $[H_2O_2]$. The plateau clearly shows that the system is now limited by the other reactant.



(b) The final extraction efficiency at $t = 10$ s. The curve shows that efficiency continues to improve even when the initial rate is saturated.

Figure 6.17: The dual role of hydrogen peroxide. The initial rate (6.17a) becomes saturated because of a bottleneck from the other reactant. However, the final efficiency (6.17b) continues to improve. This shows the benefit of stopping passivation over time.

longer limited by H_2O_2 , but by the transport of the other reactant, H^+ . The initial reaction is so fast that it consumes H^+ ions at the particle surface faster than diffusion can supply them through the developing porous crust.

A clear difference is seen when comparing the initial rate (Figure 6.17a) with the final extraction efficiency at $t = 10$ s (Figure 6.17b). Even in the plateau region where the initial rate is constant, the final efficiency continues to increase. This behavior shows the **dual benefit** of hydrogen peroxide:

1. **Kinetic Driver:** H_2O_2 directly accelerates the dissolution of the $LiCoO_2$ core through the k_3 pathway. This explains the initial rate increase, but its effect is limited when H^+ transport becomes the bottleneck.
2. **Active Passivation Suppressant:** H_2O_2 is also the key reactant that dissolves the passivating Co_3O_4 crust via the k_4 pathway. This is an effect that happens over the entire process. It prevents the build-up of a diffusion barrier that would otherwise slow down the already-limited transport of H^+ .

6.2.5. Synthesis and Engineering Implications

The last three analyses have broken down the system's kinetics. They lead to a clear set of results that are important for both industrial process optimization and for academic modeling.

The main results can be summarized as follows:

- **Kinetic Hierarchy and Control:** There is a clear order of importance for the kinetics. The peroxide-assisted pathway (k_3) controls the initial rate. However, the kinetics of crust dissolution (k_2, k_4), which are not important at $t = 0$, become more important over time to let reactants reach the core.
- **The Self-Limiting Acid Pathway:** The direct acid attack pathway (k_1) is controlled by a strong

self-limiting feedback loop. Making this pathway much faster gives high initial rates, but it has the big problem of slowing down quickly because of faster passivation. This leads to diminishing returns.

- **Peroxide as a Process Stabilizer:** Hydrogen peroxide has two benefits. Its most important job over time is not just to make the reaction faster, but to actively remove the passivating layer. This makes the process more efficient at longer time.

These findings are all connected. The central problem is a fundamental mismatch in the kinetic rates. The passivating Co_3O_4 layer is generated by the fast core-dissolution pathways (k_1, k_3) much faster than it is removed by the slower crust-dissolution pathways (k_2, k_4). This kinetic imbalance is the root cause of the self-limiting behavior. It explains why managing the crust, rather than simply increasing the initial reaction rate, is the key to process optimization.

Conclusions and Recommendations

This thesis studied the dissolution of black mass materials in a turbulent fluid. This is a critical and slow step in the hydrometallurgical recycling of lithium-ion batteries. The goal was to go beyond simple reactor-scale models and to develop a predictive understanding of the basic mechanisms that control the process efficiency at the particle level. This final chapter summarizes the main contributions of this research.

7.1. Summary of Key Scientific Findings

This section describes the main scientific contributions of this thesis. The contributions are organized to answer the research questions from Chapter 1. Each finding is supported by the results shown in the previous chapters.

7.1.1. Validation of Transport Physics (RQ1)

A necessary first step for the main investigation was the validation of the numerical model's core transport algorithms. This validation was done using the simplified "Salt Case", which is a single-component dissolution problem described in Chapter 5. This case has well-defined theoretical limits. The model's physical accuracy was confirmed because it was able to reproduce the theoretical scaling laws for two different mass transfer regimes.

In the low-turbulence, diffusion-dominated regime ($Sh \approx 2$), the model correctly predicted a linear decrease of the average particle surface area over time. This result is in exact agreement with first principles (Figure 5.2). For the high-turbulence, convection-dominated regime ($Sh \propto Re_p^{1/2}$), the model correctly captured the predicted linear decrease of the mass dissolution rate (Figure 5.3). The correct simulation of these different behaviors showed that the model has high physical accuracy. This successful validation gave a strong basis for the later studies of more complex problems.

7.1.2. The Competing Effects of Turbulence on System Performance (RQ2)

The study of how preferential concentration affects global reaction rates showed a complex relationship that is controlled by the competing effects of turbulence. The system's performance, measured by the effectiveness factor $\langle \alpha \rangle$, is very different from the behavior of an ideal, uniform system. The main finding is that the total effect of turbulence on the multi-particle system is a competition between three different physical mechanisms:

1. **Particle-Scale Convective Enhancement:** Turbulence increases the relative velocity between particles and the fluid. This improves the mass transfer for each particle, which, if no clustering is involved, improves system performance.
2. **Inter-Particle Competition (The Clustering Tax):** The same turbulence organizes particles into dense, filament-like clusters. This preferential concentration causes strong local competition for reactants, which reduces performance.

3. **Cluster-Scale Turbulent Transport:** At high turbulence intensities, turbulence can bring fresh reactants to the starved regions inside the clusters. This mechanism reduces the negative effect of clustering.

The interaction of these mechanisms controls the system's non-linear performance response to increasing turbulence intensity (Figure 6.4). At low intensities, particle-scale enhancement is the most important effect, which leads to an initial increase in performance. As turbulence gets stronger, the "clustering tax" becomes important and creates diminishing returns, causing the performance to reach a plateau. At very high intensities, cluster-scale transport becomes strong enough to overcome the local competition, which leads to another improvement in performance.

The study also confirmed that the performance penalty from clustering is largest at a resonant Kolmogorov-based Stokes Number of order one ($St_K \approx 0.23$), in clear agreement with literature benchmark studies on clustering [64]. This result connects the macroscopic system behavior directly to the basic physics of particle interaction with the smallest scales of turbulence.

Direct evidence for the "clustering tax" mechanism comes from the simulated concentration fields. Visualizations show a strong spatial anti-correlation between the particle distribution and the reactant concentration (Figure 6.6). Dense particle clusters are in the same regions as areas of strong reactant depletion. This links the macroscopic performance penalty to its microscopic cause, which is localized reactant starvation.

7.1.3. Deconstructing the Leaching Kinetic Architecture (RQ3)

The multi-step Shrinking Core Model was studied to find the most important kinetic pathways and feedback mechanisms that control the LiCoO_2 leaching process. The analysis showed that the most important kinetic step changes during the process. There is not just one single rate-determining step. At the start of the reaction ($t \rightarrow 0$), the peroxide-assisted dissolution pathway (k_3) is the most important for the process rate. As the reaction continues, a passivating product layer forms (Co_3O_4 crust). The dissolution of this crust (pathways k_2 and k_4) then becomes very important. It allows reactants to reach the unreacted core, even though these pathways were not important at the start. A key finding is increasing the k_1 rate constant increases the initial extraction rate, but it also speeds up the formation of the passivating Co_3O_4 crust. This connection causes the reaction to slow down a lot and leads to diminishing returns at high acid concentrations.

Hydrogen peroxide (H_2O_2) has two important jobs. Its main effect is seen over the whole process, not just at the start. First, it speeds up the reaction through the k_3 pathway. However, this effect is limited by the transport of H^+ . More importantly, H_2O_2 makes the process more efficient because it dissolves the passivating crust through the k_4 pathway. This continuous cleaning action stops the system from slowing itself down. It is the main reason for the slightly higher final extraction efficiencies. This explains the difference between the initial rate that becomes flat and the final efficiency that slowly keeps improving (Figure 6.17).

This analysis leads to an important conclusion. Increasing the rates of the main dissolution reactions (k_1 and k_3) is not the solution to speed up the overall process. The results show that an important rate limiting step is the slow dissolution of the passivating Co_3O_4 layer. This step is controlled by the rate constants k_2, k_4 . The model shows that the value of k_4 for H_2O_2 is very low. As a result, the passivating layer forms faster than it is removed, and it becomes the bottleneck for the entire process. Therefore, future work should not focus only on improving the main reaction kinetics. Instead, it should focus on finding ways to increase the rate of crust dissolution. This could be done by using new reducing agents that are more effective at removing the passivating layer.

7.2. Implications for Industrial Process Design

This section uses the scientific mechanisms found in this work to create a practical framework for the design and optimization of industrial multiphase reactors. The findings challenge common engineering rules and offer a more physical approach to process intensification and scale-up.

