


Politecnico di Torino


## Politecnico di Torino

Department of Industrial Engineering
Master of Science in Aerospace Engineering


# Development and Validation of a Diffusion-Based Lagrangian Spray Solver in OpenFOAM for Liquid Jet Atomization in Subsonic Crossflow


**Candidate:**
Lorenzo Proli

**Supervisors:**
Prof. Andrea Ferrero
Ing. Leonardo Stumpo



**Graduation Session:** July 2026
**Academic Year 2025–2026**

# Abstract

The atomisation of a liquid jet injected into a gaseous crossflow is a fundamental process in aerospace propulsion systems, including gas turbine combustors and liquid rocket engines. High-fidelity simulations that explicitly track the interface between the liquid and the gas can resolve the breakup of the jet in detail, but at a computational cost that remains prohibitive for parametric design studies. This work presents the development and implementation, within the open-source OpenFOAM framework, of an enhanced Lagrangian spray model in which the liquid is represented as a cloud of computational parcels rather than as a resolved interface, providing a far less expensive alternative suited to design exploration. The model couples an established treatment of the breakup of the intact liquid column with a ligament-based correction of the resulting fragment sizes, a drag law that accounts for the deformation of the droplets, and a secondary breakup mechanism that further fragments them according to the local aerodynamic loading. A well-known difficulty of this class of models is that the exchange of mass, momentum and energy between the liquid parcels and the surrounding gas becomes strongly sensitive to the computational mesh in the near-injector region, where the liquid is most concentrated; this sensitivity is mitigated through a diffusion-based smoothing technique that distributes the exchange over a finite region of space rather than concentrating it in a single computational cell.

The accuracy of the complete model is assessed against two independent experimental benchmarks that differ in injection geometry and operating conditions: a free liquid jet discharging into a transverse gas stream, examined at two values of the ratio between the momentum of the jet and that of the crossflow, and a confined configuration in which the liquid is sheared by a surrounding gas film, examined at increasing values of the aerodynamic Weber number (the dimensionless ratio between the disruptive aerodynamic forces acting on the liquid and the stabilising effect of its surface tension). For each configuration the width of the smoothing region is varied systematically and studied on its own, rather than transferring a single setup across all cases. The aim is not to treat this width as an empirical tuning parameter, but to quantify, for each configuration, the trade-off between the accuracy of the predicted droplet sizes and that of the predicted droplet velocities. The smoothing is shown to roughly halve the error on droplet size in the crossflow configuration, while in the gas-film configuration it proves essential to recover a spray consistent with the expected near-injector physics. The most severe operating regime could not be completed within the present framework, and its failure mechanism could not be isolated; it lies beyond the scope of the present comparison and is identified as a direction for future development.

# Acronyms

| Acronym | Description |
|---|---|
| CFD | Computational Fluid Dynamics |
| CLSVOF | Coupled Level-Set / Volume Of Fluid |
| DBM | Diffusion-based coarse-graining method |
| DNS | Direct Numerical Simulation |
| DPM | Discrete Phase Model |
| EL | Eulerian–Lagrangian |
| GLOP | Gas–Liquid Orifice-type Pintle injection element |
| KH | Kelvin–Helmholtz |
| LES | Large Eddy Simulation |
| LJICF | Liquid Jet In CrossFlow |
| LJIGF | Liquid Jet In subsonic Gas Film |
| LPT | Lagrangian Particle Tracking |
| NRMSE | Normalised Root-Mean-Square Error |
| PCM | Particle Centroid Method |
| PE | Pilch–Erdman |
| PIMPLE | Merged PISO–SIMPLE pressure–velocity coupling algorithm |
| RANS | Reynolds-Averaged Navier–Stokes |
| RAS | Reynolds-Averaged Simulation (OpenFOAM term for RANS) |
| SBES | Stress-Blended Eddy Simulation |
| SMD | Sauter Mean Diameter |
| SST | Shear Stress Transport turbulence model |
| VOF | Volume Of Fluid |

# Nomenclature

| Symbol | Description | Value | Unit |
|---|---|---|---|
| $a$ | Injector radius, $a = D_{\text{inj}}/2$ | – | m |
| $A_1$ | Reitz Wave spreading coefficient | 0.188 | – |
| $A_p$ | Parcel projected area | – | $\text{m}^2$ |
| $b$ | DBM smoothing bandwidth parameter | case-dependent | m |
| $B_0$ | KH/Wave child-size constant, $d_c = 2B_0\Lambda_{\text{KH}}$; Reitz default 0.61, Lambert LJICF calibration 2.44 | case-dependent | – |
| $B_1$ | KH/Wave stripping-time constant; Madabhushi LJICF default 10.0 | case-dependent | – |
| $B_{1,PE}$ | PE velocity-history coefficient used in the OpenFOAM-compatible admission check | 0.375 | – |
| $B_{2,PE}$ | PE velocity-history coefficient used in the OpenFOAM-compatible admission check | 0.2274 | – |
| $C_0$ | Wu/Madabhushi column-breakup time constant, code variable `C0_`; Madabhushi LJICF default 3.44 | case-dependent | – |
| $C_D$ | Drag coefficient | – | – |
| $C_{D,\text{blob}}$ | Intact-column/blob drag coefficient; Madabhushi LJICF default 1.48 | case-dependent | – |
| $C_{D,\text{disc}}$ | Disc-shape drag coefficient during PE deformation | 1.2 | – |
| $d_c$ | KH/Wave target child diameter | – | m |
| $d_i$ | Diameter of the $i$-th child fragment | – | m |
| $d_p$ | Parcel diameter | – | m |
| $d_{\text{ref}}$ | Deformed reference diameter during PE deformation | – | m |
| $d_{\text{SMD}}$ | Target Sauter mean diameter after PE breakup | – | m |
| $d_{\text{stable}}$ | Maximum stable fragment diameter used in PE admission checks | – | m |
| $D_0$ | Initial blob diameter, $D_0 = D_{\text{inj}}$ | – | m |
| $D_{10}$ | Arithmetic mean (number) diameter, $D_{10} = \sum N_p d / \sum N_p$ | – | m |
| $D_{\text{inj}}$ | Injector orifice diameter | case-dependent | m |
| $D_{32}$ | Sauter Mean Diameter statistic | – | m |
| $D_{p,0}$ | Parent diameter latched at PE onset | – | m |

| Symbol | Description | Value | Unit |
|---|---|---|---|
| $f$ | Deformation-ramp fraction during PE deformation, $0 \le f \le 1$ | – | – |
| $F_{\mathrm{Lig}}$ | Lambert ligament size-correction factor; Lambert LJICF default 0.4 | case-dependent | – |
| $\mathbf{F}_D$ | Aerodynamic drag force on a parcel | – | N |
| $\mathbf{g}$ | Gravitational acceleration | 9.81 | $\mathrm{m/s^2}$ |
| $J$ | Jet-to-crossflow momentum flux ratio | – | – |
| $k$ | Turbulent kinetic energy ($k$–$\omega$ SST carrier model) | – | $\mathrm{m^2/s^2}$ |
| $m_c$ | Represented mass of a child parcel | – | kg |
| $m_i$ | Represented mass of the $i$-th child fragment | – | kg |
| $m_p$ | Parcel represented mass | – | kg |
| $\dot{m}_l$ | Liquid mass-flow rate | case-dependent | kg/s |
| $m_s$ | KH/Wave stripped-mass reservoir; also reused as code-level state marker after state transitions | – | kg |
| $N_p$ | Number of physical droplets represented by a computational parcel, OpenFOAM `nParticle` | – | – |
| $N_{p,c}$ | Statistical weight assigned to a child parcel | – | – |
| $n_f$ | Number of PE child parcels generated at breakup | 5 | – |
| $Oh$ | Ohnesorge number | – | – |
| $p$ | Static (carrier) pressure | case-dependent | Pa |
| $q$ | Liquid-to-gas density ratio, $q = \rho_l/\rho_g$ | – | – |
| $Q$ | Volumetric liquid flow rate | case-dependent | $\mathrm{m^3/s}$ |
| $Re_l$ | Liquid Reynolds number based on injector radius $a$ | – | – |
| $Re_p$ | Parcel Reynolds number based on $d_p$ | – | – |
| $S_m$ | Gas-phase mass source term | – | $\mathrm{kg/(m^3\,s)}$ |
| $\mathbf{S}_u$ | Gas-phase momentum source term | – | $\mathrm{N/m^3}$ |
| $S_p$ | Semi-implicit drag/source coefficient returned by the particle-force model | – | kg/s |
| $t^*$ | PE characteristic time, $t^* = \frac{D_{p,0}}{U_{\mathrm{rel},0}}\sqrt{\rho_l/\rho_g}$ | – | s |
| $t_b$ | Total PE breakup time | – | s |
| $t_c$ | Parcel lifetime or breakup clock since injection | – | s |
| $t_{cb}$ | Madabhushi column-breakup time, code variable `tColumnBreakup` | – | s |
| $t_{\mathrm{def}}$ | PE deformation time, $t_{\mathrm{def}} = 1.6t^*$ | – | s |
| $t_{\mathrm{sec}}$ | PE onset time stored in the parcel state variable `y` | – | s |
| $Ta$ | Taylor parameter, $Ta = Oh\sqrt{We_g}$ | – | – |
| $T$ | Static (carrier) temperature | case-dependent | K |
| $\mathbf{u}_g$ | Carrier-gas velocity at the parcel position | – | m/s |
| $\mathbf{u}_p$ | Parcel velocity vector | – | m/s |
| $u_n$ | Rim-expansion velocity imparted to PE children | – | m/s |

| Symbol | Description | Value | Unit |
|---|---|---|---|
| $U_g$ | Bulk gas velocity (crossflow in LJICF, gas film in LJIGF) | case-dependent | m/s |
| $U_l$ | Bulk liquid injection velocity | case-dependent | m/s |
| $U_{\mathrm{mean}}$ | nParticle-weighted mean parcel speed $\lvert\mathbf{u}_p\rvert$ | – | m/s |
| $\mathbf{U}_{\mathrm{rel}}$ | Relative velocity vector, $\mathbf{u}_g - \mathbf{u}_p$ | – | m/s |
| $U_{\mathrm{rel}}$ | Relative velocity magnitude | – | m/s |
| $U_{\mathrm{rel},0}$ | Relative velocity latched at PE onset | – | m/s |
| $U_x$ | Streamwise (axial) droplet velocity component | – | m/s |
| $V_d$ | PE velocity-history estimate of fragment-cloud acceleration | – | m/s |
| $V_{\mathrm{cell}}$ | Eulerian cell volume | – | $\mathrm{m}^3$ |
| $W$ | Gas-film height in the LJIGF/GLOP configuration | case-dependent | m |
| $We$ | Parcel-based aerodynamic Weber number | – | – |
| $We_g$ | KH/Wave gas-side Weber number based on injector radius $a$ | – | – |
| $We_g^{D_{\mathrm{inj}}}$ | Orifice-based gas Weber number, $\rho_g U_g^2 D_{\mathrm{inj}}/\sigma$ | – | – |
| $We_g^{\mathrm{film}}$ | Gas-film Weber number, $\rho_g U_g^2 W/\sigma$ | case-dependent | – |
| $We_l$ | KH/Wave liquid Weber number | – | – |
| $We_{\mathrm{crit}}$ | Viscosity-corrected critical Weber number for PE onset | – | – |
| $\alpha_l$ | Liquid volume fraction | – | – |
| $\Delta t$ | Numerical time step or Lagrangian update step | – | s |
| $\varepsilon$ | Numerical lower bound used in stability checks | – | – |
| $\Lambda_{\mathrm{KH}}$ | KH most-unstable wavelength | – | m |
| $\mu_g$ | Gas dynamic viscosity | – | Pa s |
| $\mu_l$ | Liquid dynamic viscosity | – | Pa s |
| $\mu_t$ | Turbulent (eddy) dynamic viscosity | – | Pa s |
| $\mu_{\mathrm{eff}}$ | Effective dynamic viscosity, $\mu_{\mathrm{eff}} = \mu_g + \mu_t$ | – | Pa s |
| $\nu_t$ | Turbulent kinematic viscosity | – | $\mathrm{m}^2/\mathrm{s}$ |
| $\Omega_{\mathrm{KH}}$ | KH maximum wave growth rate | – | 1/s |
| $\omega$ | Specific turbulence dissipation rate ($k$–$\omega$ SST carrier model) | – | 1/s |
| $\rho_g$ | Gas density | – | $\mathrm{kg/m}^3$ |
| $\rho_l$ | Liquid density | – | $\mathrm{kg/m}^3$ |
| $\sigma$ | Liquid–gas surface tension | – | N/m |
| $\sigma_s^2$ | DBM diffusion/screening parameter, $\sigma_s^2 = b^2/4$ | case-dependent | $\mathrm{m}^2$ |
| $\tau_{\mathrm{KH}}$ | KH/Wave stripping timescale | – | s |
| $\boldsymbol{\tau}_{\mathrm{eff}}$ | Effective (viscous + turbulent) stress tensor | – | Pa |
| $\chi$ | Drag projected-area correction factor during PE deformation | – | – |

| Symbol | Description | Value | Unit |
|---|---|---|---|
| $\xi$ | Standard normal random variable used in root-normal child-size sampling | – | – |

# Contents

**Acronyms** . . . . . ii

**Nomenclature** . . . . . iii

**List of Figures** . . . . . ix

**List of Tables** . . . . . xiii

**1 Introduction** . . . . . 1
- 1.1 Preface . . . . . 1
- 1.2 Objectives and Contributions . . . . . 2
- 1.3 Structure of the Thesis . . . . . 3

**2 Physical Background and Lagrangian Model Formulation** . . . 5
- 2.1 Liquid-Jet-in-Crossflow Breakup Physics . . . . . 6
  - 2.1.1 Stages of Breakup . . . . . 6
  - 2.1.2 Column Stripping and Column Fracture . . . . . 7
  - 2.1.3 Controlling Parameters . . . . . 7
  - 2.1.4 Breakup-Regime Interpretation . . . . . 8
  - 2.1.5 Liquid Jet in Subsonic Gas Film Variant . . . . . 9
- 2.2 Eulerian–Lagrangian Spray Modelling Framework . . . . . 9
  - 2.2.1 Carrier-Phase Equations . . . . . 10
  - 2.2.2 Parcel Representation and Equations of Motion . . . . . 10
  - 2.2.3 Two-Way Coupling and Dilute-Spray Hypothesis . . . . . 11
  - 2.2.4 Strengths and Limitations of the EL Approach . . . . . 12
- 2.3 Blob Representation and Reitz KH/Wave Surface Stripping . . . . . 12
  - 2.3.1 Linear Stability Background . . . . . 13
  - 2.3.2 Dimensionless Groups . . . . . 14
  - 2.3.3 Most-Unstable Wavelength and Growth Rate . . . . . 14
  - 2.3.4 KH Stripping, Child Emission and Velocity Assignment . . . . . 14
- 2.4 Madabhushi Column-Breakup Model . . . . . 16
  - 2.4.1 Column-Breakup Time and State Transition . . . . . 17
- 2.5 Lambert Ligament Correction . . . . . 17
  - 2.5.1 Ligament Correction and Model Scope . . . . . 18
- 2.6 Pilch–Erdman Secondary Breakup . . . . . 18
  - 2.6.1 Onset, Latching and Timescales . . . . . 19
  - 2.6.2 Deformation, Child-Size Distribution and Fragmentation . . . . 21
- 2.7 Deformation-Dependent Drag Model . . . . . 23
  - 2.7.1 State-Dependent Drag Laws . . . . . 23
- 2.8 Source-Term Coupling and Diffusion-Based Coarse Graining . . . . . 25

2.8.1 Diffusion-Based Smoothing Formulation . . . 25
2.8.2 Bandwidth Selection and Model Role . . . 26
2.9 Complete Model Sequence, Assumptions, and Limitations . . . 27
2.9.1 Original Contribution . . . 28
2.9.2 Assumptions and Limitations . . . 28

**3 OpenFOAM Implementation . . . 30**
3.1 Solver and Library Architecture . . . 30
3.2 Injection Model and Initial Blob Parcels . . . 32
3.3 Parcel State Extension and Field Reuse . . . 32
3.4 Breakup State Machine in `Madabhushi.C` . . . 35
3.4.1 KH/Wave Stripping and Column-Breakup Latch . . . 37
3.4.2 PE Admission, Fragmentation and Lambert Flag . . . 38
3.5 Breakup-Model Child Signalling . . . 38
3.6 Child Creation and Mass Conservation in `SprayParcel.C` . . . 39
3.7 State-Dependent Drag Force . . . 40
3.8 DBM Source-Term Smoothing in `sprayFoamDBM.C` . . . 41
3.9 Time-Loop Integration and PIMPLE Coupling . . . 42
3.10 Case Dictionaries and Runtime Parameters . . . 43
3.11 Build and Execution Workflow . . . 44
3.12 Debug Outputs and Numerical Consistency Checks . . . 44
3.13 Implementation Summary . . . 45

**4 Lambert LJICF Validation: J10 Sensitivity and J20 Follow-up . 46**
4.1 Purpose of the Lambert Validation . . . 46
4.2 Numerical Setup and Post-processing Methodology . . . 47
4.3 J10: DBM Bandwidth Sensitivity on $D_{32}$ and $U_x$ Profiles . . . 49
4.4 Quantitative Error Metrics and Trade-off Analysis . . . 54
4.5 J10: Representative-Bandwidth Assessment . . . 55
4.6 Spray and Carrier-Field Visualisation (J10) . . . 56
4.7 J20 Follow-up Validation . . . 60
4.8 Summary of the Lambert Validation . . . 66

**5 Zhang LJIGF/GLOP Validation and DBM Necessity . . . 68**
5.1 Motivation: Differences Between Lambert LJICF and Zhang LJIGF . . . 68
5.2 Numerical Setup and Reference Quantities . . . 69
5.3 Carrier-to-Parcel Interpolation, PCM Stability and DBM Necessity . . . 71
5.4 DBM Bandwidth Sensitivity at Intermediate Gas-film Weber Number ($We_g^{\mathrm{film}} \approx 500$) . . . 72
5.5 Spray and Carrier-Field Visualisation at $We_g^{\mathrm{film}} \approx 500$ . . . 74
5.6 Selected DBM Bandwidth and Physical Interpretation . . . 77
5.7 Higher Gas-film Weber Number: $We_g^{\mathrm{film}} \approx 3500$ . . . 77
5.8 Spray and Carrier-Field Visualisation at $We_g^{\mathrm{film}} \approx 3500$ . . . 78
5.9 Summary of LJIGF/GLOP Findings . . . 82

**6 Conclusions . . . 84**
6.1 Summary of Results and Model Advantages . . . 84
6.2 Limitations of the Current Implementation . . . 85
6.3 Future Developments . . . 86

# List of Figures


2.1 Madabhushi breakup model and LJICF breakup structure [2, 8]. . . . . 13
2.2 Acceleration-induced secondary-breakup regimes as a function of aerodynamic Weber number, adapted from [9]. . . . . 19
2.3 Normal rim-expansion velocity geometry for PE child-parcel velocity assignment, adapted from [8]. . . . . 23

3.1 State-machine overview of the coupled Madabhushi–Lambert breakup and drag implementation. The diagram follows the state convention of Table 3.3 and the regime-tag convention of Table 3.4. . . . . 36

4.1 Lambert LJICF computational domain and boundary conditions. . . . . 47
4.2 Lambert LJICF mesh: structured hexahedral grid with near-nozzle in-plane refinement and wall-normal boundary-layer clustering. . . . . 47
4.2a Front view . . . . . 47
4.2b Isometric view . . . . . 47
4.3 $D_{32}$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J10: PCM and DBM $b \in \{D_{\text{inj}}, 3D_{\text{inj}}, 6D_{\text{inj}}\}$, with SBES and experiment. . . . . 50
4.4 $U_x$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J10; same curves as Figure 4.3. . . . . 50
4.5 Two-dimensional $D_{32}$ maps, J10: PCM, $b = D_{\text{inj}}$, $b = 3D_{\text{inj}}$, $b = 6D_{\text{inj}}$ ($y$ wall-normal, $z$ spanwise, common colour scale). . . . . 53
4.5a $D_{32}$ colormap at $x/d = 30$. . . . . 53
4.5b $D_{32}$ colormap at $x/d = 60$. . . . . 53
4.6 J10 spray fields: parcel diameter $d$ (left, in m) and parcel streamwise velocity $U_x$ (right, in m/s), PCM and $b = D_{\text{inj}}$. . . . . 58
4.6a PCM ($b = 0$): parcel $d$ . . . . . 58
4.6b PCM ($b = 0$): parcel $U_x$ . . . . . 58
4.6c $b = D_{\text{inj}}$: parcel $d$ . . . . . 58
4.6d $b = D_{\text{inj}}$: parcel $U_x$ . . . . . 58
4.7 J10 spray fields: parcel diameter $d$ (in m) and parcel streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom). . . . . 58
4.7a $b = 3D_{\text{inj}}$: parcel $d$ . . . . . 58
4.7b $b = 3D_{\text{inj}}$: parcel $U_x$ . . . . . 58
4.7c $b = 6D_{\text{inj}}$: parcel $d$ . . . . . 58
4.7d $b = 6D_{\text{inj}}$: parcel $U_x$ . . . . . 58
4.8 J10 carrier fields on the $z = 0$ plane: pressure $p$ (left, in Pa) and carrier streamwise velocity $U_x$ (right, in m/s), PCM and $b = D_{\text{inj}}$. . . . . 59
4.8a PCM ($b = 0$): carrier $p$ . . . . . 59
4.8b PCM ($b = 0$): carrier $U_x$ . . . . . 59

4.8c $b = D_{\text{inj}}$: carrier $p$ . . . 59
4.8d $b = D_{\text{inj}}$: carrier $U_x$ . . . 59
4.9 J10 carrier fields on the $z = 0$ plane: pressure $p$ (in Pa) and carrier streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom). . . . 59
4.9a $b = 3D_{\text{inj}}$: carrier $p$ . . . 59
4.9b $b = 3D_{\text{inj}}$: carrier $U_x$ . . . 59
4.9c $b = 6D_{\text{inj}}$: carrier $p$ . . . 59
4.9d $b = 6D_{\text{inj}}$: carrier $U_x$ . . . 59
4.10 $D_{32}$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J20: PCM and DBM $b \in \{D_{\text{inj}}, 3D_{\text{inj}}, 6D_{\text{inj}}\}$, with SBES and experiment. . . . 62
4.11 $U_x$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J20; same curves as Figure 4.10. . . . 62
4.12 Two-dimensional $D_{32}$ maps, J20: PCM, $b = D_{\text{inj}}$, $b = 3D_{\text{inj}}$, $b = 6D_{\text{inj}}$ ($y$ wall-normal, $z$ spanwise, same colour scale as Figure 4.5). . . . 63
4.12a $D_{32}$ colormap at $x/d = 30$. . . . 63
4.12b $D_{32}$ colormap at $x/d = 60$. . . . 63
4.13 J20 spray fields: parcel diameter $d$ (in m) and parcel streamwise velocity $U_x$ (in m/s), PCM and $b = D_{\text{inj}}$. . . . 64
4.13a PCM ($b = 0$): parcel $d$ . . . 64
4.13b PCM ($b = 0$): parcel $U_x$ . . . 64
4.13c $b = D_{\text{inj}}$: parcel $d$ . . . 64
4.13d $b = D_{\text{inj}}$: parcel $U_x$ . . . 64
4.14 J20 spray fields: parcel diameter $d$ (in m) and parcel streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom). . . . 64
4.14a $b = 3D_{\text{inj}}$: parcel $d$ . . . 64
4.14b $b = 3D_{\text{inj}}$: parcel $U_x$ . . . 64
4.14c $b = 6D_{\text{inj}}$: parcel $d$ . . . 64
4.14d $b = 6D_{\text{inj}}$: parcel $U_x$ . . . 64
4.15 J20 carrier fields on the $z = 0$ plane: pressure $p$ (in Pa) and carrier streamwise velocity $U_x$ (in m/s), PCM and $b = D_{\text{inj}}$. . . . 65
4.15a PCM ($b = 0$): carrier $p$ . . . 65
4.15b PCM ($b = 0$): carrier $U_x$ . . . 65
4.15c $b = D_{\text{inj}}$: carrier $p$ . . . 65
4.15d $b = D_{\text{inj}}$: carrier $U_x$ . . . 65
4.16 J20 carrier fields on the $z = 0$ plane: pressure $p$ (in Pa) and carrier streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom). . . . 65
4.16a $b = 3D_{\text{inj}}$: carrier $p$ . . . 65
4.16b $b = 3D_{\text{inj}}$: carrier $U_x$ . . . 65
4.16c $b = 6D_{\text{inj}}$: carrier $p$ . . . 65
4.16d $b = 6D_{\text{inj}}$: carrier $U_x$ . . . 65

5.1 Zhang LJIGF/GLOP computational domain and boundary conditions. . . . 69
5.2 Zhang GLOP mesh. . . . 71
5.2a Front view ($x$–$z$ plane) . . . 71
5.2b Isometric view . . . 71
5.3 $We_g^{\text{film}} \approx 500$: DBM bandwidth sensitivity at $\Delta x = 60$ mm. . . . 73
5.3a Equivalent droplet count . . . 73
5.3b $D_{10}$ . . . 73

5.3c $D_{32}$ . . . 73
5.3d $U_{\mathrm{mean}}$ . . . 73
5.4 $We_g^{\mathrm{film}} \approx 500$ spray fields: parcel $d$ and $|U_p|$, PCM (top) and $b = D_{\mathrm{inj}}$ (bottom). . . . 75
5.4a PCM ($b = 0$): parcel $d$ . . . 75
5.4b PCM ($b = 0$): parcel $|U_p|$ . . . 75
5.4c $b = D_{\mathrm{inj}}$: parcel $d$ . . . 75
5.4d $b = D_{\mathrm{inj}}$: parcel $|U_p|$ . . . 75
5.5 $We_g^{\mathrm{film}} \approx 500$ spray fields: parcel $d$ and $|U_p|$, $b = 3D_{\mathrm{inj}}$ (top) and $b = 6D_{\mathrm{inj}}$ (bottom). . . . 75
5.5a $b = 3D_{\mathrm{inj}}$: parcel $d$ . . . 75
5.5b $b = 3D_{\mathrm{inj}}$: parcel $|U_p|$ . . . 75
5.5c $b = 6D_{\mathrm{inj}}$: parcel $d$ . . . 75
5.5d $b = 6D_{\mathrm{inj}}$: parcel $|U_p|$ . . . 75
5.6 $We_g^{\mathrm{film}} \approx 500$ carrier fields on the $y = 0$ plane: pressure $p$ (left) and velocity magnitude $|U|$ (right), PCM and $b = D_{\mathrm{inj}}$. . . . 76
5.6a PCM ($b = 0$): carrier $p$ . . . 76
5.6b PCM ($b = 0$): carrier $|U|$ . . . 76
5.6c $b = D_{\mathrm{inj}}$: carrier $p$ . . . 76
5.6d $b = D_{\mathrm{inj}}$: carrier $|U|$ . . . 76
5.7 $We_g^{\mathrm{film}} \approx 500$ carrier fields on the $y = 0$ plane: $p$ and $|U|$, $b = 3D_{\mathrm{inj}}$ (top) and $b = 6D_{\mathrm{inj}}$ (bottom). . . . 76
5.7a $b = 3D_{\mathrm{inj}}$: carrier $p$ . . . 76
5.7b $b = 3D_{\mathrm{inj}}$: carrier $|U|$ . . . 76
5.7c $b = 6D_{\mathrm{inj}}$: carrier $p$ . . . 76
5.7d $b = 6D_{\mathrm{inj}}$: carrier $|U|$ . . . 76
5.8 $We_g^{\mathrm{film}} \approx 3500$: DBM bandwidth sensitivity at $\Delta x = 60$ mm. Reference data from Zhang et al. [3], Fig. 24. . . . 78
5.8a Equivalent droplet count . . . 78
5.8b $D_{10}$ . . . 78
5.8c $D_{32}$ . . . 78
5.8d $U_{\mathrm{mean}}$ . . . 78
5.9 $We_g^{\mathrm{film}} \approx 3500$ spray fields on the $y = 0$ symmetry plane: parcel diameter $d$ (left) and velocity magnitude $|U_p|$ (right), PCM and $b = D_{\mathrm{inj}}$. . . . 80
5.9a PCM ($b = 0$): parcel $d$ . . . 80
5.9b PCM ($b = 0$): parcel $|U_p|$ . . . 80
5.9c $b = D_{\mathrm{inj}}$: parcel $d$ . . . 80
5.9d $b = D_{\mathrm{inj}}$: parcel $|U_p|$ . . . 80
5.10 $We_g^{\mathrm{film}} \approx 3500$ spray fields: parcel $d$ and $|U_p|$, $b = 3D_{\mathrm{inj}}$ (top) and $b = 6D_{\mathrm{inj}}$ (bottom). . . . 80
5.10a $b = 3D_{\mathrm{inj}}$: parcel $d$ . . . 80
5.10b $b = 3D_{\mathrm{inj}}$: parcel $|U_p|$ . . . 80
5.10c $b = 6D_{\mathrm{inj}}$: parcel $d$ . . . 80
5.10d $b = 6D_{\mathrm{inj}}$: parcel $|U_p|$ . . . 80
5.11 $We_g^{\mathrm{film}} \approx 3500$ carrier fields on the $y = 0$ plane: pressure $p$ (left) and velocity magnitude $|U|$ (right), PCM and $b = D_{\mathrm{inj}}$. . . . 81
5.11a PCM ($b = 0$): carrier $p$ . . . 81

5.11b PCM ($b = 0$): carrier $|U|$ . . . . . 81
5.11c $b = D_{\text{inj}}$: carrier $p$ . . . . . 81
5.11d $b = D_{\text{inj}}$: carrier $|U|$ . . . . . 81
5.12 $We_g^{\text{film}} \approx 3500$ carrier fields on the $y = 0$ plane: $p$ and $|U|$, $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom). . . . . 81
5.12a $b = 3D_{\text{inj}}$: carrier $p$ . . . . . 81
5.12b $b = 3D_{\text{inj}}$: carrier $|U|$ . . . . . 81
5.12c $b = 6D_{\text{inj}}$: carrier $p$ . . . . . 81
5.12d $b = 6D_{\text{inj}}$: carrier $|U|$ . . . . . 81

# List of Tables


3.1 Main modified or added source components. . . . . 31
3.2 Reuse of existing parcel fields in the Madabhushi–Lambert implementation. 33
3.3 Meaning of the `ms` state variable. . . . . 34
3.4 Meaning of the `user` parcel-regime flag. . . . . 35
3.5 State-dependent drag selection in `MadabhushiDragForce`. . . . . 41
3.6 Time-loop sequence of the coupled solver. . . . . 42

4.1 Lambert LJICF numerical setup and dispersed-phase model parameters. 48
4.2 J10 error metrics . . . . . 54
4.3 J20 error metrics . . . . . 61
4.4 Common-support NRMSE . . . . . 67

5.1 Zhang LJIGF/GLOP numerical setup . . . . . 70

# 1 Introduction

## 1.1 Preface

Liquid atomisation is a key physical process in a wide range of aerospace propulsion systems. In gas turbine combustors, the quality of fuel atomisation directly governs flame stability, combustion efficiency, and pollutant emissions. In propulsion-relevant injection systems, the interaction between a liquid jet and a surrounding gaseous stream governs atomisation quality, mixing, evaporation rate, and ultimately combustion performance. A thorough understanding of the underlying spray dynamics is therefore essential for the design and optimisation of modern propulsion components [1].