7.2.1. Defining the Optimal Energy Input for Mixing

An important conclusion of this research is that the simple idea "more energy is always better" is wrong. It is bad for both cost and physics. The relationship between mixing energy input, represented by the turbulent velocity fluctuation u_{rms} , and process efficiency, measured by the effectiveness factor $\langle \alpha \rangle$, is non-linear (Figure 6.4). This behavior, which is caused by the competing mechanisms from Section 7.1.2, creates an optimal operating window. Operating outside this window is not efficient. At low energy input, the process is limited by particle-scale convection. In the plateau, it is slowed down by the "clustering tax". Only at very high energy input transport overcome this penalty.

The practical recommendation is clear: it is necessary to find this performance curve for a system. This helps to find the point where the benefit of more mixing is smaller than the cost of energy and the penalty from clustering.

7.2.2. A Predictive Framework for Clustering Risk and Scale-Up

An important contribution of this work is a practical engineering model based on the physics of resonant clustering (St_{crit}). The model connects the particle Stokes number to standard mixer design equations. It predicts the critical particle diameter, $d_{p,crit}$, that is most likely to have bad performance from clustering for a given set of operating conditions (Figure 6.8). This model gives direct, quantitative help for process design and scale-up.

Impeller Selection The model shows an important trade-off in impeller design. Impellers with high efficiency (low Power Number N_p) must turn faster to give a target power input. This higher speed moves the clustering risk to smaller particles. For this reason, for processes with fine powders, a less efficient impeller (high N_p) that turns at a lower speed can be a better choice to avoid the resonant clustering penalty (Figure 6.8).

Scale-Up Strategy The analysis shows that simple scale-up rules, like keeping a constant specific power (P/V), are wrong. For a given P/V , a larger reactor needs a higher impeller speed, which changes the turbulent timescales. This change moves the critical clustering diameter, $d_{p,crit}$, to smaller particle sizes (Figure 6.8). A process that works well in the lab can fail when scaled up if the new operating conditions move the system into the resonant penalty zone. This model should be used to make sure that the dynamic similarity, which is given by the operational Stokes number, is the same at different scales.

7.2.3. The Optimal Reducing Agent

This research shows that an optimization strategy based only on maximizing the acid concentration is incorrect. This approach is not optimal because it leads to self-passivation. The key to an efficient process is to correctly manage the dual role of hydrogen peroxide. While H_2O_2 is a primary reactant for the dissolution reaction (the k_3 pathway), its more important function is to suppress the passivating layer through the k_4 pathway.

This leads to a clear conclusion for process design. The concentration of hydrogen peroxide must be high enough to perform two functions: first, to act as a primary reactant for core dissolution, and second, to continuously dissolve the passivating product layer. The practical recommendation would be to add enough H_2O_2 not just to drive the main reaction, but to prevent the process from becoming self-limiting. However, this solution is difficult in practice. This is because the rate constant for crust dissolution, k_4 , is naturally much smaller than the rate constants for the main reactions, k_1 and k_3 . Therefore, a better alternative is to find a more efficient reducing agent. This agent should be selected for its ability to dissolve the crust, even if it is less effective at accelerating the core dissolution.

Another possible solution is to use a mechanical method to break this crust. For example, after a thick crust has formed, the particles could pass through a mechanical shredder. This would break the particles into smaller pieces, expose the unreacted core, and reduce the negative effect of the crust on the total process rate.

7.3. Limitations of the Current Study

The numerical model in this thesis uses some simplifying assumptions. These assumptions were needed to study the specific physical effects for this research. It is important to understand these assumptions because they show the limits of this study and give ideas for future work.

The main simplifications are in the model for the turbulent flow and its interaction with the particles. There are two main simplifications:

- First, the model uses a static, pre-computed DNS velocity field. This creates a **frozen flow** environment. This method makes it possible to study the effect of a fixed turbulent field on the particles, without the flow also changing.
- Second, the model uses a strict **one-way momentum coupling**. Particles are moved by the fluid, but the fluid is not affected by the particles. This means the effect of **turbulence modulation** is not included. For example, the drag from particles in dense clusters can reduce the turbulent energy. This simplification limits the model's use to the dilute or semi-dilute systems ($\Phi_v \ll 10^{-3}$) studied in this work. Some simulations in this study exceed this limit. In these cases, the model is expected to predict correct physical trends, but with lower accuracy for the quantitative values. The one-way coupling is the main limitation for using the model for denser systems, which are common in industry.

All simulations were also done on a two-dimensional slice of a three-dimensional, isotropic turbulent flow. Because of this, the model does not include **vortex stretching**. This is an important mechanism in the turbulent energy cascade that moves energy from large to small scales. This could change how particles spread and the shape of the filamentary clusters. The 2D approach was a practical choice to reduce the computational cost. This made it possible to do the many parametric studies in this work, which would have been too slow in 3D.

The model also simplifies the properties of the solid particles to make the calculations easier. Specifically, all particles are assumed to be perfect **spheres**. The simulation tracks a single particle size, which makes it a **monodisperse system**, instead of a full particle size distribution (PSD). The model also does not include changes in particle porosity or complex changes in shape, like fragmentation, during dissolution. Real industrial "black mass" is a complex material with particles of different sizes and shapes that are not spheres. Because of these simplifications, the model should be used to understand physical trends, not to get exact predictions for a specific industrial process.

Finally, the entire model assumes **isothermal conditions**. This means several effects are not included: the heat from the reaction, the change of reaction rates with temperature (Arrhenius relation), and the change of fluid properties (like viscosity) with temperature. This simplification was made to focus only on the chemo-mechanical transport phenomena. In industrial reactors, heat can change over the process. In these cases, non-isothermal effects could be important and could create strong thermal feedback loops.

7.4. Recommendations for Future Work

The limitations from the last section show a clear path for future research. The next recommendations describe a plan to remove the main assumptions of the current model. This will make the model more predictive and closer to real industrial systems.

7.4.1. Targeted Experimental Validation

The numerical predictions from this thesis are clear hypotheses that should be tested with experiments. This research should have two goals:

1. **To validate the main chemo-physical mechanisms**, like the non-monotonic Stokes penalty and the localized reactant starvation. Example tools are: Particle Image Velocimetry (PIV) and Planar Laser-Induced Fluorescence (PLIF).
2. **To validate the practical engineering design model**, specifically the predicted critical clustering diameter ($d_{p,crit}$). Techniques like high-speed imaging could directly measure particle clustering and connect it to the total dissolution rates under different mixer parameters.

A single experimental program that studies both the basic physics and the engineering model would connect the numerical predictions of this thesis to real industrial use.

7.4.2. Enhancing Fidelity with Dynamic and Higher Resolution Turbulence

This study uses a pre-computed DNS flow field from the Johns Hopkins Turbulence Database (JHTDB). This method gives high-resolution turbulence data for a low computational cost. The next logical step is to improve the model by using more features of the JHTDB dataset to remove the current simplifications. As said, one of the main simplification in this work is the "frozen turbulence" assumption. This means that particles move through a static flow field. Two areas for future work can remove this limitation:

1. **Use Higher-Resolution DNS Grids:** The JHTDB has velocity fields with much finer grids than the one used in this study. Running the simulations on these grids would resolve the small-scale eddies that are similar in size to the particles. This would show more accurately how these small structures affect reactant transport to the particle surface.
2. **Implement a Time-Evolving Flow Field:** The most important improvement is to replace the static field with a dynamic one. The JHTDB has time-series data that can be used to create a flow field that changes over time. In this setup, particles would move through the complete life cycle of turbulent eddies. This would allow the study of the process history for each particle. For example, it would be possible to see if particles trapped in vortices dissolve slower than others. This method is more realistic than a static field, but it still avoids the high cost of a fully coupled simulation.

These two improvements would use the full capability of the pre-computed DNS method. They would give a more accurate understanding of the interaction between turbulent transport and surface reactions.

7.4.3. Extension to Three-Dimensional Simulations

It is recommended to do some targeted three-dimensional simulations. This will measure the effect of the 2D domain assumption on the results. It is not necessary to re-run the entire parametric study. A focused study on a few important cases is enough, like the resonant Stokes number case where clustering is strongest. The goal of this work would be to study the role of vortex stretching and other 3D turbulent structures on particle clustering and on the effectiveness factor $\langle \alpha \rangle$. This study would measure the error from the 2D simplification. It would also check if using 2D for large parametric studies is a good choice.