Among the configurations of practical interest, the liquid jet in crossflow (LJICF) has attracted considerable attention in both the experimental and computational communities [2]. In this configuration, a liquid jet is injected transversely into a gaseous stream, undergoing a complex sequence of deformation, primary breakup into ligaments and large drops, and subsequent secondary atomisation into finer droplets. The resulting spray is characterised by a strong coupling between inertial, surface tension, and aerodynamic forces, whose relative importance is captured by dimensionless groups such as the aerodynamic Weber number and the momentum flux ratio.

A closely related configuration is the liquid jet in subsonic gas film (LJIGF), in which the liquid is injected into a finite-height subsonic gas film rather than a uniform crossflow. This setup is representative of gas-liquid orifice type pintle (GLOP) simplified injection elements used in throttleable liquid rocket engines [3], and introduces additional complexity through the coupling between the jet column and the confined gas-film boundary layer. The LJIGF configuration has received comparatively less modelling attention than the canonical LJICF, yet its relevance to modern propulsion hardware motivates its inclusion in the present work. High-fidelity computational approaches, such as Volume-of-Fluid (VOF) methods combined with Direct Numerical Simulation (DNS) or Large Eddy Simulation (LES), can resolve the liquid–gas interface with great accuracy and provide detailed insight into the breakup mechanisms [4]. However, the associated computational cost scales prohibitively with the physical domain size and the range of relevant length scales, making these methods unsuitable for routine engineering design iterations or parametric studies.

Lagrangian spray models offer a pragmatic alternative. In this framework, the dispersed liquid phase is represented by a statistical sample of computational parcels whose evolution is governed by sub-models for drag, breakup, evaporation, and turbulent dispersion, while the carrier gas phase is solved with conventional Reynolds-Averaged

Navier–Stokes (RANS) or LES approaches. The predictive fidelity of such models depends critically on the accuracy of the sub-models for primary and secondary breakup, and on the proper treatment of the near-injector source terms, whose mesh sensitivity can otherwise compromise the solution quality.

## 1.2 Objectives and Contributions

The present work aims to develop, implement, and validate an enhanced Lagrangian spray model within the open-source OpenFOAM framework, targeting both LJICF and LJIGF configurations. The main objectives and original contributions are summarised as follows.

1. **Model development in OpenFOAM.** A new Lagrangian injection and breakup library is built within the `lagrangian/intermediate` solver infrastructure of OpenFOAM [5], extending the existing cloud and parcel classes to accommodate the LJICF-specific physics described in the items below.

2. **Integration of the Madabhushi–Lambert primary breakup model.** The Madabhushi model [6] treats the liquid column through a blob injection approach, originally rooted in the Wave atomisation framework of Reitz [7]: parcels are initialised at the nozzle exit with size equal to the injector diameter, and Kelvin–Helmholtz (KH) wave instability governs the shedding of child droplets from the column surface. Column breakup time and penetration distance are predicted by empirical correlations derived from LJICF experiments. Lambert's subsequent enhancement [8] introduces a ligament breakup mechanism after column fragmentation, intended to reduce the overprediction of child droplet diameters that arises in the intermediate ligament-shedding regime. Both contributions are coupled within a unified breakup-model state machine.

3. **Custom drag model for deforming droplets.** A deformation-dependent drag correlation is implemented to account for the progressive shape distortion of large droplets under aerodynamic loading, improving the accuracy of droplet trajectory predictions in the near-field region of the jet.

4. **Secondary breakup via the Pilch–Erdman model.** Secondary atomisation of child droplets is governed by acceleration-induced breakup. The Pilch–Erdman model [9] predicts the total breakup time and the maximum stable fragment size from velocity-history data and the instantaneous aerodynamic Weber number, covering multiple breakup regimes above the critical Weber number.

5. **Source-term smoothing via diffusion-based coarse-graining (DBM).** In the near-injector region, the effective volume and momentum exchange associated with the injected parcels may be too large compared with the refined wall-adjacent cells required to resolve the SST boundary layer. A standard point-source projection of the Lagrangian source terms may therefore produce unrealistically large local source/sink terms, leading to depletion of the continuous phase, density collapse, numerical divergence, and distorted crossflow penetration. The diffusion-based coarse-graining method of Sun and Xiao [10] is implemented to redistribute the

Lagrangian source terms over neighbouring cells through a conservative screened-diffusion operator, reducing the excessive mesh sensitivity of the coupled Eulerian–Lagrangian solution. Grid refinement in the near-injector region remains necessary for proper boundary-layer resolution; DBM addresses the source-term concentration artefact that arises from this refinement.

6. **Assessment against the Lambert LJICF benchmark.** The accuracy of the complete model chain is assessed against the experimental LJICF dataset of Lambert [8].

7. **Per-campaign source-term projection study.** The DBM bandwidth $b$ is swept over $b/D_{\mathrm{inj}} \in \{0, 1, 3, 6\}$ on each of the four validation campaigns of this work—Lambert J10, Lambert J20, and the two Zhang gas-film cases at $We_g^{\mathrm{film}} \simeq 500$ and $We_g^{\mathrm{film}} \simeq 3500$. The value $b = 0$ recovers the bare PCM projection. The introduction of the diffusion-based smoothing was originally motivated by the Zhang $We_g^{\mathrm{film}} \simeq 500$ case, where the combination of low gas velocity and large parcel diameters causes local gas-cell depletion of the continuous phase. The objective is not to tune $b$ as a physical closure, but to identify, for each campaign, the trade-off between droplet-size statistics, parcel velocity profiles and spray penetration.

8. **Lambert J10 and J20 PCM and DBM analysis.** Both Lambert configurations are analysed with the per-campaign DBM sweep above, using the `cell` carrier-to-parcel variant for sampling the gas state at the parcel position. The two configurations are treated independently, so that the trade-off between $D_{32}$ and $U_x$ is identified per momentum-flux ratio rather than being inherited from a single reference case.

9. **Extension to the Zhang LJIGF/GLOP configuration.** The transferability of the model is explored on the liquid jet in subsonic gas-film configuration studied by Zhang et al. [3], representing a gas-liquid orifice type pintle simplified injection element. The two Zhang cases at $We_g^{\mathrm{film}} \simeq 500$ and $We_g^{\mathrm{film}} \simeq 3500$ are run with the same DBM bandwidth sweep used for the Lambert campaigns. The PCM interpolation is restricted here to the `cellPoint` variant, since the local gas-cell depletion first observed in the $We_g^{\mathrm{film}} \simeq 500$ case makes the `cell` variant unviable. The bare-PCM limit ($b = 0$) is retained as a reference setup, and is shown to produce a spray distribution inconsistent with the expected near-injector physics. The $We_g^{\mathrm{film}} \simeq 100$ and $We_g^{\mathrm{film}} \simeq 8500$ cases reported in the Zhang campaign are not included in the present comparison; the former is outside the validated operating range, while the latter could not be completed within the present framework and its failure mechanism was not isolated. Extending the screened-diffusion treatment to additional source terms is identified as a future development, to be assessed through dedicated tests.

## 1.3 Structure of the Thesis

The remainder of this thesis is organised as follows.

Chapter 2 presents the physical background of LJICF and LJIGF atomisation and the mathematical formulation of the Lagrangian spray model, integrating literature, governing physics, and thesis-specific modelling choices for each model component.

Chapter 3 details the OpenFOAM implementation, covering the software architecture, the extended parcel and cloud class hierarchy, the injection boundary conditions, the breakup and drag state-machine logic, and the numerical procedures adopted for source-term smoothing.

Chapter 4 presents the validation of the implemented model against the Lambert LJICF benchmark. Both the J10 and J20 configurations are analysed with a per-campaign source-term projection study using the `cell` PCM interpolation variant and the DBM bandwidth $b/D_{\text{inj}} \in \{0, 1, 3, 6\}$, where $b = 0$ corresponds to the bare PCM projection. Combined profile comparisons, error metrics and parcel-field visualisations are used to identify, for each Lambert configuration, the trade-off between $D_{32}$ and $U_x$.

Chapter 5 explores the transferability of the model to the Zhang LJIGF/GLOP configuration. The same bandwidth sweep used in the Lambert campaigns is applied to the two Zhang cases at $We_g^{\text{film}} \simeq 500$ and $We_g^{\text{film}} \simeq 3500$, restricted to the `cellPoint` interpolation, since the $We_g^{\text{film}} \simeq 500$ case originally exposed the local gas-cell depletion that motivated the diffusion-based smoothing and made the `cell` variant unviable. The chapter therefore quantifies the necessary role of DBM in preventing source-term concentration and unphysical jet behaviour in the confined gas-film geometry.

Chapter 6 summarises the main findings and contributions of the work, discusses the limitations of the current implementation, and outlines directions for future extensions and improvements.

# 2 Physical Background and Lagrangian Model Formulation

This chapter establishes the physical context and mathematical framework of the Lagrangian spray model implemented in this thesis. The liquid jet in crossflow is described through a sequence of sub-models which reproduce the main stages of the atomisation process: liquid-column motion and wave shedding, column breakup, ligament-corrected primary-fragment formation, secondary aerodynamic breakup, and state-dependent droplet drag. For each component, the relevant physical background, governing relations, and modelling choices are presented together, so that the theoretical formulation can be directly connected to the OpenFOAM implementation discussed in Chapter 3.

## Notation Policy

The notation used in this chapter follows the original references whenever possible, but it is adapted where necessary to remain consistent with the model variables adopted in this thesis. This is particularly important because the formulation combines correlations from Madabhushi [6], Reitz [7], Pilch and Erdman [9], and Lambert et al. [8], whose symbols are not always mutually consistent.

- The Madabhushi column-breakup constant is denoted by $C_0$. The column-breakup time is denoted by $t_{cb}$, consistently with the notation used in the present model formulation.
- The Reitz Wave model is written using the injector radius $a = D_{\text{inj}}/2$ as characteristic length scale. Reitz's liquid and gas Weber numbers $We_1$ and $We_2$ are denoted here by $We_l$ and $We_g$, respectively. The Ohnesorge number $Z$ is written as $Oh$, while the Taylor parameter is denoted by $Ta = Oh\sqrt{We_g}$.
- The Pilch–Erdman characteristic time is denoted by $t^*$, as in the original formulation. The quantities latched at the onset of a Pilch–Erdman breakup event are written as $D_{p,0}$ and $U_{\text{rel},0}$, corresponding to the parent diameter and relative velocity at the beginning of the breakup event.
- The parcel lifetime since injection is denoted by $t_c$, while the onset time of a Pilch–Erdman event is denoted by $t_{\text{sec}}$. These quantities are used to distinguish the global parcel residence time from the local PE deformation clock.

A specific ambiguity concerns the Weber number notation. In this thesis, $We_g$ is used

only when referring to the gas-side Weber number of the Kelvin–Helmholtz Wave model, based on the injector radius $a$. The aerodynamic Weber number controlling Pilch–Erdman deformation and breakup is instead denoted by $We$, because it is based on the current parcel diameter. When discussing the Zhang et al. [3] gas-film configuration, the gas-film Weber number is explicitly identified as $We_g^{\mathrm{film}}$ or as the case parameter used in the reference paper, since it is based on the gas film height rather than on either the injector radius or the droplet diameter. The orifice-based gas Weber number used to discuss the Wu LJICF regimes is

$$We_g^{D_{\mathrm{inj}}} = \frac{\rho_g U_g^2 D_{\mathrm{inj}}}{\sigma}, \tag{2.1}$$

whereas the gas-film Weber number used for the Zhang LJIGF/GLOP configuration is

$$We_g^{\mathrm{film}} = \frac{\rho_g U_g^2 W}{\sigma}, \tag{2.2}$$

with $U_g$ the bulk gas-film velocity and $W$ the gas-film height.

## 2.1 Liquid-Jet-in-Crossflow Breakup Physics

The liquid jet in crossflow (LJICF) is a canonical atomisation configuration of direct relevance to aerospace propulsion systems, where liquid fuel is injected transversely into a high-speed gaseous stream. A liquid jet of bulk velocity $U_l$ issues from a circular orifice of diameter $D_{\mathrm{inj}}$ and interacts with a gaseous crossflow of velocity $U_g$ and density $\rho_g$. The initially transverse liquid momentum is progressively redirected by aerodynamic loading, the column bends downstream, surface waves develop along the gas–liquid interface, and the jet eventually disintegrates into ligaments and droplets.

From a modelling standpoint, LJICF atomisation is particularly challenging because several mechanisms coexist over short length and time scales. The intact column behaves as a coherent liquid structure, but it simultaneously sheds droplets from its surface, deforms under aerodynamic loading, fractures into large fragments and ligaments, and finally produces a dilute spray whose evolution is governed by drag, turbulent dispersion, and secondary breakup. This sequence is the physical basis of the Madabhushi model adopted in this work: the liquid core is represented by Lagrangian parcels with the injector diameter, surface stripping is represented through a Kelvin–Helmholtz Wave model, column breakup triggers a Pilch–Erdman type aerodynamic breakup stage, and the Lambert correction accounts for the fact that the first fragments after column breakup are ligament-like rather than spherical.

### 2.1.1 Stages of Breakup

Three spatial regions are conventionally distinguished in LJICF atomisation. In the liquid-column region, the jet remains connected to the injector and behaves as a coherent structure, although its surface is already affected by aerodynamic waves and stripping. Wu et al. [2] showed that, for larger injection velocity and momentum-flux-ratio

conditions, liquid jets penetrate farther into the crossflow and exhibit surface breakup before the column fractures. The liquid-column penetration is mainly controlled by the liquid-to-air momentum flux ratio, while column fracture can be interpreted through an aerodynamic-breakup time scale.