7.4.4. Adding Realistic Particle Properties

To make the model's predictions better for real "black mass", future versions must include more realistic particle properties. This can be done with several research projects:

- **Polydisperse Systems:** The current model can be connected to a Population Balance Model (PBM). This would allow the simulation to track a full Particle Size Distribution (PSD) and how it changes over time.
- **Non-Spherical Morphology:** Models for non-spherical particles can be added. This needs to include drag laws that depend on orientation and more complex surface area calculations. This can be done using shape factors to model the real shape of the particles.
- **Dynamic Internal Structure:** A more advanced model for the inside of the particle can be developed. In this model, the porosity of the product crust would be a variable. This variable would be calculated from the rates of crust formation (k_1) and dissolution (k_2, k_4). This would create a more realistic, non-linear diffusion resistance inside the particle.

7.4.5. Development of a Non-Isothermal Model

An important next step is to develop a non-isothermal model. This would capture the thermodynamic effects that are common in industrial reactors. This needs an extra transport equation for energy (temperature) on the Eulerian grid. This new equation must be coupled with the reaction kinetics solver. This is done by using temperature-dependent (Arrhenius) rate constants for all four reactions ($k_i(T)$). The main goal would be to study thermal feedback loops, like local heating from exothermic reactions. This heating could speed up the local kinetics. The current model does not capture these effects.



Shrinking Core Model Rate Equations

A.1. Derivation of the Unified Rate Equation

This appendix shows the step-by-step derivation of the overall rate equation for the Shrinking Core Model (SCM), which is Equation (3.4) in Chapter 3. The main idea is the pseudo-steady-state assumption. This means the total molar rate of the fluid reactant, $\dot{N}_A$ (in mol/s), is constant through the three resistance layers, which are in series.

The derivation defines the concentration drop across each layer. This drop is written in terms of $\dot{N}_A$ and a layer-specific resistance, similar to Ohm's Law.

Step 1: Mass Transfer Through the Outer Fluid Film

The molar rate of reactant A from the bulk fluid (C_{bulk}) to the particle's outer surface ($C_{surface}$ at radius r_p) is controlled by convective mass transfer:

$$\dot{N}_A = k_L(4\pi r_p^2)(C_{bulk} - C_{surface}) \quad (A.1)$$

where k_L is the film mass transfer coefficient. The concentration drop across this film can be written as:

$$\Delta C_{film} = C_{bulk} - C_{surface} = \dot{N}_A \cdot \left(\frac{1}{k_L(4\pi r_p^2)} \right) \quad (A.2)$$

The term in parentheses is the resistance of the fluid film.

Step 2: Diffusion Through the Product Layer

The reactant must then diffuse through the porous product layer from the outer surface to the unreacted core (radius r_c , concentration C_{core}). For diffusion in a spherical shell, Fick's first law is:

$$\dot{N}_A = -D_{eff}A(r)\frac{dC_A}{dr} = -D_{eff}(4\pi r^2)\frac{dC_A}{dr} \quad (A.3)$$

where D_{eff} is the effective diffusivity. To find the concentration drop, we separate variables and integrate from the core to the outer surface:

$$\dot{N}_A \int_{r_c}^{r_p} \frac{dr}{4\pi r^2} = -D_{eff} \int_{C_{core}}^{C_{surface}} dC_A \quad (A.4)$$

$$\frac{\dot{N}_A}{4\pi} \left[-\frac{1}{r} \right]_{r_c}^{r_p} = -D_{eff} [C_A]_{C_{core}}^{C_{surface}} \quad (A.5)$$

$$\frac{\dot{N}_A}{4\pi} \left(\frac{1}{r_c} - \frac{1}{r_p} \right) = D_{eff}(C_{surface} - C_{core}) \quad (A.6)$$

Solving for the concentration drop across the product layer gives:

$$\Delta C_{\text{product layer}} = C_{\text{surface}} - C_{\text{core}} = \dot{N}_A \cdot \left(\frac{r_p - r_c}{4\pi D_{eff} r_p r_c} \right) \quad (\text{A.7})$$

The term in parentheses is the resistance of the product layer.

Step 3: Surface Reaction at the Core

Assuming a first-order irreversible reaction at the surface of the unreacted core, the rate of consumption of A is:

$$\dot{N}_A = k_s(4\pi r_c^2)C_{\text{core}} \quad (\text{A.8})$$

where k_s is the first-order surface reaction rate constant. The concentration drop driving the reaction is from C_{core} to an assumed zero concentration:

$$\Delta C_{\text{reaction}} = C_{\text{core}} - 0 = \dot{N}_A \cdot \left(\frac{1}{k_s(4\pi r_c^2)} \right) \quad (\text{A.9})$$

The term in parentheses is the resistance of the surface reaction.

Step 4: Summation of Resistances

The total concentration drop from the bulk fluid to the point of reaction is the sum of the individual drops:

$$C_{\text{bulk}} = (C_{\text{bulk}} - C_{\text{surface}}) + (C_{\text{surface}} - C_{\text{core}}) + (C_{\text{core}} - 0) \quad (\text{A.10})$$

Substituting the expressions for each concentration drop from Equations (A.2), (A.7), and (A.9):

$$C_{\text{bulk}} = \dot{N}_A \left(\frac{1}{k_L(4\pi r_p^2)} \right) + \dot{N}_A \left(\frac{r_p - r_c}{4\pi D_{eff} r_p r_c} \right) + \dot{N}_A \left(\frac{1}{k_s(4\pi r_c^2)} \right) \quad (\text{A.11})$$

Factoring out the constant molar rate $\dot{N}_A$ and solving for it yields the final expression:

$$\dot{N}_A = \frac{C_{\text{bulk}}}{\frac{1}{k_L(4\pi r_p^2)} + \frac{r_p - r_c}{D_{eff}(4\pi r_p r_c)} + \frac{1}{k_s(4\pi r_c^2)}} \quad (\text{A.12})$$

This completes the derivation of the governing rate equation, which matches Equation (3.4) in the main text.

A.2. Derivation of the Core Shrinkage Rate Equation

This section derives the governing differential equation for the radius of the unreacted core, $r_c(t)$, as presented in Equation (3.5). The derivation connects the molar rate of the fluid reactant A to the rate of consumption of the solid reactant B.

Step 1: Relate Moles of Solid to Core Radius

The number of moles of the solid reactant B in the unreacted core, N_B , is the product of the core's volume (V_c) and the solid's molar density, $\rho_{B,molar}$ (in mol/m³):

$$N_B = V_c \cdot \rho_{B,molar} = \left(\frac{4}{3}\pi r_c^3 \right) \rho_{B,molar} \quad (\text{A.13})$$

Step 2: Relate Rate of Molar Change to Shrinkage Rate

The rate of change of moles of B in the core is found by differentiating with respect to time using the chain rule:

$$\frac{dN_B}{dt} = \frac{d}{dt} \left(\frac{4}{3}\pi r_c^3 \rho_{B,molar} \right) = (4\pi r_c^2) \rho_{B,molar} \frac{dr_c}{dt} \quad (\text{A.14})$$

Since the core is shrinking, its contents are being consumed. The rate of consumption of B is therefore $-\frac{dN_B}{dt}$.

Step 3: Apply Reaction Stoichiometry

The reaction stoichiometry is given by $A(\text{fluid}) + b \cdot B(\text{solid}) \rightarrow \text{Products}$. The rate of consumption of solid B is stoichiometrically related to the rate of consumption of fluid reactant A, $\dot{N}_A$:

$$-\frac{dN_B}{dt} = b \cdot \dot{N}_A \quad (\text{A.15})$$

Step 4: Combine and Finalize

By equating the two expressions for the rate of consumption of B, we can solve for the core shrinkage rate:

$$-\left(4\pi r_c^2\right)\rho_{B,molar}\frac{dr_c}{dt} = b \cdot \dot{N}_A \quad (\text{A.16})$$

$$\frac{dr_c}{dt} = -\frac{b \cdot \dot{N}_A}{(4\pi r_c^2)\rho_{B,molar}} \quad (\text{A.17})$$

Finally, substituting the full expression for $\dot{N}_A$ from section A.1 gives the complete governing equation for the rate of core shrinkage:

$$\frac{dr_c}{dt} = -\frac{b}{(4\pi r_c^2)\rho_{B,molar}} \left(\frac{C_{bulk}}{\frac{1}{k_L(4\pi r_p^2)} + \frac{r_p - r_c}{D_{eff}(4\pi r_p r_c)} + \frac{1}{k_s(4\pi r_c^2)}} \right) \quad (\text{A.18})$$

Simplifying the denominator by distributing the $(4\pi r_c^2)$ term yields the final form:

$$\frac{dr_c}{dt} = -\frac{bC_{bulk}/\rho_{B,molar}}{\frac{r_c^2}{r_p^2 k_L} + \frac{r_c(r_p - r_c)}{r_p D_{eff}} + \frac{1}{k_s}} \quad (\text{A.19})$$

This result matches Equation (3.5) in Chapter 3.