Downstream of the intact column, a ligament and primary-fragment region is formed. This region contains large, highly deformed liquid structures generated by column fracture, together with droplets stripped from the column surface. It is precisely this region that motivates the extension proposed by Lambert et al. [8]: the original Madabhushi formulation tends to overestimate the child-droplet diameter immediately after column breakup, because the first post-column fragments are not spherical droplets but ligament-like structures. The Lambert correction therefore reduces the first child-droplet diameters generated after column breakup by a ligament factor $F_{\mathrm{Lig}}$.

Farther downstream, the spray becomes progressively more dilute and the liquid phase is represented by dispersed droplets. In this region, parcel trajectories are mainly controlled by aerodynamic drag and turbulent dispersion, while droplet size continues to evolve through secondary aerodynamic breakup when the local Weber number remains above the stability threshold.

### 2.1.2 Column Stripping and Column Fracture

Two liquid-removal mechanisms must be distinguished. The first one is surface stripping, where interfacial waves grow on the liquid column and shed droplets or ligaments from its periphery. Wu et al. [2] observed that, under sufficiently strong aerodynamic loading, droplets can be stripped from the deflected edges of the liquid column before the development of large-scale column waves. This mechanism is represented in the Madabhushi formulation by the Wave breakup model, which continuously removes mass from the parent column parcel and creates wave-shed child parcels.

The second mechanism is column fracture, which terminates the intact-core representation and marks the transition to the primary-fragment and ligament region. In the Madabhushi–Lambert formulation, the column does not break at an empirically imposed spatial location. Instead, it remains in the column state until the parcel residence time reaches the column-breakup time $t_{cb}$, defined by Equation (2.35). This is consistent with the physical interpretation of Wu et al. [2], who related the fracture process of the liquid column to the time scale of aerodynamic secondary breakup of an analogous droplet.

### 2.1.3 Controlling Parameters

The global deflection of a transverse liquid jet is primarily governed by the jet-to-crossflow momentum flux ratio,

$$J = \frac{\rho_l U_l^2}{\rho_g U_g^2}, \tag{2.3}$$

where $U_l$ is the bulk injection velocity and $U_g$ is the crossflow velocity. At fixed geometry and fluid properties, increasing $J$ increases the penetration of the liquid column and

delays its bending toward the streamwise direction. In the Wu et al. experiments, the column trajectory and fracture height were correlated using the liquid-to-air momentum flux ratio $\bar{q}$, which is denoted as $J$ in the present thesis.

No empirical trajectory correlation is imposed in the present model: the liquid penetration and breakup location emerge from the Lagrangian parcel dynamics, the blob-drag treatment of the intact column, and the time-based column-breakup criterion.

The aerodynamic Weber number based on the current parcel diameter is defined as

$$We = \frac{\rho_g U_{\mathrm{rel}}^2 d_p}{\sigma}, \tag{2.4}$$

where $U_{\mathrm{rel}}$ is the magnitude of the slip velocity between the parcel and the surrounding gas, $d_p$ is the current parcel diameter, and $\sigma$ is the liquid surface tension. This Weber number is the local quantity used to determine whether a droplet or fragment is aerodynamically unstable. It is therefore distinct from the orifice-based gas Weber number used to characterise the experimental operating condition and from the gas-side Weber number used in the Reitz Wave model.

Liquid viscosity enters through the Ohnesorge number,

$$Oh = \frac{\mu_l}{\sqrt{\rho_l d_p \sigma}}, \tag{2.5}$$

or, equivalently, through the ratio $\sqrt{We_l}/Re_l$ when liquid Weber and Reynolds numbers are defined on the same diameter and velocity scale. In the secondary-breakup model, $Oh$ modifies the critical Weber number and the child-size correlation, reducing the tendency of highly viscous droplets to break up.

A key statistical quantity used to compare numerical and experimental sprays is the Sauter Mean Diameter,

$$D_{32} = \frac{\sum_i N_i d_i^3}{\sum_i N_i d_i^2}, \tag{2.6}$$

where $d_i$ is the droplet diameter and $N_i$ is the statistical weight of the corresponding parcel. This definition is consistent with the Lagrangian post-processing used later in the validation chapters, where each computational parcel represents a finite number of physical droplets.

### 2.1.4 Breakup-Regime Interpretation

The qualitative breakup mode depends on the level of aerodynamic loading. Wu et al. [2] interpreted liquid-column breakup in crossflow as closely related to aerodynamic secondary breakup of droplets: the column first deforms, then flattens, and finally disintegrates under gas loading. In their experiments, an orifice-based gas Weber number $We_g^{D_{\mathrm{inj}}}$ of about 60 produced several bag-shaped membranes, similar to multimode droplet breakup, whereas larger $We_g^{D_{\mathrm{inj}}}$ conditions produced acceleration-wave and shear-like breakup processes. The definition of $We_g^{D_{\mathrm{inj}}}$ is given in Equation (2.1).

In the present thesis this regime description is used only as physical background. The implemented model does not switch between fixed LJICF regimes according to an orifice-based Weber-number map. Instead, the active sub-model is determined by the parcel state: intact-core parcels undergo Wave stripping until the column-breakup time $t_{cb}$, post-column fragments enter a Pilch–Erdman-type breakup sequence, and child droplets are re-tested against the local stability criterion before any further breakup event is allowed.

### 2.1.5 Liquid Jet in Subsonic Gas Film Variant

A closely related configuration is the liquid jet in subsonic gas film (LJIGF), which appears in gas–liquid orifice-type pintle injection elements. In this case, the gas stream has a finite height $W$, and the liquid jet can penetrate a region whose thickness is comparable with the gas-film height. The resulting flow is therefore not equivalent to an unbounded LJICF: the gas film can cover only part of the liquid structure, the lower wall and gas-film boundary layer affect the breakup, and the atomisation morphology depends strongly on the local momentum ratio and on the gas-film Weber number $We_g^{\mathrm{film}}$, defined in Equation (2.2).

Zhang et al. [3] showed that, at approximately fixed local momentum ratio, increasing the gas-film Weber number changes the dominant instability and the downstream droplet distribution. Low gas-film Weber numbers produce larger unstable liquid structures and poorer atomisation, whereas higher values generate denser and smaller droplets through stronger shear and surface-wave activity. At the same time, the column-breakup location can remain nearly unchanged when the local momentum ratio is held fixed. This behaviour is interpreted as the result of competing effects between stronger gas loading, liquid acceleration, and the breakup time scale. The Zhang GLOP case considered in this thesis is therefore used as a complementary reference to assess the behaviour of the implemented Lagrangian model outside the classical unbounded LJICF configuration.

## 2.2 Eulerian–Lagrangian Spray Modelling Framework

The dispersed two-phase flow is described within an Eulerian–Lagrangian (EL) framework, commonly referred to as the Discrete Phase Model (DPM) in commercial CFD terminology. The carrier gas is solved on a fixed Eulerian finite-volume mesh, whereas the liquid phase is tracked in a Lagrangian sense through a statistical ensemble of computational parcels. Each parcel represents a finite number of identical physical droplets or liquid fragments, whose position, velocity, diameter and thermophysical state are advanced along their trajectory.

This modelling choice is particularly convenient for liquid-jet-in-crossflow simulations once the liquid phase is represented as a dispersed spray. The continuous gas field remains Eulerian, while the liquid structures generated by primary and secondary breakup are represented by parcels whose diameters and statistical weights evolve according to the breakup model. The approach therefore does not resolve the gas–

liquid interface explicitly; instead, the unresolved atomisation processes are represented through drag, breakup and phase-coupling sub-models.

### 2.2.1 Carrier-Phase Equations

In the solver used in this thesis, the carrier phase is advanced by a transient finite-volume formulation derived from the compressible Navier–Stokes equations. For the cold-flow, low-Mach-number cases considered here, density variations are limited, but the numerical formulation remains that of the compressible spray solver. In compact conservative form, the carrier-phase equations can be written as

$$\frac{\partial \rho_g}{\partial t} + \nabla \cdot \left( \rho_g \mathbf{u}_g \right) = S_m, \tag{2.7}$$

$$\frac{\partial \rho_g \mathbf{u}_g}{\partial t} + \nabla \cdot \left( \rho_g \mathbf{u}_g \mathbf{u}_g \right) = -\nabla p + \nabla \cdot \boldsymbol{\tau}_{\mathrm{eff}} + \rho_g \mathbf{g} + \mathbf{S}_u, \tag{2.8}$$

where $\rho_g$, $\mathbf{u}_g$, and $p$ are the gas density, velocity and pressure, respectively. The tensor $\boldsymbol{\tau}_{\mathrm{eff}}$ contains molecular and turbulent viscous stresses. Under the Boussinesq eddy viscosity closure it may be written as

$$\boldsymbol{\tau}_{\mathrm{eff}} = \mu_{\mathrm{eff}} \left[ \nabla \mathbf{u}_g + \left( \nabla \mathbf{u}_g \right)^{\top} - \frac{2}{3} \left( \nabla \cdot \mathbf{u}_g \right) \mathbf{I} \right], \qquad \mu_{\mathrm{eff}} = \mu_g + \mu_t, \tag{2.9}$$

where $\mu_t$ is the turbulent eddy viscosity. The term $\mathbf{S}_u$ is the momentum source associated with the dispersed Lagrangian phase. In the non-evaporating configurations considered in this work, the mass source $S_m$ is negligible, whereas the momentum coupling can be significant in the dense near-injector region.

Turbulence closure is provided by the SST $k$–$\omega$ model of Menter [11]. This choice is appropriate for wall-bounded crossflow configurations because it combines near-wall sensitivity with robustness in separated and adverse-pressure-gradient regions, which are relevant near the injector wall and in the jet wake.

### 2.2.2 Parcel Representation and Equations of Motion

A computational parcel represents $N_p$ identical physical droplets or liquid fragments of diameter $d_p$, liquid density $\rho_l$, dynamic viscosity $\mu_l$, and surface tension $\sigma$. The parcel mass is therefore

$$m_p = N_p \rho_l \frac{\pi}{6} d_p^3, \tag{2.10}$$

where $N_p$ is the statistical parcel weight. At injection, $N_p$ is determined from the prescribed liquid mass flow rate, the parcel injection rate and the initial blob diameter. During deformation, the represented liquid mass is conserved, while the aerodynamic response may change through the drag model. During breakup and state transitions, the diameter, parcel mass and $N_p$ are updated consistently so that the total physical liquid mass represented by the parent and child parcels is conserved.

The Lagrangian trajectory of each parcel is governed by

$$\frac{\mathrm{d}\mathbf{x}_p}{\mathrm{d}t} = \mathbf{u}_p, \tag{2.11}$$

$$m_p \frac{\mathrm{d}\mathbf{u}_p}{\mathrm{d}t} = \mathbf{F}_D + m_p \mathbf{g}, \tag{2.12}$$

where $\mathbf{x}_p$ and $\mathbf{u}_p$ are the parcel position and velocity, $\mathbf{F}_D$ is the aerodynamic drag force, and $\mathbf{g}$ is the gravitational acceleration. Other hydrodynamic forces, such as added mass and pressure-gradient forces, are neglected in the present formulation. This is consistent with the large liquid-to-gas density ratio of the water–gas systems considered here, for which aerodynamic drag is the dominant coupling mechanism controlling parcel acceleration.

The local carrier velocity interpolated at the parcel position determines the relative velocity

$$\mathbf{U}_{\mathrm{rel}} = \mathbf{u}_g(\mathbf{x}_p, t) - \mathbf{u}_p, \qquad U_{\mathrm{rel}} = |\mathbf{U}_{\mathrm{rel}}|. \tag{2.13}$$

The same relative velocity enters both the drag law and the local Weber number used by the breakup criteria.

### 2.2.3 Two-Way Coupling and Dilute-Spray Hypothesis

The dispersed phase is treated as dilute in the sense that the liquid volume fraction inside each carrier cell remains small enough for parcel–parcel collisions and finite-volume displacement of the carrier phase to be neglected. Under this hypothesis, the primary feedback of the liquid phase on the gas phase is through momentum exchange. Formally, the liquid volume fraction in a cell of volume $V_{\mathrm{cell}}$ can be estimated as

$$\alpha_l = \frac{1}{V_{\mathrm{cell}}} \sum_{p \in \mathrm{cell}} N_p \frac{\pi}{6} d_p^3, \qquad \alpha_l \ll 1. \tag{2.14}$$

The gas–parcel momentum exchange is accumulated by the Lagrangian cloud and introduced into the carrier-phase momentum equation as a source term. In semi-implicit form, the source contribution can be interpreted as the sum of an explicit momentum-transfer term and an implicit linearised drag coefficient,

$$\mathbf{S}_u \simeq \mathbf{S}_u^{\mathrm{exp}} - S_p^{\mathrm{imp}} \mathbf{u}_g. \tag{2.15}$$

The exact finite-volume implementation of this source term is discussed in Chapter 3.

A standard way to map Lagrangian source terms onto the Eulerian mesh is the Particle Centroid Method (PCM), where the parcel contribution is deposited into the cell containing the parcel centroid. Although straightforward and conservative, this approach may generate excessively concentrated source terms in small cells, especially near the injector where the parcel number density is high and the carrier mesh is refined. In the present thesis, this issue is addressed by applying a diffusion-based coarse-graining procedure to the accumulated momentum source fields before they enter the Eulerian momentum solution. The method is discussed in Section 2.8.

Although the dilute hypothesis allows parcel–parcel collisions to be neglected in the momentum equations, a collision model is retained to account for droplet coalescence in the denser near-injector region. The `trajectory` stochastic collision model of Nordin [12], as implemented in OpenFOAM [5], is activated with coalescence enabled: collision probability is evaluated stochastically from the relative parcel trajectories within each Eulerian cell, and a coalescence outcome is applied when the collision Weber number falls below the coalescence threshold. The model operates on the same parcel representation as the breakup and drag models, and its contribution enters the two-way coupling through the standard source-term accumulation.

### 2.2.4 Strengths and Limitations of the EL Approach

The main advantage of the Eulerian–Lagrangian approach is that it allows practical spray simulations over engineering-scale domains without resolving the full gas–liquid interface. This makes it suitable for parametric studies of LJICF and LJIGF configurations, provided that the unresolved physics of primary breakup, secondary breakup, drag and turbulent dispersion are represented by appropriate sub-models. Its main limitation is that the predicted spray structure depends directly on these closures and on the mapping of Lagrangian source terms to the Eulerian mesh.

Interface-resolved approaches, such as VOF, CLSVOF or hybrid VOF–LPT methods, can describe the near-field liquid topology with much greater detail, but at a significantly higher computational cost. The Zhang reference configuration, for example, was generated with a VOF–LPT strategy [3], in which continuous liquid structures are resolved before sufficiently small droplets are converted into Lagrangian particles. In contrast, the present work focuses on an engineering EL formulation and improves its predictive behaviour through a consistent breakup state machine, a state-dependent drag model, and a diffusion-based treatment of the Eulerian source terms.