B

Derivation of the Quiescent Concentration Profile and $Sh=2$ Limit

This appendix derives from first principles the steady-state concentration profile around a single spherical particle in a quiescent fluid. It then shows how this profile leads directly to the theoretical Sherwood number limit of $Sh = 2$. This is a fundamental result in mass transfer theory.

B.1. Governing Equation and Assumptions

The derivation starts from the general conservation of species equation (Fick's Second Law):

$$\frac{\partial C}{\partial t} + \nabla \cdot \mathbf{J} = 0 \quad (\text{B.1})$$

where the flux vector $\mathbf{J}$ includes both diffusive and convective components. For this specific case, we apply two key simplifying assumptions:

1. **Steady State:** The concentration profile is time-invariant, thus $\frac{\partial C}{\partial t} = 0$.
2. **Quiescent Medium:** The fluid is perfectly still, meaning there is no convective transport ($\mathbf{u} = 0$). The flux is therefore purely diffusive, given by Fick's First Law: $\mathbf{J} = -D\nabla C$.

Substituting these assumptions into the conservation equation yields Laplace's equation for concentration:

$$\nabla^2 C = 0 \quad (\text{B.2})$$

B.2. Solution in Spherical Coordinates

Given the geometry of the problem, we use spherical coordinates. We assume the system is spherically symmetric, meaning the concentration C is a function of only the radial distance r , i.e., $C = C(r)$. In this case, the Laplacian operator simplifies to:

$$\nabla^2 C = \frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{dC}{dr} \right) \quad (\text{B.3})$$

Setting this equal to zero gives the governing Ordinary Differential Equation (ODE):

$$\frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{dC}{dr} \right) = 0 \quad (\text{B.4})$$

This ODE can be solved by integrating twice with respect to r . The first integration yields:

$$r^2 \frac{dC}{dr} = A \quad (\text{B.5})$$

where A is the first constant of integration. Rearranging and integrating a second time gives the general solution for the concentration profile:

$$C(r) = -\frac{A}{r} + B \quad (\text{B.6})$$

where B is the second constant of integration.

B.3. Application of Boundary Conditions

The constants A and B are determined by applying the physical boundary conditions (BCs) of the system:

- **BC 1 (Particle Surface):** At the particle surface ($r = R_p$), the concentration is the saturation concentration, C_s .

$$C(r = R_p) = C_s$$

- **BC 2 (Far Field):** At an infinite distance from the particle ($r \rightarrow \infty$), the concentration returns to the bulk fluid concentration, C_{bulk} .

$$C(r \rightarrow \infty) = C_{bulk}$$

Applying BC 2 to the general solution (Eq. B.6):

$$C_{bulk} = -\frac{A}{\infty} + B \implies B = C_{bulk}$$

Next, applying BC 1 and substituting $B = C_{bulk}$:

$$C_s = -\frac{A}{R_p} + C_{bulk} \implies A = -R_p(C_s - C_{bulk})$$

Substituting the solved constants A and B back into the general solution yields the final expression for the steady-state concentration profile around the sphere:

$$C(r) = C_{bulk} + (C_s - C_{bulk}) \frac{R_p}{r} \quad (\text{B.7})$$

This is the same as Eq. 2.3 in the main text.

B.4. Derivation of the Sherwood Number Limit ($Sh = 2$)

The Sherwood number (Sh) is defined in terms of the mass transfer coefficient, k_L . We can derive k_L by equating the fundamental diffusive flux with the phenomenological definition of mass transfer.

The total molar dissolution rate, $\dot{n}$ (in mol/s), from the particle surface is the diffusive flux at the surface ($J_{r=R_p}$) multiplied by the particle's surface area ($A_p = 4\pi R_p^2$):

$$\dot{n} = A_p \cdot J_{r=R_p} = (4\pi R_p^2) \left(-D \frac{dC}{dr} \right)_{r=R_p} \quad (\text{B.8})$$

From our first integration step, we know $\frac{dC}{dr} = \frac{A}{r^2}$. Substituting the expression for A :

$$\frac{dC}{dr} = -\frac{R_p(C_s - C_{bulk})}{r^2}$$

Evaluating this gradient at the surface $r = R_p$:

$$\left(\frac{dC}{dr} \right)_{r=R_p} = -\frac{C_s - C_{bulk}}{R_p}$$

The total molar rate is therefore:

$$\dot{n} = (4\pi R_p^2) \left(-D \left(-\frac{C_s - C_{bulk}}{R_p} \right) \right) = 4\pi D R_p (C_s - C_{bulk}) \quad (\text{B.9})$$

The molar rate is also defined using the mass transfer coefficient k_L :

$$\dot{n} = k_L A_p (C_s - C_{bulk}) = k_L (4\pi R_p^2) (C_s - C_{bulk}) \quad (\text{B.10})$$


By equating Eq. B.9 and Eq. B.10, we can solve for k_L :

$$k_L (4\pi R_p^2) (C_s - C_{bulk}) = 4\pi D R_p (C_s - C_{bulk}) \implies k_L = \frac{D}{R_p}$$

Finally, we use the definition of the Sherwood number ($Sh = k_L d_p / D$, where $d_p = 2R_p$):

$$Sh = \frac{(D/R_p) \cdot (2R_p)}{D} = 2 \quad (\text{B.11})$$

This completes the derivation, proving that for pure diffusion from an isolated spherical particle into a quiescent, infinite medium, the Sherwood number is exactly 2.



Numerical Solver for the MassTransferOnly Mode

The fundamental governing equation describes the rate of change of a particle's radius, r , as a function of the mass transfer rate:

$$\frac{dr}{dt} = -\frac{Sh(Re_p)D_{mol}}{2\rho_p r}(C_s - C_{bulk}) \quad (C.1)$$

The key numerical challenge arises from the $1/r$ term on the right-hand side, which causes the ODE to become infinitely stiff as $r \rightarrow 0$. A standard explicit time-stepping method would be forced to use impractically small time steps to maintain stability.

To overcome this stiffness, a semi-implicit numerical scheme is employed. The method applies a first-order, A-stable Implicit Euler discretization to the governing dr/dt equation. The term "semi-implicit" is used here because, as will be shown, a subsequent linearization will allow for a direct algebraic solution, avoiding a full non-linear iterative solve.

The Implicit Euler discretization of Eq. (C.1) is formulated by evaluating the right-hand side at the future time step $n + 1$:

$$\frac{r^{n+1} - r^n}{\Delta t_{sim}} = -\frac{Sh(r^{n+1})D_{mol}(C_s - C_{bulk})}{2\rho_p r^{n+1}} \quad (C.2)$$

This equation is a non-linear algebraic equation for the unknown new radius r^{n+1} , as r^{n+1} appears in both the denominator and implicitly within the $Sh(r^{n+1})$ term. To make this equation tractable without resorting to a slow, iterative root-finder, a physically justified linearization assumption is introduced: it is assumed that for a sufficiently small time step Δt_{sim} , the Sherwood number does not change significantly. Therefore, it is approximated using the value from the beginning of the time step: $Sh(r^{n+1}) \approx Sh(r^n)$. This is a robust approximation, as Sh is a relatively weak function of r (proportional to $\sqrt{r}$), and the change in r is small within a single Δt_{sim} .

Applying this linearization to Eq. (C.2) yields:

$$\frac{r^{n+1} - r^n}{\Delta t_{sim}} \approx -\frac{Sh(r^n)D_{mol}(C_s - C_{bulk})}{2\rho_p r^{n+1}} \quad (C.3)$$

This is now a much simpler non-linear equation for r^{n+1} . To solve it, we perform an algebraic rearrangement to cast it into a standard polynomial form. Multiplying both sides by $2\rho_p r^{n+1}\Delta t_{sim}$ gives:

$$2\rho_p r^{n+1}(r^{n+1} - r^n) \approx -Sh(r^n)D_{mol}(C_s - C_{bulk})\Delta t_{sim}$$

Expanding the left-hand side and dividing the entire equation by $2\rho_p$ yields:

$$(r^{n+1})^2 - r^n r^{n+1} \approx -\frac{Sh(r^n)D_{mol}(C_s - C_{bulk})\Delta t_{sim}}{2\rho_p}$$

Finally, moving all terms to one side, we arrive at a standard quadratic equation for the unknown r^{n+1} :

$$(r^{n+1})^2 - r^n r^{n+1} + \left[\frac{Sh(r^n)D_{mol}(C_s - C_{bulk})\Delta t_{sim}}{2\rho_p} \right] = 0 \quad (C.4)$$

Equation (C.4) is a quadratic equation of the form $ax^2 + bx + c = 0$, where:

- $x = r^{n+1}$ (the unknown new radius)
- $a = 1$
- $b = -r^n$ (the previous radius)
- $c = A_{coeff} = \left[\frac{Sh(r^n)D_{mol}(C_s - C_{bulk})\Delta t_{sim}}{2\rho_p} \right]$

This exact form is what is solved non-iteratively in the MATLAB code using the standard quadratic formula: $x = (-b \pm \sqrt{b^2 - 4ac})/2a$.