## 2.3 Blob Representation and Reitz KH/Wave Surface Stripping

The intact liquid column emerging from the injector is not resolved as a continuous liquid–gas interface. Resolving the near-nozzle liquid core and its surface topology would require an interface-resolved method and a mesh resolution that is not compatible with the engineering-scale Eulerian– Lagrangian framework adopted in this thesis. Instead, the *blob-injection paradigm* introduced by Reitz [7] and adopted by Madabhushi [6] represents the liquid column by a stream of Lagrangian parcels, or *blobs*, whose initial diameter is equal to the injector orifice diameter:

$$D_0 = D_{\text{inj}}. \tag{2.16}$$

Each blob is injected at the liquid bulk velocity $U_l$, directed transversely to the crossflow. During the liquid-column stage, the blob is not interpreted as an isolated spherical droplet, but as a Lagrangian surrogate for a portion of the intact liquid column. The effects of surface stripping, column breakup and subsequent secondary breakup are then introduced through sub-models acting on the parcel state.

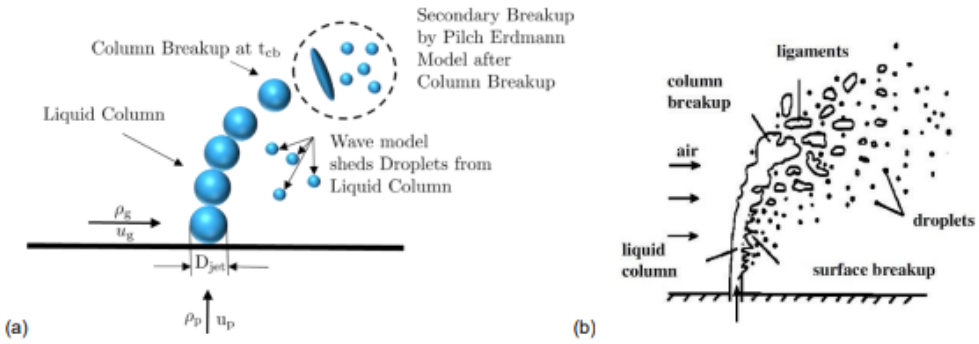



**Figure 2.1:** *Madabhushi breakup model and LJICF breakup structure [2, 8].*

As shown in Figure 2.1, the model separates the liquid-column stage from the post-column breakup stage. Before the column-breakup time $t_{cb}$, the parent blob remains in the liquid-column regime and undergoes Wave stripping from its surface. Once $t_{cb}$ is reached, the remaining core parcel enters a Pilch–Erdman-type secondary-breakup sequence. The droplets shed from the column surface and the fragments generated after column breakup are then tracked as child parcels, whose subsequent evolution depends on the local drag and breakup criteria.

### 2.3.1 Linear Stability Background

The Wave model is based on the temporal stability analysis of a cylindrical liquid jet subject to aerodynamic forcing by the surrounding gas. In the formulation of Reitz [7], infinitesimal disturbances are imposed on the liquid surface and a dispersion relation is obtained for the disturbance growth rate as a function of wavenumber. The full dispersion relation is not solved during the CFD calculation. In physical terms, the liquid surface is perturbed by modes of the form $\eta \propto \exp(ikz + \omega t)$, and the dominant disturbance is selected from the wavenumber that maximises the growth rate, $k^* = \arg\max_k \mathrm{Re}[\omega(k)]$. Algebraic fits of this most unstable solution are used to provide the wavelength $\Lambda_{\mathrm{KH}} = 2\pi/k^*$ and the growth rate $\Omega_{\mathrm{KH}}$ of the dominant surface wave.

In the Madabhushi formulation, these correlations are used to represent continuous droplet shedding from the surface of the liquid column. The most unstable wavelength defines the characteristic size of the droplets stripped from the column, while the corresponding growth rate enters the breakup timescale controlling the reduction of the parent blob diameter.

### 2.3.2 Dimensionless Groups

The KH/Wave correlations are written in terms of the injector radius $a = D_{\mathrm{inj}}/2$ and of the local gas–parcel relative velocity $U_{\mathrm{rel}}$. The dimensionless groups are

$$We_g = \frac{\rho_g \, U_{\mathrm{rel}}^2 \, a}{\sigma}, \tag{2.17}$$

$$We_l = \frac{\rho_l \, U_{\mathrm{rel}}^2 \, a}{\sigma}, \tag{2.18}$$

$$Re_l = \frac{\rho_l \, U_{\mathrm{rel}} \, a}{\mu_l}, \tag{2.19}$$

$$Oh = \frac{\sqrt{We_l}}{Re_l} = \frac{\mu_l}{\sqrt{\rho_l \, a \, \sigma}}, \qquad Ta = Oh \sqrt{We_g}. \tag{2.20}$$

Here $We_g$ is the gas-side Weber number driving the surface-wave instability, $We_l$ is the liquid Weber number, $Re_l$ is the liquid Reynolds number, $Oh$ is the Ohnesorge number and $Ta$ is the Taylor parameter. With respect to the original notation of Reitz [7], the present $We_g$, $We_l$, $Oh$, and $Ta$ correspond to $We_2$, $We_1$, $Z$, and $T$, respectively.

The symbol $We_g$ in Equation (2.17) must not be confused with the gas-film Weber number used in the Zhang GLOP/LJIGF configuration. The present definition is based on the injector radius and on the local parcel slip velocity, whereas the Zhang gas-film Weber number is based on the gas-film height and on the bulk gas-film velocity. Both are also distinct from the parcel-diameter-based aerodynamic Weber number $We$ used by the Pilch–Erdman breakup model.

### 2.3.3 Most-Unstable Wavelength and Growth Rate

The empirical correlations used by Reitz [7] and by Madabhushi [6] give the most unstable wavelength as

$$\Lambda_{\mathrm{KH}} = 9.02 \, a \, \frac{\left(1 + 0.45\sqrt{Oh}\right)\left(1 + 0.4\, Ta^{0.7}\right)}{\left(1 + 0.87\, We_g^{1.67}\right)^{0.6}}, \tag{2.21}$$

and the corresponding growth rate as

$$\Omega_{\mathrm{KH}} = \frac{0.34 + 0.38\, We_g^{1.5}}{(1 + Oh)\left(1 + 1.4\, Ta^{0.6}\right)} \sqrt{\frac{\sigma}{\rho_l \, a^3}}. \tag{2.22}$$

Increasing $We_g$ increases the aerodynamic forcing of the surface disturbance and reduces the characteristic wavelength, promoting finer surface stripping. Increasing $Oh$ introduces stronger viscous damping and therefore reduces the growth of the instability.

### 2.3.4 KH Stripping, Child Emission and Velocity Assignment

The characteristic Wave stripping timescale is

$$\tau_{\mathrm{KH}} = \frac{3.726 \, B_1 \, a}{\Lambda_{\mathrm{KH}} \Omega_{\mathrm{KH}}}, \tag{2.23}$$

where $B_1$ is the Wave breakup time constant. The Madabhushi/Reitz default $B_1 = 10.0$ is retained for the Lambert LJICF cases, while the Zhang LJIGF/GLOP cases adopt a separately calibrated value reported in Chapter 5.

Wave stripping is admitted only when the local gas-side Weber number exceeds the activation threshold

$$We_g > 6.0. \tag{2.24}$$

This is an implementation-level activation gate for the Reitz–Madabhushi Wave stripping model [6, 7], following the reference OpenFOAM Wave model [5]; it is not an explicit threshold in the original references. The characteristic child-droplet diameter associated with the most unstable wave is

$$d_c = 2B_0\Lambda_{\mathrm{KH}}, \tag{2.25}$$

where $B_0$ is the KH/Wave child-size coefficient. Although the Reitz–Wave calibration commonly uses $B_0 = 0.61$, the coefficient is retained as an adjustable model parameter in this thesis. In the present LJICF validation cases, the Lambert calibration $B_0 = 2.44$ is adopted. The value $B_0 = 2.44$ should therefore be interpreted as an effective post-coalescence diameter coefficient calibrated by Lambert et al. on LJICF child-size statistics, rather than as a literal radius coefficient in the Reitz sense. The Zhang LJIGF/GLOP configuration requires a separate calibration because the finite gas-film geometry modifies the near-field stripping and atomisation behaviour. The adopted values are reported in the corresponding validation chapters.

A small numerical lower bound is imposed in the simulations to avoid degenerate child diameters. A wave-shedding event is allowed only if

$$d_c \le D_{\mathrm{inj}}, \qquad d_c < d_p. \tag{2.26}$$

These constraints prevent the Wave model from creating droplets larger than the injected column diameter or larger than the current parent parcel. This constraint is particularly relevant when $B_0 > 1$, as in the Lambert calibration $B_0 = 2.44$ adopted for the LJICF validation cases: in the early phase after injection, the most unstable wavelength $\Lambda_{\mathrm{KH}}$ is comparable with or larger than $D_{\mathrm{inj}}$, so $2B_0\Lambda_{\mathrm{KH}}$ may initially exceed $D_{\mathrm{inj}}$. Wave shedding is therefore suppressed until the gas-side aerodynamic loading has reduced $\Lambda_{\mathrm{KH}}$ below $D_{\mathrm{inj}}/(2B_0)$.

When the Wave model is active and $d_p > d_c$, the parent blob diameter relaxes toward the target child diameter. In diameter form, the relaxation can be written as

$$\frac{\mathrm{d}d_p}{\mathrm{d}t} = -\frac{d_p - d_c}{\tau_{\mathrm{KH}}}. \tag{2.27}$$

This is equivalent to the original Reitz formulation written in terms of the parent radius, since $r = d/2$ and $r_c = d_c/2$ imply

$$\frac{\mathrm{d}r}{\mathrm{d}t} = -\frac{r - r_c}{\tau_{\mathrm{KH}}} \quad \Longleftrightarrow \quad \frac{\mathrm{d}d_p}{\mathrm{d}t} = -\frac{d_p - d_c}{\tau_{\mathrm{KH}}}. \tag{2.28}$$

In the implemented time integration, the diameter update is

$$d_p^{n+1} = \frac{(\Delta t/\tau_{\mathrm{KH}})\, d_c + d_p^n}{1 + \Delta t/\tau_{\mathrm{KH}}}. \tag{2.29}$$

The mass removed from the parent during the time step is accumulated in the stripped-mass reservoir $m_s$:

$$\Delta m_s = N_p \rho_l \frac{\pi}{6} \left[ \left( d_p^n \right)^3 - \left( d_p^{n+1} \right)^3 \right]. \tag{2.30}$$

After this update, the parent parcel mass is made consistent with the reduced diameter and the statistical parcel weight.

A wave-shed child parcel is created when the accumulated stripped mass exceeds five percent of the initial wave-stage parcel mass:

$$m_s > 0.05\, m_{0,\mathrm{wave}}. \tag{2.31}$$

This criterion follows the computational shedding rule used in the Madabhushi model. The child diameter is set to $d_c$, while the child mass is set by the emission threshold. The child statistical weight is therefore computed from mass consistency:

$$N_p^{\mathrm{child}} = \frac{m_{\mathrm{child}}}{\rho_l (\pi/6) d_c^3}. \tag{2.32}$$

After the child parcel is created, its mass is subtracted from the stripped-mass reservoir so that the parent–child mass budget remains consistent.

The wave-shed child inherits the parent velocity plus a transverse perturbation associated with the surface-wave motion. In vector form, this can be written as

$$\mathbf{u}_p^{\mathrm{child}} = \mathbf{u}_p + A_1 \Lambda_{\mathrm{KH}} \Omega_{\mathrm{KH}}\, \hat{\mathbf{n}}, \tag{2.33}$$

where $A_1 = 0.188$ is the empirical angular-spread coefficient and $\hat{\mathbf{n}}$ is a randomly sampled unit direction in the plane normal to the parent velocity. The product $\Lambda_{\mathrm{KH}} \Omega_{\mathrm{KH}}$ has dimensions of velocity and sets the scale of the transverse perturbation. Equivalently, the perturbation can be expressed through the half-angle tangent

$$\tan \theta_{1/2} = \frac{A_1 \Lambda_{\mathrm{KH}} \Omega_{\mathrm{KH}}}{\left| \mathbf{u}_p \right|}. \tag{2.34}$$

## 2.4 Madabhushi Column-Breakup Model

Surface stripping, described in Section 2.3, removes liquid from the periphery of the blob parcel but does not by itself terminate the intact liquid core. Column breakup represents the separate event by which the coherent liquid-column representation is ended and the remaining core parcel is passed to an aerodynamic secondary-breakup sequence. In the model proposed by Madabhushi [6], the liquid jet in crossflow is represented within a Lagrangian framework that combines the Reitz KH/Wave stripping model with an empirical column-breakup timescale derived from LJICF observations. This allows the model to describe, within a single parcel-based formulation, the bending of the intact liquid column, continuous surface shedding, column fracture and subsequent secondary breakup.

### 2.4.1 Column-Breakup Time and State Transition

The column-breakup time is written as

$$t_{cb} = C_0 \frac{D_{\text{inj}}}{\left|\mathbf{u}_g(\mathbf{x}_p, t)\right|} \sqrt{\frac{\rho_l}{\rho_g(\mathbf{x}_p, t)}}, \tag{2.35}$$

where $C_0$ is the empirical column-breakup constant of Madabhushi [6]. The default value $C_0 = 3.44$ is adopted for the Lambert LJICF cases, while the Zhang LJIGF/GLOP cases use a separately calibrated value reported in Chapter 5. The structure of Equation (2.35) reflects the aerodynamic nature of the breakup process: increasing the gas velocity shortens the time required to fracture the column, while increasing the liquid-to-gas density ratio delays the liquid response to the imposed gas loading.

This local-state evaluation is a deliberate departure from the original Madabhushi formulation, where $U_g$ and $\rho_g$ are interpreted as uniform free-stream crossflow values. The choice is consistent with the Eulerian–Lagrangian coupling logic adopted here, where every parcel sees the carrier-phase state interpolated at its current position. Two consequences follow: first, in confined LJIGF/GLOP geometries, the gas-film boundary layer can locally reduce $|U_g|$ near the wall, increasing $t_{cb}$ and delaying column breakup in that region; second, two-way momentum coupling feeds back on the breakup clock, since strong local gas deceleration in the recirculation behind the jet modifies the local value of $t_{cb}$. This behaviour cannot be represented by a purely free-stream-based column-breakup time.

When the parcel lifetime $t_c$ reaches $t_{cb}$, the core parcel is no longer treated as an intact liquid-column parcel. In the present model, this event is handled as a state transition rather than as an instantaneous multi-fragment breakup. First, any residual stripped mass accumulated during the KH/Wave stage is reabsorbed into the parent parcel, so that the total represented liquid mass remains consistent. The parent diameter is then recomputed from the updated mass and statistical weight.

After this mass-consistency update, the parcel is latched into the first post-column Pilch–Erdman sequence by storing the parent diameter $D_{p,0}$, the relative velocity $U_{\text{rel},0}$, and the PE onset time. No new child parcel is created at the instant $t_c = t_{cb}$; the actual fragmentation into child parcels occurs later, when the latched Pilch–Erdman sequence reaches its breakup time $t_b$. This distinction separates the end of the intact-column representation from the subsequent aerodynamic fragmentation process.

## 2.5 Lambert Ligament Correction

Lambert et al. [8] identified a limitation of the original Madabhushi formulation in the region immediately after column breakup. The liquid structures produced by the breakup of the column are better interpreted as ligaments or large deformed fragments than as compact spherical droplets. Applying the standard Pilch–Erdman size correlation directly to these structures can therefore overestimate the first post-column child-droplet diameter and lead to an overprediction of the Sauter Mean Diameter.

### 2.5.1 Ligament Correction and Model Scope

The Pilch–Erdman correlations are secondary-breakup correlations for aerodynamically deformed drops. In LJICF atomisation, however, the region immediately after column fracture contains ligament-like structures generated by the primary breakup of the liquid column. These structures are not well represented by a compact spherical parent droplet. The Lambert correction therefore modifies only the first post-column PE breakup event, where the ligament interpretation is most relevant.

The correction is introduced as a multiplicative factor applied to the nominal PE child diameter:

$$\bar{d}_{\mathrm{child}}^{\mathrm{Lambert}} = F_{\mathrm{Lig}}\, \bar{d}_{\mathrm{child}}^{\mathrm{PE}}, \tag{2.36}$$

where $F_{\mathrm{Lig}} \in (0, 1)$. In the present implementation, the default value is

$$F_{\mathrm{Lig}} = 0.4, \tag{2.37}$$

consistent with the Lambert calibration used for the LJICF J10/J20 reference cases [8]. The Zhang LJIGF/GLOP cases adopt a separately calibrated value reported in Chapter 5, because transferability of $F_{\mathrm{Lig}}$ to the confined gas-film configuration is not assumed a priori.

The correction is applied only to the first post-column PE breakup of the core parcel. After this first post-column event, all subsequent child parcels follow the standard Pilch–Erdman breakup logic without further ligament correction. This restriction prevents wave-shed children and later secondary fragments from incorrectly receiving the Lambert factor.

The advantage of the Lambert correction is its simplicity: it improves the initial post-column child-size distribution without modifying the KH/Wave stripping model or the column-breakup time. Its limitation is that $F_{\mathrm{Lig}}$ remains empirical and calibrated for a specific LJICF operating range. It does not explicitly model ligament stretching, ligament diameter, Rayleigh–Plateau breakup or a cascade of ligament-fragmentation events. Its transferability to the confined LJIGF/GLOP configuration is therefore not guaranteed a priori and is assessed through comparison with the Zhang reference case.

## 2.6 Pilch–Erdman Secondary Breakup

After column breakup, and also for sufficiently unstable wave-shed child parcels, aerodynamic secondary breakup is described using the Pilch–Erdman model [9]. The model addresses acceleration-induced breakup of liquid drops exposed to a relative flow and provides correlations for critical Weber number, breakup time, velocity history and maximum stable fragment size. In the present formulation, it is not used as a primary column-breakup model; it is activated only after the Madabhushi column-breakup latch or when a child parcel satisfies the standard PE admission criteria.

The original Pilch–Erdman correlations are based on collected experimental data for drops suddenly exposed to high relative velocities. Their use in a continuously evolving spray therefore requires an onset-latching procedure: once the parcel first satisfies the breakup admission criterion, the parent diameter and relative velocity at that instant

are stored and used as reference quantities throughout the PE event. This avoids continuously changing the PE timescale and fragment-size correlations during the deformation process.

### 2.6.1 Onset, Latching and Timescales

Pilch and Erdman distinguished several acceleration-induced breakup mechanisms, including vibrational breakup, bag breakup, bag-and-stamen breakup, sheet stripping, and wave-crest stripping followed by catastrophic breakup (Figure 2.2). These mechanisms are associated with increasing Weber number and describe the qualitative morphology of the breakup process. The piecewise correlations used for the total breakup time, however, should be interpreted as empirical calibration ranges for $t_b/t^*$, rather than as sharp universal morphological regime boundaries. The implemented model therefore uses the correlations to compute breakup timing and fragment properties, but does not explicitly switch between separate geometric breakup morphologies.

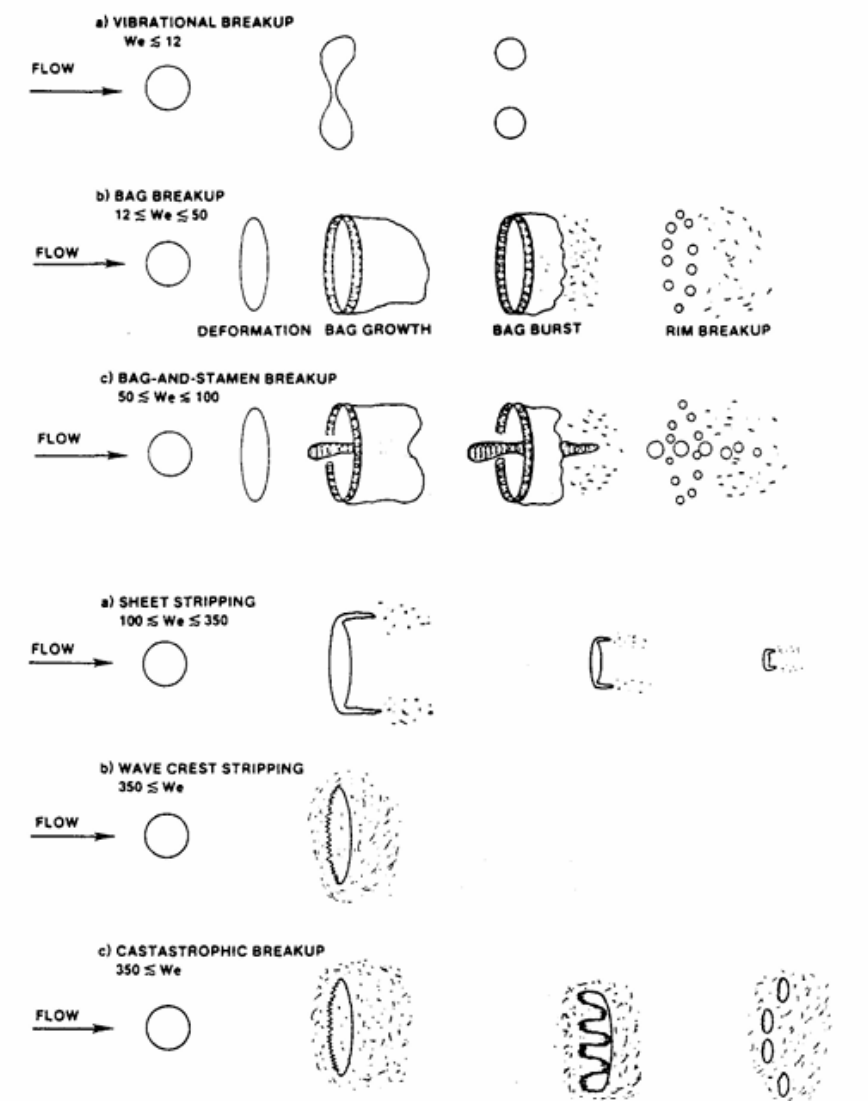



**Figure 2.2:** *Acceleration-induced secondary-breakup regimes as a function of aerodynamic Weber number, adapted from [9].*

At each Lagrangian time step, the parcel evaluates the aerodynamic Weber number using its current diameter and relative velocity:

$$We = \frac{\rho_g \, U_{\mathrm{rel}}^2 \, d_p}{\sigma}. \tag{2.38}$$

The onset of a PE breakup event is admitted when this Weber number exceeds the viscosity-corrected critical value

$$We > We_{\mathrm{crit}}, \qquad We_{\mathrm{crit}} = 12\left(1 + 1.077\, Oh^{1.6}\right). \tag{2.39}$$

The correction increases the critical Weber number for viscous droplets and recovers the inviscid critical value $We_{\mathrm{crit}} = 12$ as $Oh \to 0$.

At the onset instant, the model stores

$$D_{p,0}, \qquad U_{\mathrm{rel},0}, \qquad We_0, \qquad Oh_0. \tag{2.40}$$

These quantities are then held fixed during the corresponding PE event.

Wave-shed and post-column child parcels are not automatically forced to enter a new PE cycle. Before latching a standard PE event, the fragment is tested against both the critical Weber number and the maximum stable fragment diameter. The velocity-history estimate is

$$V_d = U_{\mathrm{rel}} \sqrt{\frac{\rho_g}{\rho_l}} \left(B_{1,PE}\, \bar{t}_b + B_{2,PE}\, \bar{t}_b^{\,2}\right), \qquad B_{1,PE} = 0.375, \quad B_{2,PE} = 0.2274, \tag{2.41}$$

where $\bar{t}_b = t_b/t^*$. The coefficients $B_{1,PE} = 0.375$ and $B_{2,PE} = 0.2274$ correspond to the non-oscillatory Pilch–Erdman velocity-history approximation used by the reference OpenFOAM `PilchErdman` model [5].[1] This second-order velocity-history estimate is sufficient for the child-fragment admission check used here. The corresponding maximum stable diameter, evaluated as in the reference OpenFOAM `PilchErdman` implementation [5], is computed as

$$d_{\mathrm{stable}} = \frac{12\,\sigma}{\rho_g U_{\mathrm{rel}}^2 \max\left[\left(\max\left(1 - \frac{V_d}{U_{\mathrm{rel}}}, 0\right)\right)^2, \varepsilon\right]}. \tag{2.42}$$

This expression is obtained by imposing the inviscid critical Weber number $We_c = 12$ on the residual relative velocity after fragment acceleration, $U_{\mathrm{rel}} \max(1 - V_d/U_{\mathrm{rel}}, 0)$:

$$\frac{\rho_g \left[U_{\mathrm{rel}} \max\left(1 - \frac{V_d}{U_{\mathrm{rel}}}, 0\right)\right]^2 d_{\mathrm{stable}}}{\sigma} = 12. \tag{2.43}$$

The inner clamp $\max(1 - V_d/U_{\mathrm{rel}}, 0)$ is essential: when the velocity-history correlation predicts $V_d \geq U_{\mathrm{rel}}$, which occurs near the onset threshold where $t_b/t^*$ diverges as $We_0 \to$

[1] The values $B_{1,PE} = 0.375$ and $B_{2,PE} = 0.2274$ follow directly from the Pilch & Erdman (1987) velocity-history correlation Eq. (20), $V_d/(V\varepsilon^{1/2}) = \frac{3}{4} C_d\, \bar{t}_b + 3B\, \bar{t}_b^{\,2}$, evaluated with the incompressible-flow constants $C_d = 0.5$ and $B = 0.0758$ of their Eq. (23): this gives $\frac{3}{4} C_d = 0.375$ and $3B = 0.2274$. The same evaluation is adopted in the reference OpenFOAM `PilchErdman` model. The constants are retained here to keep the child-fragment admission criterion consistent with the standard implementation, which only requires an order-of-magnitude estimate of the residual relative velocity at $t = t_b$.

$12^+$, the fragment cloud has fully decelerated within $t_b$ and emerges at near-zero relative velocity. In this limit the maximum stable diameter must tend to infinity, meaning that all fragments are stable. Without the inner clamp, squaring a negative residual velocity factor would produce a spuriously small $d_{\text{stable}}$ and could trigger near-threshold fragmentation of already-stable children. The outer $\varepsilon$ safeguard prevents division by zero in finite-precision arithmetic. The viscous correction is retained in the PE admission criterion, whereas the maximum stable diameter follows the original inviscid PE stability threshold. A child parcel is admitted to a standard PE cycle only if

$$We > We_{\text{crit}} \qquad \text{and} \qquad d_p > d_{\text{stable}}. \tag{2.44}$$

This condition prevents already-stable fragments from being repeatedly re-latched into artificial PE breakup cycles.

The Pilch–Erdman characteristic time is

$$t^* = \frac{D_{p,0}}{U_{\text{rel},0}} \sqrt{\frac{\rho_l}{\rho_g}}, \tag{2.45}$$

and the deformation time is

$$t_{\text{def}} = 1.6\, t^*. \tag{2.46}$$

The total breakup time is

$$t_b = \left(\frac{t_b}{t^*}\right) t^*, \tag{2.47}$$

where the dimensionless ratio follows the piecewise Pilch–Erdman correlation

$$\frac{t_b}{t^*} = \begin{cases} 6.0\,(We_0 - 12)^{-1/4}, & 12 < We_0 \le 18, \\ 2.45\,(We_0 - 12)^{1/4}, & 18 < We_0 \le 45, \\ 14.1\,(We_0 - 12)^{-1/4}, & 45 < We_0 \le 351, \\ 0.766\,(We_0 - 12)^{1/4}, & 351 < We_0 \le 2670, \\ 5.5, & We_0 > 2670. \end{cases} \tag{2.48}$$

[2] The argument $(We_0 - 12)$ of the piecewise $t_b/t^*$ correlations is retained even when viscous effects modify the admission threshold. The original Pilch–Erdman fits are calibrated against experimental data using the inviscid critical Weber number 12, and replacing 12 by the viscosity-corrected $We_{\text{crit}}$ of Equation (2.39) would shift the argument outside the calibration domain. The viscous correction is therefore applied only to the admission criterion $We_0 > We_{\text{crit}}$, while the breakup-time fits are evaluated on $(We_0 - 12)$.

### 2.6.2 Deformation, Child-Size Distribution and Fragmentation

During PE deformation, the parent fragment is represented as an increasingly deformed body. The reference deformation diameter and post-breakup SMD follow the

[2] The piecewise correlations are continuous at the interval boundaries (e.g., $6 \cdot 6^{-1/4} = 2.45 \cdot 6^{1/4} \approx 3.83$ at $We_0 = 18$), so the choice between strict and inclusive inequalities at the lower bound has negligible numerical impact. The convention adopted here matches the OpenFOAM `PilchErdman` implementation.

Madabhushi–Lambert implementation [6, 8]. The reference deformation diameter is

$$d_{\mathrm{ref}} = \begin{cases} D_{p,0}\left(1 + 0.19\sqrt{We_0}\right), & We_0 < 100, \\ 2.9\, D_{p,0}, & We_0 \geq 100. \end{cases} \tag{2.49}$$

This relation is used as the deformation reference scale for the drag and child-size treatment.

The viscosity-corrected Weber number based on this deformed reference scale is

$$Oh_{\mathrm{ref}} = \frac{\mu_l}{\sqrt{\rho_l\, d_{\mathrm{ref}}\, \sigma}}, \tag{2.50}$$

$$We_{\mathrm{corr}} = \frac{We_0}{1 + 1.077\, Oh_{\mathrm{ref}}^{1.6}}. \tag{2.51}$$

The target Sauter mean diameter after breakup is then

$$d_{\mathrm{SMD}} = 1.5\, \frac{Oh_{\mathrm{ref}}^{0.2}}{We_{\mathrm{corr}}^{0.25}} D_{p,0}. \tag{2.52}$$

Thus, stronger aerodynamic loading decreases the target child size, whereas larger liquid viscosity increases it.