A critical implementation detail is the selection of the correct physical root from the $\pm$ possibilities. The code implements the '+' root:

$$r^{n+1} = \frac{r^n + \sqrt{(r^n)^2 - 4A_{coeff}}}{2} \quad (C.5)$$

This choice is physically motivated. In the limit of a very small time step ($\Delta t_{sim} \rightarrow 0$), the coefficient $A_{coeff} \rightarrow 0$, and the discriminant $(r^n)^2 - 4A_{coeff} \rightarrow (r^n)^2$.

- The '+' root yields $(r^n + \sqrt{(r^n)^2})/2 = (r^n + r^n)/2 = r^n$. This is the correct physical behavior: for a zero time step, there should be no change in radius.
- The '-' root would yield $(r^n - \sqrt{(r^n)^2})/2 = (r^n - r^n)/2 = 0$. This is an unphysical solution, suggesting the particle vanishes instantly.

Therefore, the '+' root is rigorously identified as the only physically meaningful solution. The implementation includes a final robustness check, which check against potential floating-point errors that could cause the term under the square root to become slightly negative, ensuring a real-valued result for the new radius and preventing simulation failure. This semi-implicit, linearized method results in a non-iterative quadratic solve, providing an exceptionally fast, unconditionally stable, and first-order accurate update for the particle radius.

D

Derivation of the SCM Core Surface Concentration Polynomial

This appendix provides a detailed derivation of the special cubic polynomial used to solve for the reactant concentration at the core-crust interface, C_{core,H^+} , within the Shrinking Core Model (SCM). This value is not an independent variable but is determined by a quasi-steady-state balance between the rate of reactant diffusion to the core and the rate of reactant consumption by the surface and crust reactions.

D.1. The Physical Principle: Quasi-Steady-State Flux Balance

The derivation begins with the fundamental physical principle that, under quasi-steady-state conditions, the total molar flux of the primary reactant (H^+) diffusing from the bulk fluid to the core surface must be equal to the total molar flux of reactant being consumed by all reactions.

$$\dot{N}_{diffusive} = \dot{N}_{reactive} \quad (D.1)$$

where $\dot{N}$ represents the total molar rate (moles/s).

D.2. Formulation of the Diffusive Flux

The diffusive flux through the porous spherical crust is modeled using Fick's first law for a spherical shell. The total molar rate of diffusion is given by:

$$\dot{N}_{diffusive} = 4\pi D_{eff,H} \frac{R_p r_c}{R_p - r_c} (C_{bulk,H^+} - C_{core,H^+}) \quad (D.2)$$

where $D_{eff,H}$ is the effective diffusivity of H^+ in the crust, R_p is the particle's outer radius, r_c is the unreacted core radius, and C_{bulk,H^+} and C_{core,H^+} are the bulk and core surface concentrations, respectively. The model for the effective diffusivity, $D_{eff,H}$, is based on the empirical correlation developed by Cerrillo-Gonzalez et al. [50] to account for the effects of porosity and tortuosity in the product layer.

For convenience, we define a diffusive transport coefficient, A_0 :

$$A_0 = \frac{4\pi D_{eff,H} R_p r_c}{R_p - r_c} \quad (D.3)$$

Equation (D.2) can then be written more compactly as:

$$\dot{N}_{diffusive} = A_0 (C_{bulk,H^+} - C_{core,H^+}) \quad (D.4)$$

D.3. Formulation of the Reactive Flux

The total rate of H^+ consumption is the sum of consumption from all four SCM reactions (r_1, r_2, r_3, r_4), weighted by their respective stoichiometric coefficients for H^+ :

$$\dot{N}_{reactive} = 2r_1 + 6r_2 + 3r_3 + 6r_4 \quad (D.5)$$

The individual reaction rates are defined as:

$$\begin{aligned} r_1 &= k_1 C_{core, H^+} (4\pi r_c^2) \\ r_2 &= k_2 n_{crust} (C_{bulk, H^+})^{2/3} \\ r_3 &= k_3 (C_{core, H^+})^{1/3} (C_{bulk, H_2O_2})^2 (4\pi r_c^2) \\ r_4 &= k_4 n_{crust} (C_{bulk, H^+})^{2/3} (C_{bulk, H_2O_2})^2 \end{aligned}$$

where k_i are the reaction rate constants, n_{crust} is the moles of crust material, and $SA_{core} = 4\pi r_c^2$ is the core surface area.

D.4. Algebraic Rearrangement to Polynomial Form

By equating the diffusive and reactive fluxes from Eq. (D.1) and substituting the detailed expressions, we obtain:

$$A_0(C_{bulk, H^+} - C_{core, H^+}) = 2r_1 + 6r_2 + 3r_3 + 6r_4 \quad (D.6)$$

Substituting the rate laws and moving all terms to one side yields:

$$\begin{aligned} A_0 C_{core, H^+} + 2k_1 (4\pi r_c^2) C_{core, H^+} + 3k_3 (4\pi r_c^2) (C_{bulk, H_2O_2})^2 (C_{core, H^+})^{1/3} \\ + (6k_2 n_{crust} + 6k_4 n_{crust} (C_{bulk, H_2O_2})^2) (C_{bulk, H^+})^{2/3} - A_0 C_{bulk, H^+} = 0 \end{aligned} \quad (D.7)$$

To transform this into a standard polynomial, we introduce the substitution $x = (C_{core, H^+})^{1/3}$, which implies $x^3 = C_{core, H^+}$. Grouping terms by powers of x :

$$\begin{aligned} \underbrace{(A_0 + 2k_1 SA_{core}) x^3}_P + \underbrace{(3k_3 SA_{core} (C_{bulk, H_2O_2})^2) x}_Q \\ + \underbrace{[(6k_2 n_{crust} + 6k_4 n_{crust} (C_{bulk, H_2O_2})^2) (C_{bulk, H^+})^{2/3} - A_0 C_{bulk, H^+}]}_R = 0 \end{aligned} \quad (D.8)$$

This is a cubic equation of the form $Px^3 + Qx + R = 0$, where the coefficients P, Q, R are functions of the known physical parameters and state variables. This is the exact equation solved in the code to determine the core surface concentration at each time step.

E

Derivation of the Analytical Jacobian for the SCM

This appendix provides a detailed derivation of the analytical Jacobian matrix, $\mathbf{J}$, for the Shrinking Core Model (SCM) ODE system. The use of an analytical Jacobian is critical for the performance and stability of the Rosenbrock stiff ODE solver detailed in Section 4.5.2.

E.1. System Definition

The state of a single particle is described by the vector $\mathbf{y} = [y_1, y_2]^T$, where $y_1 = r_c$ (the unreacted core radius) and $y_2 = n_{crust}$ (the moles of solid crust material, Co_3O_4). The evolution of the system is governed by the ODE system $d\mathbf{y}/dt = \mathbf{f}(\mathbf{y})$, where $\mathbf{f} = [f_1, f_2]^T$:

$$f_1(\mathbf{y}) = \frac{dr_c}{dt} = -\frac{M_{\text{LiCoO}_2}}{\rho_{\text{LiCoO}_2}} \frac{r_1(\mathbf{y}) + r_3(\mathbf{y})}{4\pi r_c^2} \quad (\text{E.1})$$

$$f_2(\mathbf{y}) = \frac{dn_{crust}}{dt} = \frac{1}{6}r_1(\mathbf{y}) - r_2(\mathbf{y}) - r_4(\mathbf{y}) \quad (\text{E.2})$$

The reaction rates, r_i , are functions of the state $\mathbf{y}$ and local fluid concentrations. Specifically, r_1 and r_3 depend on the concentration at the core surface, C_{core,H^+} , which is itself an implicit function of $\mathbf{y}$.