When the elapsed PE time reaches $t_b$, the parent parcel fragments. In the present implementation, the breakup event produces $n_f$ fragments, with $n_f = 5$ used as the default number of child parcels. Following Madabhushi [6] and Lambert et al. [8], fragment diameters are drawn from a root-normal distribution whose cumulative distribution function is

$$F(D) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{\sqrt{D} - \sqrt{\bar{d}_{\mathrm{SMD}}}}{0.238\sqrt{\bar{d}_{\mathrm{SMD}}}}\right)\right], \tag{2.53}$$

where $\bar{d}_{\mathrm{SMD}}$ is the target Sauter mean diameter of Equation (2.52). Sampling with a standard normal variate $\xi \sim \mathcal{N}(0,1)$ gives the nominal fragment diameter:

$$\bar{d}_{\mathrm{child}} = 1.2\, d_{\mathrm{SMD}}\, (1 + 0.238\, \xi)^2, \qquad \xi \sim \mathcal{N}(0,1). \tag{2.54}$$

When the event corresponds to the first PE breakup of the core parcel, the Lambert factor of Equation (2.36) is applied to the sampled child diameters.

Mass conservation is imposed by assigning a common statistical weight to the new fragments:

$$N_{p,c} = \frac{m_{\mathrm{break}}}{(\rho_l \pi/6) \sum_{i=1}^{n_f} d_i^3}, \tag{2.55}$$

so that

$$m_i = N_{p,c}\, \rho_l\, \frac{\pi}{6} d_i^3, \qquad \sum_{i=1}^{n_f} m_i = m_{\mathrm{break}}. \tag{2.56}$$

The normal rim-expansion velocity assigned to each child is

$$u_n = \frac{5\, D_{p,0}}{t_b - t_{\mathrm{def}}}, \tag{2.57}$$

with the direction sampled in the plane normal to the parent velocity $\vec{u}_p$ (Figure 2.3), consistently with the formulation of Madabhushi [6] and Lambert et al. [8]. In the implementation, no physical cap is applied to this value; only a small numerical denominator safeguard is retained to prevent division by zero.

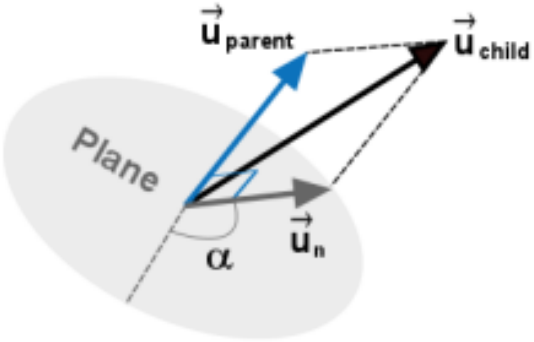



**Figure 2.3:** *Normal rim-expansion velocity geometry for PE child-parcel velocity assignment, adapted from [8].*

## 2.7 Deformation-Dependent Drag Model

The drag force is the main coupling mechanism controlling the trajectory of Lagrangian parcels. During the intact-column stage, drag determines the bending and penetration of the blob representation of the liquid jet. After column breakup, it controls the deceleration and dispersion of ligaments, fragments and child droplets. For this reason, the drag law is made dependent on the breakup state of the parcel.

### 2.7.1 State-Dependent Drag Laws

For a spherical parcel, the aerodynamic drag acceleration can be written as

$$\frac{\mathrm{d}\mathbf{u}_p}{\mathrm{d}t} = \frac{3}{4} C_D \frac{\rho_g U_{\mathrm{rel}}}{\rho_l d_p} \mathbf{U}_{\mathrm{rel}}. \tag{2.58}$$

In OpenFOAM-type Lagrangian solvers, the force model is commonly returned in semi-implicit coefficient form:

$$S_p = m_p \frac{3}{4} C_D \frac{Re_p \mu_g}{\rho_l d_p^2}, \tag{2.59}$$

where

$$Re_p = \frac{\rho_g U_{\mathrm{rel}} d_p}{\mu_g}. \tag{2.60}$$

For PE-active parcels, the same semi-implicit structure is retained, but the coefficient is multiplied by an area correction factor accounting for droplet deformation.

During the liquid-column stage, the core parcel uses the constant Madabhushi blob-drag coefficient

$$C_{D,\mathrm{blob}} = 1.48. \tag{2.61}$$

The value 1.48 is the Madabhushi/Lambert LJICF reference for Jet-A1 fuel. For the Zhang LJIGF/GLOP water cases, the selection of $C_{D,\mathrm{blob}}$ was guided by Wu et al. [2], who analysed the column trajectory and drag of water jets in crossflow and reported an averaged column drag coefficient of 1.696, of the same order of magnitude as the drag of a solid circular cylinder. For the Zhang configuration, a value of $C_{D,\mathrm{blob}} = 2.16$ was adopted, exceeding this reference value, to account for the larger projected frontal area characteristic of the confined gas-film geometry in the LJIGF/GLOP configuration, which increases the effective aerodynamic loading on the column relative to an unconfined crossflow. The full rationale and its effect on the simulated spray are discussed in Chapter 5. The coefficient governs the aerodynamic loading on the non-spherical liquid column while it is still represented as a blob parcel.

Wave-shed child parcels and post-column fragments that have not entered a PE deformation stage use the standard OpenFOAM sphere-drag correlation:

$$C_{D,\mathrm{sph}}(Re_p) = \begin{cases} \dfrac{24}{Re_p}\left(1 + \dfrac{1}{6}Re_p^{2/3}\right), & Re_p \le 1000, \\ 0.424, & Re_p > 1000. \end{cases} \tag{2.62}$$

This is the standard OpenFOAM sphere-drag law [5], corresponding to the Putnam-type correlation [13] used in OpenFOAM rather than to the classical Schiller–Naumann form [14].

Once a parcel is PE-active, its aerodynamic response is modified to account for deformation. During the interval from PE onset to $t_{\mathrm{def}}$, the drag coefficient is ramped from the spherical value to the disc value:

$$C_{D,\mathrm{eff}} = (1 - f)\, C_{D,\mathrm{sph}} + f\, C_{D,\mathrm{disc}}, \qquad f = \frac{t - t_{\mathrm{sec}}}{t_{\mathrm{def}}}, \qquad 0 \le f \le 1, \tag{2.63}$$

where $t_{\mathrm{sec}}$ is the PE onset time. This linear interpolation is an implementation-level simplification of the more general deformation-dependent drag treatment of Liu, Mather and Reitz [15]. Instead of solving an additional deformation oscillator, the deformation parameter is approximated as increasing linearly over the interval $0 \le t - t_{\mathrm{sec}} \le t_{\mathrm{def}}$, keeping the drag response synchronised with the Pilch–Erdman deformation time. After the deformation time, the drag coefficient is held at

$$C_{D,\mathrm{disc}} = 1.2. \tag{2.64}$$

The drag reference diameter is ramped consistently during the same interval:

$$d_{\mathrm{ref,drag}} = \begin{cases} D_{p,0}\left(1 + 0.19\sqrt{We_0}\, f\right), & We_0 < 100, \\ D_{p,0}\,(1 + 1.9 f), & We_0 \ge 100. \end{cases} \tag{2.65}$$

At $f = 1$, this expression recovers the fully deformed reference diameter of Equation (2.49). The increased projected area of the deformed droplet is represented through the area correction factor

$$\chi = \left(\frac{d_{\mathrm{ref,drag}}}{d_p}\right)^2. \tag{2.66}$$

In semi-implicit form, the PE-active drag coefficient returned to the parcel equation of motion is

$$S_p^{\mathrm{def}} = m_p \frac{3}{4} C_{D,\mathrm{eff}} \chi \frac{Re_p \mu_g}{\rho_l d_p^2}, \tag{2.67}$$

with $C_{D,\mathrm{eff}}$ from Equation (2.63) and $\chi$ from Equation (2.66). At the deformation-time endpoint ($f = 1$), $C_{D,\mathrm{eff}}$ recovers $C_{D,\mathrm{disc}} = 1.2$ and $\chi$ reaches its fully-deformed value, consistent with Lambert et al. [8].

## 2.8 Source-Term Coupling and Diffusion-Based Coarse Graining

Two-way coupling requires that the momentum exchanged between each parcel and the carrier gas be introduced as a source term in the Eulerian momentum equation Equation (2.8). As anticipated in Section 2.2, the standard Particle Centroid Method (PCM) deposits the full contribution of a Lagrangian parcel into the Eulerian cell containing the parcel centroid.

In the standard PCM projection, these source fields remain concentrated in the host cells of the parcels. This is locally conservative, but it makes the corresponding source density depend strongly on the local cell volume. In near-injector regions, where the mesh is refined to resolve the wall boundary layer and the shear layer around the jet, the carrier cells can be much smaller than the effective volume associated with the injected blob parcels. Depositing the entire parcel–gas exchange into one or a few such cells can therefore produce excessively concentrated source terms, leading to non-physical local perturbations of the carrier velocity, pressure and density. In the present solver, this issue is addressed by smoothing the accumulated Lagrangian source fields before they enter the Eulerian momentum solution.

### 2.8.1 Diffusion-Based Smoothing Formulation

The diffusion-based coarse-graining method (DBM) proposed by Sun and Xiao [10] provides a convenient way to replace a cell-centred PCM source field by a smooth, conservative counterpart on the same Eulerian mesh. For a generic scalar or vector source field $\phi_{\mathrm{PCM}}$, the smoothed field $\phi_s$ is obtained by solving

$$\phi_s - \sigma_s^2 \nabla^2 \phi_s = \phi_{\mathrm{PCM}}, \tag{2.68}$$

with zero-gradient boundary conditions,

$$\nabla \phi_s \cdot \hat{\mathbf{n}} = 0 \qquad \text{on } \partial\Omega. \tag{2.69}$$

The parameter $\sigma_s^2$ has dimensions of area and controls the spatial extent of the smoothing. The zero-gradient condition is used so that the diffusive redistribution does not create or remove the total integrated source over the computational domain.

Integrating Equation (2.68) over $\Omega$ and applying the divergence theorem gives

$$\int_\Omega \phi_s \, \mathrm{d}V = \int_\Omega \phi_{\mathrm{PCM}} \, \mathrm{d}V, \tag{2.70}$$

provided that the boundary condition Equation (2.69) is satisfied. DBM therefore acts as a conservative spatial redistribution of the source term, rather than as a modification of the total momentum exchanged between the Eulerian and Lagrangian phases. With different boundary conditions, such as a Dirichlet treatment on an inlet or outlet, this global conservation property would no longer follow directly from the same integral argument; this is why zero-gradient conditions are retained for the smoothed source fields.

The smoothing operation is closely related to Gaussian-kernel coarse-graining, which is the theoretical interpretation used by Sun and Xiao for the Diffusion-Based Method [10]. In their formulation, the Gaussian equivalence is obtained through a transient diffusion process, whose unbounded-domain kernel is

$$G_G(\mathbf{r}) = \frac{1}{\left(4\pi\sigma_s^2\right)^{n/2}} \exp\left(-\frac{|\mathbf{r}|^2}{4\sigma_s^2}\right), \tag{2.71}$$

where $n$ is the spatial dimension.

The present OpenFOAM implementation uses the implicit screened-diffusion solve of Equation (2.68) as a robust finite-volume smoothing operator. The exact unbounded Green's function of the operator $1 - \sigma_s^2\nabla^2$ is not Gaussian; in three dimensions it has the Yukawa form

$$G_Y(r) = \frac{\exp(-r/\sigma_s)}{4\pi\sigma_s^2 r}. \tag{2.72}$$

Nevertheless, the screened-diffusion solve is used here as a Gaussian-equivalent coarse-graining operation in the numerical sense: it spreads the initially cell-centred PCM source over a finite neighbourhood while preserving the integrated source through the zero-gradient boundary condition. The sensitivity to the smoothing bandwidth is therefore assessed explicitly in the Lambert and Zhang validation chapters, where the strength of the source-term concentration artefact and the role of the DBM regularisation are different.

### 2.8.2 Bandwidth Selection and Model Role

The smoothing intensity is controlled by a user-prescribed bandwidth $b$. In the present formulation, the diffusion parameter is defined as

$$\sigma_s^2 = \frac{b^2}{4}. \tag{2.73}$$

With this choice, the Gaussian form in Equation (2.71) can be written as

$$G_G(\mathbf{r}) \propto \exp\left(-\frac{r^2}{b^2}\right),$$

so $b$ acts as a practical e-folding bandwidth for the smoothed source. The corresponding one-dimensional standard deviation of this Gaussian kernel is $b/\sqrt{2}$. In the present work, $b$ is therefore treated as a numerical smoothing bandwidth, not as a directly measurable physical spray length scale.

The bandwidth $b$ controls the compromise between source localisation and source over-smearing. If $b$ is too small, the method approaches the original PCM projection and may retain the near-injector source-concentration artefact. If $b$ is too large, the parcel–gas momentum exchange is distributed over an excessively wide region, weakening the local two-way coupling and potentially distorting the predicted jet trajectory. In practice, the effective sensitivity to $b$ is case- and mesh-dependent. In configurations where source-term concentration is not the dominant error mechanism, even small non-zero values of $b$ may already produce the regularised DBM response, leading to a weak dependence on the precise bandwidth value. Conversely, in confined configurations with strong near-injector coupling, such as the LJIGF/GLOP case considered later, the choice of $b$ becomes more critical because PCM source localisation can directly alter the gas-film structure and jet deflection.

The bandwidth values are selected as multiples or fractions of the injector diameter and are reported separately for each validation case. This avoids introducing $b$ as a universal closure constant and allows the numerical role of the smoothing operator to be evaluated in relation to the local mesh size, injector diameter and source-term concentration. The carrier-to-parcel interpolation scheme also matters, because the `cell` and `cellPoint` variants can produce different parcel velocity histories and therefore different breakup statistics even when DBM is disabled. This becomes decisive in the confined Zhang campaigns, where the `cell` variant is unviable and the `cellPoint` scheme is required, as discussed in Chapter 5.

DBM does not replace physical mesh refinement. The near-wall and shear-layer regions still require adequate spatial resolution for the carrier phase and for the turbulence model. The role of DBM is instead to remove the artificial dependence of the Lagrangian source deposition on very small host-cell volumes, thereby improving the robustness of the coupled Eulerian–Lagrangian solution.

## 2.9 Complete Model Sequence, Assumptions, and Limitations

The model described in the previous sections is a coupled state-dependent Lagrangian formulation. The active breakup and drag laws are not selected from a global regime map, but from the current parcel state, local gas properties and latched breakup variables. At each Lagrangian update, the model sequence can be summarised as follows:

1. Inject core/blob parcels at the injector face with initial diameter $D_0 = D_{\text{inj}}$, velocity directed along the injector axis with magnitude $U_l$, and initial core state.
2. Interpolate the local carrier-gas velocity $\mathbf{u}_g(\mathbf{x}_p, t)$ and density $\rho_g(\mathbf{x}_p, t)$ at the parcel position, and compute the relative velocity $\mathbf{U}_{\text{rel}}$ from Equation (2.13).
3. For intact-core parcels, evaluate KH/Wave stripping (Section 2.3); update the stripped-mass accumulator and create a wave-shed child parcel when the accumulated mass exceeds the shedding threshold.
4. Evaluate the column-breakup clock using Equation (2.35). When $t_c \geq t_{cb}$, terminate the intact-column state and latch the core parcel into the first post-column Pilch–

Erdman sequence.
5. For post-column fragments and wave-shed children, evaluate the Pilch–Erdman admission criteria. If the parcel is unstable, latch $D_{p,0}$, $U_{\mathrm{rel},0}$, $We_0$, and $Oh_0$.
6. Advance the PE deformation process. When the elapsed PE time reaches $t_b$, fragment the parent into $n_f = 5$ children, with diameters sampled from Equation (2.54) and rim-expansion velocity from Equation (2.57).
7. Apply the Lambert ligament correction of Equation (2.36) only if the event is the first PE breakup of the original core parcel.
8. Assign the drag response according to the parcel state: blob drag for intact-core parcels, spherical drag for non-deforming children, and deformation/disc drag for PE-active parcels.
9. Integrate the parcel kinematic and momentum equations *Equations* (2.11)*and* (2.12), and accumulate the corresponding Eulerian momentum source terms.
10. Apply DBM smoothing to the accumulated source fields before the Eulerian momentum equation is solved.