E.2. Jacobian Structure

The Jacobian matrix is defined as:

$$\mathbf{J} = \frac{\partial \mathbf{f}}{\partial \mathbf{y}} = \begin{pmatrix} \frac{\partial f_1}{\partial y_1} & \frac{\partial f_1}{\partial y_2} \\ \frac{\partial f_2}{\partial y_1} & \frac{\partial f_2}{\partial y_2} \end{pmatrix} = \begin{pmatrix} J_{11} & J_{12} \\ J_{21} & J_{22} \end{pmatrix} \quad (\text{E.3})$$

E.3. The Chain Rule and Implicit Differentiation

The primary challenge in deriving the Jacobian elements is that the rates r_1 and r_3 depend on C_{core,H^+} , which is not an independent variable but is determined by a non-linear algebraic flux-balance equation, $G(C_{core,H^+}, y_1, y_2) = 0$. To find the derivatives of C_{core,H^+} with respect to the state variables, we use the implicit function theorem:

$$\frac{\partial C_{core,H^+}}{\partial y_j} = -\frac{\partial G / \partial y_j}{\partial G / \partial C_{core,H^+}} \quad (\text{E.4})$$

This is the central mathematical tool for the derivation. The derivatives of the reaction rates r_i with respect to a state variable y_j are then found using the chain rule, for example:

$$\frac{\partial r_1}{\partial y_j} = \frac{\partial r_1}{\partial C_{core,H^+}} \frac{\partial C_{core,H^+}}{\partial y_j} + \frac{\partial r_1}{\partial SA_{core}} \frac{\partial SA_{core}}{\partial y_j} \quad (\text{E.5})$$

where $SA_{core} = 4\pi r_c^2 = 4\pi y_1^2$ is the surface area of the core.

E.4. Derivation of Jacobian Components

The following sections derive each component of the Jacobian matrix by applying the principles above.

E.4.1. Derivatives of Reaction Rates

First, we find the partial derivatives of the reaction rates r_i with respect to the state variables $y_1 = r_c$ and $y_2 = n_{crust}$.

$$\begin{aligned}\frac{\partial r_1}{\partial y_1} &= k_1 \left(\frac{\partial C_{core,H^+}}{\partial y_1} SA_{core} + C_{core,H^+} (8\pi y_1) \right) \\ \frac{\partial r_1}{\partial y_2} &= k_1 \frac{\partial C_{core,H^+}}{\partial y_2} SA_{core} \\ \frac{\partial r_2}{\partial y_1} &= 0 \quad ; \quad \frac{\partial r_2}{\partial y_2} = k_2 (C_{bulk,H^+})^{2/3} \\ \frac{\partial r_3}{\partial y_1} &= k_3 (C_{bulk,H_2O_2})^2 \left(\frac{1}{3} C_{core,H^+}^{-2/3} \frac{\partial C_{core,H^+}}{\partial y_1} SA_{core} + C_{core,H^+}^{1/3} (8\pi y_1) \right) \\ \frac{\partial r_3}{\partial y_2} &= k_3 (C_{bulk,H_2O_2})^2 \left(\frac{1}{3} C_{core,H^+}^{-2/3} \frac{\partial C_{core,H^+}}{\partial y_2} SA_{core} \right) \\ \frac{\partial r_4}{\partial y_1} &= 0 \quad ; \quad \frac{\partial r_4}{\partial y_2} = k_4 (C_{bulk,H^+})^{2/3} (C_{bulk,H_2O_2})^2\end{aligned}$$

E.4.2. Assembly of Final Jacobian Elements

Using the product and quotient rules on Eqs. (E.1) and (E.2):

$$\begin{aligned}J_{11} &= \frac{\partial f_1}{\partial y_1} = -\frac{M_{LiCoO_2}}{\rho_{LiCoO_2}} \left[\frac{(\frac{\partial r_1}{\partial y_1} + \frac{\partial r_3}{\partial y_1})(4\pi y_1^2) - (r_1 + r_3)(8\pi y_1)}{(4\pi y_1^2)^2} \right] \\ J_{12} &= \frac{\partial f_1}{\partial y_2} = -\frac{M_{LiCoO_2}}{\rho_{LiCoO_2}} \frac{(\frac{\partial r_1}{\partial y_2} + \frac{\partial r_3}{\partial y_2})}{4\pi y_1^2} \\ J_{21} &= \frac{\partial f_2}{\partial y_1} = \frac{1}{6} \frac{\partial r_1}{\partial y_1} - \frac{\partial r_2}{\partial y_1} - \frac{\partial r_4}{\partial y_1} = \frac{1}{6} \frac{\partial r_1}{\partial y_1} \\ J_{22} &= \frac{\partial f_2}{\partial y_2} = \frac{1}{6} \frac{\partial r_1}{\partial y_2} - \frac{\partial r_2}{\partial y_2} - \frac{\partial r_4}{\partial y_2}\end{aligned}$$

E.5. Special Case: The No-Crust Limit

A critical edge case occurs at the beginning of the simulation when the crust thickness $\tau = R_p - r_c$ is effectively zero. In this limit, the diffusive resistance is nil, and the concentration at the core surface is equal to the bulk concentration, i.e., $C_{core,H^+} = C_{bulk,H^+}$. Since C_{bulk,H^+} is a parameter and not a function of the state variables y , we have:

$$\frac{\partial C_{core,H^+}}{\partial y_1} = 0 \quad \text{and} \quad \frac{\partial C_{core,H^+}}{\partial y_2} = 0 \quad (\text{for } \tau \rightarrow 0) \quad (E.6)$$

Substituting these into the derivatives of the reaction rates simplifies them greatly:

$$\begin{aligned}\frac{\partial r_1}{\partial y_1} &= k_1 C_{bulk,H^+} (8\pi y_1) & \frac{\partial r_1}{\partial y_2} &= 0 \\ \frac{\partial r_3}{\partial y_1} &= k_3 (C_{bulk,H_2O_2})^2 C_{bulk,H^+}^{1/3} (8\pi y_1) & \frac{\partial r_3}{\partial y_2} &= 0\end{aligned}$$

The derivatives of r_2 and r_4 remain unchanged. This leads to a much simpler Jacobian matrix for the no-crust case, which is implemented in the code to ensure numerical stability at $t = 0$. For instance, it is clear that J_{12} and J_{22} will evaluate to simpler expressions involving only the derivatives of r_2 and r_4 , as the terms dependent on r_1 and r_3 vanish. This robust handling of the initial condition is essential for the solver's stability.

Bibliography

- [1] U. Saleem, A. Wilhelms, J. Sottmann, H. K. Knuutila, and S. Bandyopadhyay, "Direct lithium extraction (DLE) methods and their potential in Li-ion battery recycling", en, *Separation and Purification Technology*, vol. 361, p. 131 315, Jul. 2025, issn: 13835866. doi: 10.1016/j.seppur.2024.131315.
- [2] R. Kumar *et al.*, "Sustainable recovery of high-valued resources from spent lithium-ion batteries: A review of the membrane-integrated hybrid approach", *Chemical Engineering Journal*, vol. 470, p. 144 169, Aug. 2023, issn: 1385-8947. doi: 10.1016/j.cej.2023.144169.
- [3] Y. Yao, M. Zhu, Z. Zhao, B. Tong, Y. Fan, and Z. Hua, "Hydrometallurgical Processes for Recycling Spent Lithium-Ion Batteries: A Critical Review", en, *ACS Sustainable Chemistry & Engineering*, vol. 6, no. 11, pp. 13 611–13 627, Nov. 2018, issn: 2168-0485, 2168-0485. doi: 10.1021/acssuschemeng.8b03545.
- [4] M. Cerrillo-Gonzalez, M. Villen-Guzman, C. Vereda-Alonso, J. Rodriguez-Maroto, and J. Paz-Garcia, "Towards Sustainable Lithium-Ion Battery Recycling: Advancements in Circular Hydrometallurgy", en, *Processes*, vol. 12, no. 7, p. 1485, Jul. 2024, issn: 2227-9717. doi: 10.3390/pr12071485.
- [5] "Global Supply Chains of EV Batteries", en, *IEA.org*, 2022.
- [6] "EV Battery Supply Chain Sustainability", en, *IEA.org*, 2024.
- [7] A. Racu and J. Poliscanova, "From Waste to Value", *Transport & Environment*, Dec. 2024.
- [8] European Union, *Regulation (EU) 2023/1542 of the European Parliament and of the Council of 12 July 2023 concerning batteries and waste batteries, amending Directive 2008/98/EC and Regulation (EU) 2019/1020 and repealing Directive 2006/66/EC*, 2023.
- [9] E. Asadi Dalini, G. Karimi, S. Zandevakili, and M. Goodarzi, "A Review on Environmental, Economic and Hydrometallurgical Processes of Recycling Spent Lithium-ion Batteries", en, *Mineral Processing and Extractive Metallurgy Review*, vol. 42, no. 7, pp. 451–472, Oct. 2021, issn: 0882-7508, 1547-7401. doi: 10.1080/08827508.2020.1781628.
- [10] A. Chagnes and B. Pospiech, "A brief review on hydrometallurgical technologies for recycling spent lithium-ion batteries", en, *Journal of Chemical Technology & Biotechnology*, vol. 88, no. 7, pp. 1191–1199, Jul. 2013, issn: 0268-2575, 1097-4660. doi: 10.1002/jctb.4053.
- [11] T. H. Wines, "Recycling of Lithium-Ion-Batteries: Hydrometallurgy Process", PALL Corporation, Tech. Rep., Jun. 2024.
- [12] The MathWorks, Inc., *MATLAB*, Natick, Massachusetts, United States, 2024.
- [13] O. Levenspiel, "Chemical reaction engineering. Hauptbd.", en, in 3. ed, Num Pages: 668, New York Weinheim: Wiley, 1999, isbn: 978-0-471-25424-9.
- [14] Y. Li *et al.*, "A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence", *Journal of Turbulence*, vol. 9, N31, Jan. 2008, arXiv:0804.1703 [physics], issn: 1468-5248. doi: 10.1080/14685240802376389.
- [15] C. Crowe, J. Schwarzkopf, M. Sommerfeld, and Y. Tsuji, *Multiphase Flows with Droplets and Particles, Second Edition*. Taylor & Francis, 2011, isbn: 978-1-4398-4050-4.
- [16] R. B. Bird, W. E. Stewart, and E. N. Lightfoot, *Transport Phenomena*, en. John Wiley & Sons, Dec. 2006, Google-Books-ID: L5FnNIIaGfcC, isbn: 978-0-470-11539-8.
- [17] H. Schlichting and K. Gersten, *Boundary-Layer Theory*, en. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, isbn: 978-3-662-52917-1 978-3-662-52919-5. doi: 10.1007/978-3-662-52919-5.
- [18] A. J. Bard and L. R. Faulkner, *Electrochemical methods: fundamentals and applications*, en, 2. edition. New York Weinheim: Wiley, 2001, isbn: 978-0-471-04372-0.
- [19] R. Clift, J. R. Grace, and M. E. Weber, *Bubbles, Drops, and Particles*, en. Courier Corporation, Jan. 2005, Google-Books-ID: UUrOmD8niUQC, isbn: 978-0-486-44580-9.
- [20] N. Frossling, "Über die Verdunstung fallender Tropfen", 1938.
- [21] W. E. Ranz and W. R. Marshall, "Evaporation from drops: Part 2", en, *Chemical Engineering Progress*, vol. 48, no. 4, pp. 173–180, 1952, issn: 0360-7275.
- [22] E. Karchniwy, N. E. L. Haugen, A. Klimanek, Ø. Langørgen, and S. Sladek, "The effect of turbulence on mass transfer in solid fuel combustion: RANS model", en, *Combustion and Flame*, vol. 227, pp. 65–78, May 2021, issn: 00102180. doi: 10.1016/j.combustflame.2020.12.040.
- [23] F. Montagnaro, "Sherwood (Sh) Number in Chemical Engineering Applications—A Brief Review", en, *Energies*, vol. 17, no. 17, p. 4342, Aug. 2024, issn: 1996-1073. doi: 10.3390/en17174342.
- [24] A. A. Noyes and W. R. Whitney, "The rate of solution of solid substances in their own solutions", en, *Journal of the American Chemical Society*, vol. 19, no. 12, pp. 930–934, Dec. 1897, issn: 0002-7863, 1520-5126. doi: 10.1021/ja02086a003.
- [25] K. Zhang, Z.-Q. Zhou, P.-D. Han, Z.-G. Sun, and G. Xi, "Dissolution characteristics of solutes with different shapes using the moving particle semi-implicit method", en, *Physics of Fluids*, vol. 34, no. 11, p. 117 104, Nov. 2022, issn: 1070-6631, 1089-7666. doi: 10.1063/5.0120966.

- [26] A. Uzi, G. Beit Halevy, and A. Levy, "CFD-DEM Modeling of soluble NaCl particles conveyed in brine", en, *Powder Technology*, vol. 360, pp. 1278–1294, Jan. 2020, issn: 00325910. doi: 10.1016/j.powtec.2019.09.039.
- [27] L. Pauling, *General Chemistry*. New York, N.Y: Dover Publications, INC, 1988, isbn: 978-0-486-13465-9.
- [28] S. B. Pope, *Turbulent flows*, en, 1. publ., 12. print. Cambridge: Cambridge Univ. Press, 2015, isbn: 978-0-521-59125-6 978-0-521-59886-6.
- [29] F. Toschi and E. Bodenschatz, "Lagrangian Properties of Particles in Turbulence", en, *Annual Review of Fluid Mechanics*, vol. 41, no. 1, pp. 375–404, Jan. 2009, issn: 0066-4189, 1545-4479. doi: 10.1146/annurev.fluid.010908.165210.
- [30] N. E. L. Haugen, J. Krüger, D. Mitra, and T. Løvås, "The effect of turbulence on mass transfer rates of small inertial particles with surface reactions", en, *Journal of Fluid Mechanics*, vol. 836, pp. 932–951, Feb. 2018, issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2017.820.
- [31] E. L. Paul, V. A. Atiemo-Obeng, and S. M. Kresta, Eds., *Handbook of industrial mixing: science and practice*, en. Hoboken, N.J: Wiley-Interscience, 2010, isbn: 978-0-471-26919-9 978-0-471-45144-0 978-0-471-45145-7 978-1-60119-414-5.
- [32] H. Hartmann, J. Derksen, and H. Van Den Akker, "Numerical simulation of a dissolution process in a stirred tank reactor", en, *Chemical Engineering Science*, vol. 61, no. 9, pp. 3025–3032, May 2006, issn: 00092509. doi: 10.1016/j.ces.2005.10.058.
- [33] Z. Li, Y. Wang, K. Li, B. Zhang, and H. Chen, "CFD-DEM modeling of particle dissolution behavior in stirred tanks", en, *Powder Technology*, vol. 442, p. 119 894, Jun. 2024, issn: 00325910. doi: 10.1016/j.powtec.2024.119894.
- [34] F. Grisafi, A. Brucato, G. Caputo, S. Lima, and F. Scargiali, "Modelling particle dissolution in stirred vessels", en, *Chemical Engineering Research and Design*, vol. 195, pp. 662–672, Jul. 2023, issn: 02638762. doi: 10.1016/j.cherd.2023.06.026.
- [35] M. Ge and G. Zheng, "Fluid–Solid Mixing Transfer Mechanism and Flow Patterns of the Double-Layered Impeller Stirring Tank by the CFD-DEM Method", en, *Energies*, vol. 17, no. 7, p. 1513, Mar. 2024, issn: 1996-1073. doi: 10.3390/en17071513.
- [36] T. Hörmann, D. Suzzi, and J. G. Khinast, "Mixing and Dissolution Processes of Pharmaceutical Bulk Materials in Stirred Tanks: Experimental and Numerical Investigations", en, *Industrial & Engineering Chemistry Research*, vol. 50, no. 21, pp. 12 011–12 025, Nov. 2011, issn: 0888-5885, 1520-5045. doi: 10.1021/ie2002523.
- [37] M. Ishii and T. Hibiki, *Thermo-fluid dynamics of two-phase flow*, en. New York, N.Y: Springer Science+Business Media, 2006, OCLC: ocm63789873, isbn: 978-0-387-28321-0 978-0-387-29187-1.
- [38] J. W. Holbeach and M. R. Davidson, "Eulerian model for dissolution and dispersion of particles entering a liquid", en, *Third International Conference on CFD in the Minerals and Process Industries*, CSIRO, Melbourne, Australia, Dec. 2003.
- [39] D. Ramkrishna, *Population balances: theory and applications to particulate systems in engineering*, en. San Diego, CA: Academic Press, 2000, isbn: 978-0-12-576970-9.
- [40] S. Zhan, M. Li, J. Zhou, J. Yang, and Y. Zhou, "CFD simulation of dissolution process of alumina in an aluminum reduction cell with two-particle phase population balance model", en, *Applied Thermal Engineering*, vol. 73, no. 1, pp. 805–818, Dec. 2014, issn: 13594311. doi: 10.1016/j.applthermaleng.2014.08.040.
- [41] S. Balachandar, "A scaling analysis for point-particle approaches to turbulent multiphase flows", en, *International Journal of Multiphase Flow*, vol. 35, no. 9, pp. 801–810, Sep. 2009, issn: 03019322. doi: 10.1016/j.ijmultiphaseflow.2009.02.013.
- [42] M. R. Maxey and J. J. Riley, "Equation of motion for a small rigid sphere in a nonuniform flow", en, *The Physics of Fluids*, vol. 26, no. 4, pp. 883–889, Apr. 1983, issn: 0031-9171. doi: 10.1063/1.864230.
- [43] S. Elghobashi, "On predicting particle-laden turbulent flows", en, *Applied Scientific Research*, vol. 52, no. 4, pp. 309–329, Jun. 1994, issn: 1573-1987. doi: 10.1007/BF00936835.
- [44] H. Zhu, Z. Zhou, R. Yang, and A. Yu, "Discrete particle simulation of particulate systems: Theoretical developments", en, *Chemical Engineering Science*, vol. 62, no. 13, pp. 3378–3396, Jul. 2007, issn: 00092509. doi: 10.1016/j.ces.2006.12.089.
- [45] G. Tryggvason, *Direct Numerical Simulations of Gas–Liquid Multiphase Flows*, en. Cambridge: Cambridge University Press, 2011, isbn: 978-0-521-78240-1 978-1-139-15697-4.
- [46] R. Safa, A. Soltani Goharrizi, S. Jafari, and E. Jahanshahi Javaran, "Simulation of particles dissolution in the shear flow: A combined concentration lattice Boltzmann and smoothed profile approach", en, *Computers & Mathematics with Applications*, vol. 79, no. 3, pp. 603–622, Feb. 2020, issn: 08981221. doi: 10.1016/j.camwa.2019.07.015.
- [47] A. Nienow, "Agitated vessel particle-liquid mass transfer: A comparison between theories and data", en, *The Chemical Engineering Journal*, vol. 9, no. 2, pp. 153–160, Jan. 1975, issn: 03009467. doi: 10.1016/0300-9467(75)80007-4.
- [48] H. Cao, C. Amador, X. Jia, Y. Li, and Y. Ding, "A modelling framework for bulk particles dissolving in turbulent regime", en, *Chemical Engineering Research and Design*, vol. 114, pp. 108–118, Oct. 2016, issn: 02638762. doi: 10.1016/j.cherd.2016.08.012.
- [49] C. Carletti, S. Bikić, G. Montante, and A. Paglianti, "Mass Transfer in Dilute Solid–Liquid Stirred Tanks", en, *Industrial & Engineering Chemistry Research*, vol. 57, no. 18, pp. 6505–6515, May 2018, issn: 0888-5885, 1520-5045. doi: 10.1021/acs.iecr.7b04730.
- [50] M. Cerrillo-Gonzalez, M. Villen-Guzman, C. Vereda-Alonso, J. Rodriguez-Maroto, and J. Paz-Garcia, "Acid leaching of LiCoO2 enhanced by reducing agent. Model formulation and validation", en, *Chemosphere*, vol. 287, p. 132 020, Jan. 2022, issn: 00456535. doi: 10.1016/j.chemosphere.2021.132020.
- [51] M. M. Cerrillo-Gonzalez, M. Villen-Guzman, C. Vereda-Alonso, C. Gomez-Lahoz, J. M. Rodriguez-Maroto, and J. M. Paz-Garcia, "Recovery of Li and Co from LiCoO2 via Hydrometallurgical–Electrodialytic Treatment", en, *Applied Sciences*, vol. 10, no. 7, p. 2367, Jan. 2020, Number: 7 Publisher: Multidisciplinary Digital Publishing Institute, issn: 2076-3417. doi: 10.3390/app10072367.

- [52] M. D. M. Cerrillo-Gonzalez, J. M. Paz-Garcia, M. Muñoz-Espinosa, J. M. Rodriguez-Maroto, and M. Villen-Guzman, "Extraction and selective precipitation of metal ions from LiCoO₂ cathodes using citric acid", en, *Journal of Power Sources*, vol. 592, p. 233 870, Feb. 2024, issn: 03787753. doi: 10.1016/j.jpowsour.2023.233870.
- [53] M. d. M. Cerrillo-Gonzalez, M. Villen-Guzman, L. F. Acedo-Bueno, J. M. Rodriguez-Maroto, and J. M. Paz-Garcia, "Hydrometallurgical Extraction of Li and Co from LiCoO₂ Particles-Experimental and Modeling", en, *Applied Sciences*, vol. 10, no. 18, p. 6375, Jan. 2020, Number: 18 Publisher: Multidisciplinary Digital Publishing Institute, issn: 2076-3417. doi: 10.3390/app10186375.
- [54] C. Liu *et al.*, "Highly efficient lithium-ion battery cathode material recycling using deep eutectic solvent based nanofluids", en, *RSC Sustainability*, vol. 1, no. 2, pp. 270–281, 2023, issn: 2753-8125. doi: 10.1039/D2SU000047D.
- [55] N. Wongnaree, T. Patcharawit, T. Yingnakorn, and S. Khumkoa, "Leaching Kinetics of Valuable Metals from Calcined Material of Spent Lithium-Ion Batteries", en, *ACS Omega*, vol. 9, no. 47, pp. 46 822–46 833, Nov. 2024, issn: 2470-1343, 2470-1343. doi: 10.1021/acsomega.4c05086.
- [56] D. Amalia, P. Singh, W. Zhang, and A. N. Nikoloski, "The Effect of a Molasses Reductant on Acetic Acid Leaching of Black Mass from Mechanically Treated Spent Lithium-Ion Cylindrical Batteries", en, *Sustainability*, vol. 15, no. 17, p. 13 171, Sep. 2023, issn: 2071-1050. doi: 10.3390/su151713171.
- [57] K. D. Squires and J. K. Eaton, "Preferential concentration of particles by turbulence", en, *Physics of Fluids A: Fluid Dynamics*, vol. 3, no. 5, pp. 1169–1178, May 1991, issn: 0899-8213. doi: 10.1063/1.858045.
- [58] *JHU Turbulence Database: Isotropic Turbulence (1024-cubed)*, 2008. doi: 10.7281/T1KK98XB.
- [59] M. Jin and Q. Chen, "Improvement of fast fluid dynamics with a conservative semi-Lagrangian scheme", en, *International Journal of Numerical Methods for Heat & Fluid Flow*, vol. 25, no. 1, pp. 2–18, Jan. 2015, issn: 0961-5539. doi: 10.1108/HFF-04-2013-0119.
- [60] P. K. Yeung and S. Pope, "An algorithm for tracking fluid particles in numerical simulations of homogeneous turbulence", *Journal of Computational Physics*, vol. 79, pp. 373–416, 1988. doi: 10.1016/0021-9991(88)90022-8.
- [61] M. A. T. Van Hinsberg, J. H. M. T. T. Boonkamp, F. Toschi, and H. J. H. Clercx, "Optimal interpolation schemes for particle tracking in turbulence", en, *Physical Review E*, vol. 87, no. 4, p. 043 307, Apr. 2013, issn: 1539-3755, 1550-2376. doi: 10.1103/PhysRevE.87.043307.
- [62] R. Hockney and J. Eastwood, *Computer Simulation Using Particles*. CRC Press, 2021, isbn: 978-1-4398-2205-0.
- [63] E. Hairer, S. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems (Solving Ordinary Differential Equations II: Stiff and Differential-algebraic Problems)*. Springer, 1993, isbn: 978-3-540-60452-5.
- [64] J. K. Eaton and J. Fessler, "Preferential concentration of particles by turbulence", *International Journal of Multiphase Flow*, vol. 20, no. 1, pp. 169–209, Aug. 1994. doi: 10.1016/0301-9322(94)90072-8.
- [65] D. Marchisio and R. Fox, "Reacting Flows and the Interaction between Turbulence and Chemistry", en, in *Reference Module in Chemistry, Molecular Sciences and Chemical Engineering*, Elsevier, 2016, B9780124095472115264, isbn: 978-0-12-409547-2. doi: 10.1016/B978-0-12-409547-2.11526-4.