### 2.9.1 Original Contribution

The individual components used in the formulation are established in the literature: the Reitz Wave model, the Madabhushi LJICF column-breakup framework, the Pilch–Erdman secondary-breakup correlations, the Lambert ligament correction and the diffusion-based coarse-graining method. The contribution of this thesis is their integration into an open-source OpenFOAM implementation with a consistent state machine, state-dependent drag treatment, mass-conservative child creation and DBM smoothing of the Lagrangian momentum source terms.

This is relevant because the Lambert enhancement was originally assessed in a commercial-solver context, whereas the present work provides an extensible OpenFOAM implementation and tests its behaviour not only on Lambert LJICF reference cases, but also on the confined Zhang LJIGF/GLOP configuration. The latter assessment is used to examine the transferability of the Madabhushi–Lambert model outside the classical unbounded-crossflow setting.

### 2.9.2 Assumptions and Limitations

The model rests on several simplifying assumptions that define its range of applicability:

- **Unresolved liquid core.** The intact liquid column is not resolved as a continuous gas–liquid interface. It is represented through blob parcels whose initial diameter is set equal to the injector diameter.
- **Semi-empirical breakup closures.** The KH/Wave correlations, the column-breakup constant $C_0$, the Pilch–Erdman time and size correlations, and the Lambert factor $F_{\mathrm{Lig}}$ are empirical or semi-empirical. Their transferability outside their calibration range must therefore be assessed case by case.
- **Case-dependent KH child-size coefficient.** The coefficient $B_0$ controlling the KH/Wave child size is retained as an adjustable parameter. The Lambert LJICF cases use values

consistent with the Lambert reference setup, while the Zhang LJIGF/GLOP simulations require a separate calibration because the confined gas-film configuration changes the near-field stripping behaviour.

- **Local carrier-state dependence.** The KH/Wave stripping rate, column-breakup clock, PE onset criterion and drag law use gas properties interpolated from the Eulerian solution at the parcel position. The model is therefore responsive to local flow acceleration, recirculation and two-way coupling, but also sensitive to mesh resolution, turbulence closure and source-term mapping.

- **Latched PE evolution.** Once a Pilch–Erdman event is admitted, the onset quantities $D_{p,0}$, $U_{\mathrm{rel},0}$, $We_0$, and $Oh_0$ are frozen for the duration of that event. This is consistent with the calibration logic of the PE correlations, but it remains an idealisation of a continuously evolving turbulent crossflow.

- **LJICF–LJIGF transferability.** The Madabhushi–Lambert model was developed for LJICF. Its application to the finite-height gas-film geometry of the Zhang GLOP case is not guaranteed a priori and is assessed in Chapter 5.

- **DBM bandwidth and interpolation sensitivity.** The DBM bandwidth $b$ is a numerical smoothing parameter. It reduces source-concentration artefacts, but excessive smoothing can distort the local parcel–gas coupling. Its influence is quantified separately in the Lambert and Zhang validation chapters. The carrier-to-parcel interpolation scheme is a further sensitivity, because the `cell` and `cellPoint` schemes can modify the local gas velocity sampled by the parcels and therefore affect both velocity statistics and breakup history; this becomes decisive in the confined Zhang campaigns (Chapter 5), where the `cell` variant is unviable.

- **Dilute-spray hypothesis.** Parcel–parcel collisions and finite-volume displacement effects of the dispersed phase are neglected. This assumption is consistent with an Eulerian–Lagrangian dilute-spray formulation and must be checked a posteriori through the predicted liquid loading.

- **No resolved interface topology.** The model cannot provide the instantaneous liquid-column shape, ligament topology, cavity formation or detailed near-nozzle interface dynamics that would require an interface-resolved VOF or VOF–LPT approach.

# 3 OpenFOAM Implementation

This chapter describes the OpenFOAM implementation of the Madabhushi–Lambert Lagrangian breakup and drag model formulated in Chapter 2. The aim is not to reproduce the complete source code, but to document the implementation at a level that makes the solver architecture, the parcel-state logic and the numerical coupling procedure reproducible and auditable. Particular attention is given to the modified libraries, the reuse of parcel-state variables, the multi-child creation mechanism, the state-dependent drag selection, the diffusion-based smoothing of the Lagrangian source terms and their placement within the solver time loop.

The implementation is designed to preserve the standard OpenFOAM Lagrangian spray infrastructure as much as possible. The additional behaviour required by the present model is introduced only where it is strictly needed: persistent breakup-state tracking inside the parcel data, multi-fragment generation during Pilch–Erdman breakup, synchronisation between breakup and drag regimes, and diffusion-based coarse-graining of the Eulerian source terms. This strategy avoids rewriting the full parcel-tracking framework and keeps the solver compatible with the standard OpenFOAM cloud-evolution sequence, while still allowing the LJICF-specific physics described in Chapter 2 to be activated through runtime-selectable models. The complete solver is publicly available as open-source software [16].

## 3.1 Solver and Library Architecture

The implementation starts from the standard OpenFOAM spray-solver architecture, in which the carrier phase is advanced on an Eulerian finite-volume mesh and the liquid phase is represented by a Lagrangian spray cloud. The custom solver used in this work, `sprayFoamDBM`, follows the structure of `sprayFoam`, but inserts an additional diffusion-based smoothing step after the Lagrangian cloud evolution and before the Eulerian PIMPLE solution. In this way, the cloud first tracks parcels, evaluates breakup and drag, creates children and accumulates the raw `UTrans` and `UCoeff` source fields; these source fields are then regularised before they are used by the carrier-phase momentum equation.

Development and debugging were carried out mainly with OpenFOAM v2406. The same modelling logic can be compiled against other OpenFOAM versions, provided that the modified parcel class, breakup library, drag-force library and solver are built consistently against the target installation. This point is important because minor API differences between OpenFOAM versions can affect parcel creation, clone construction,

runtime selection and library linking. The implementation therefore keeps the standard OpenFOAM runtime-selection mechanism for both breakup models and particle forces, while adding only the model-specific state handling required by the Madabhushi–Lambert formulation.

The main added or modified components are summarised in Table 3.1. The table is organised by responsibility rather than by file-system location, because the same source files can be arranged differently depending on whether the model is compiled as a local case library or as a reusable OpenFOAM user library.

| Component | Role in the implementation |
|---|---|
| `BreakupModel.H` | Base-class header extended with virtual methods for multi-child signalling (`pendingChildren`, `childMsInit`, `childUserInit`) required by the multi-fragment PE breakup. |
| `Madabhushi.C/H` | Runtime-selectable breakup model. It implements KH/Wave stripping, column-breakup latching, Pilch–Erdman evolution, the Lambert ligament correction, child signalling and pending-child storage. |
| `SprayParcel.C` | Parcel-level child creation, compiled in the custom `liblagrangianSpray.so`. It extends the standard clone-based creation mechanism to assign breakup flags, reset child state variables and consume the pending-child buffer. |
| `MadabhushiDragForce.C/H` | Runtime-selectable particle force. It selects blob, spherical or deformation/disc drag according to the persistent parcel state. |
| `sprayFoamDBM.C` | Solver-level extension of `sprayFoam`. It calls the DBM smoothing routine after cloud evolution and before the Eulerian PIMPLE solution. |
| `smoothSourceTerms.H` | Finite-volume implementation of the screened-diffusion smoothing applied to the accumulated `UTrans` and `UCoeff` source fields. |
| `createDBMFields.H` | Creation and initialisation of DBM work fields and dictionary parameters, including the smoothing bandwidth. |

**Table 3.1:** *Main modified or added source components.*

At class level, the breakup and drag extensions follow the standard OpenFOAM inheritance pattern:

$$\begin{aligned} \texttt{BreakupModel<CloudType>} &\longleftarrow \texttt{Madabhushi<CloudType>}, \\ \texttt{ParticleForce<CloudType>} &\longleftarrow \texttt{MadabhushiDragForce<CloudType>}. \end{aligned} \tag{3.1}$$

The breakup class determines the evolution of the parcel diameter, represented mass and breakup state. The parcel class is responsible for instantiating new OpenFOAM parcel objects when the breakup model requests child creation. The drag class reads the same persistent parcel state and returns the appropriate semi-implicit force coefficient. Finally, the solver-level DBM routine smooths the accumulated Lagrangian source terms

before the Eulerian momentum equation is solved. This separation of responsibilities keeps the breakup physics, child creation, drag selection and Eulerian source-term treatment modular while ensuring that they remain synchronised through the same parcel-state variables.

## 3.2 Injection Model and Initial Blob Parcels

No custom injection class is required for the Madabhushi–Lambert model. The liquid jet is introduced through the standard OpenFOAM spray-cloud injection dictionary, while the custom breakup model interprets the injected parcels as liquid-column blobs. The essential modelling requirement is that the initial parcel diameter is prescribed equal to the injector orifice diameter,

$$D_0 = D_{\mathrm{inj}}, \tag{3.2}$$

so that each injected parcel represents a Lagrangian surrogate of a finite portion of the unresolved liquid core, as discussed in Section 2.3.

For a prescribed liquid mass-flow rate $\dot{m}_l$, parcel injection rate $\dot{N}_{\mathrm{par}}$, and initial blob diameter $D_0$, the represented mass assigned to one computational parcel is

$$m_{p,0} = \frac{\dot{m}_l}{\dot{N}_{\mathrm{par}}}, \tag{3.3}$$

and the corresponding statistical parcel weight is

$$N_{p,0} = \frac{m_{p,0}}{\rho_l(\pi/6)D_0^3}. \tag{3.4}$$

In OpenFOAM this statistical weight corresponds to `nParticle`. The injected velocity is assigned by the selected injection model and is aligned with the injector axis. For the Lambert LJICF cases this direction is transverse to the crossflow, whereas for the Zhang LJIGF/GLOP configuration it is defined consistently with the gas-film geometry. Therefore, the injection patch, liquid mass-flow rate, injection duration, parcel injection rate and initial parcel diameter are case-setup quantities; the custom breakup library only assumes that the injected parcels enter the model as blob parcels with $D_0 = D_{\mathrm{inj}}$.

The custom breakup model starts from this injected state. Newly injected core parcels enter the state machine with `user` = 0, no active Pilch–Erdman latch, and `ms` available as the KH/Wave stripped-mass accumulator. Subsequent changes to the parcel state are controlled by the breakup model and by the child-creation logic described in the following sections. In this way, the standard OpenFOAM injection infrastructure is kept separate from the LJICF-specific breakup, drag and state-management logic.

## 3.3 Parcel State Extension and Field Reuse

The Madabhushi–Lambert model requires persistent information that must remain attached to each parcel between Lagrangian time steps. This information includes the

parcel breakup clock, the onset time of a Pilch–Erdman event, the parent diameter and relative velocity latched at PE onset, the KH/Wave stripped-mass reservoir, and the regime flags required to distinguish core parcels from wave-shed and post-column fragments. Rather than extending the fundamental OpenFOAM parcel header with new data members, the implementation reuses scalar fields already available in the spray parcel class. This preserves compatibility with the existing cloud infrastructure, parcel I/O, clone construction and runtime-selection mechanisms.

The consequence is that some OpenFOAM identifiers are reused with a model-specific meaning. This convention is documented in Table 3.2. The labels are inherited from the OpenFOAM parcel API, but their default TAB/KH interpretation is not used by the active Madabhushi–Lambert model. In particular, the TAB oscillation equation is not solved in the present formulation; therefore `y`, `yDot` and `KHindex` are used as persistent storage slots for the PE state variables.

| **Code field** | **Pre-existing field / default OpenFOAM use** | **Use in this thesis** |
|---|---|---|
| `tc` | Breakup clock or parcel breakup-time variable | Parcel lifetime and breakup clock $t_c$, used by the KH/Wave, column-breakup and Pilch–Erdman timing logic. |
| `y` | TAB deformation coordinate in TAB-type models | PE onset time $t_{\mathrm{sec}}$. A value $y \leq 0$ indicates that no PE event is currently active. |
| `yDot` | TAB deformation-rate variable in TAB-type models | Dual role. Before PE onset, it stores the initial wave-stage core mass used by the five-percent KH shedding threshold. After PE onset, it stores the latched parent diameter $D_{p,0}$. |
| `KHindex` | KH-related event/index storage in KH-type models | Latched relative velocity magnitude $U_{\mathrm{rel},0}$ at PE onset. It is reset for newly created children so that they do not inherit stale PE information from the parent. |
| `ms` | Breakup-model scalar storage; used here also for stripped mass | Dual use. Before column breakup it is the accumulated KH/Wave stripped mass; after a state transition it is used as a discrete state marker for PE-active and sentinel parcels. |
| `user` | Generic user scalar attached to the parcel | Persistent regime tag distinguishing the original core lineage, wave-shed children and post-column/post-PE children. It is mainly used to keep breakup and drag-regime selection synchronised. |

**Table 3.2:** *Reuse of existing parcel fields in the Madabhushi–Lambert implementation.*

The primary mathematical mapping used after a PE latch is

$$\begin{aligned} \texttt{tc} &\leftrightarrow t_c, \\ \texttt{y} &\leftrightarrow t_{\text{sec}}, \\ \texttt{yDot} &\leftrightarrow D_{p,0}, \\ \texttt{KHindex} &\leftrightarrow U_{\text{rel},0}. \end{aligned} \tag{3.5}$$

Before PE onset, `yDot` has a different role and stores the initial wave-stage mass used by the KH shedding threshold. After PE onset, `y`, `yDot` and `KHindex` define the frozen reference state used to compute the Pilch–Erdman characteristic time, deformation diameter, post-breakup size distribution and deformation-dependent drag.

The variable `ms` has the most delicate role, because it is used both as a physical stripped-mass accumulator and as a state marker. The convention is summarised in Table 3.3.

| Value | Meaning |
|---|---|
| $\texttt{ms} \geq 0$ | Accumulated stripped mass during the KH/Wave stage of an intact core parcel. |
| $\texttt{ms} = -1$ | First PE event of the original core parcel. This marker identifies the event to which the Lambert ligament correction must be applied. |
| $\texttt{ms} = 2$ | Standard PE-active parcel. The parcel has already passed the PE admission check and is evolving through a standard Pilch–Erdman deformation event. |
| $\texttt{ms} = -10$ | Wave-shed child awaiting the standard PE admission check. |
| $\texttt{ms} = -20$ | Post-column or post-PE child awaiting the standard PE admission check. If the admission test is not satisfied, the fragment remains in this sentinel path and is transported without a new PE latch. |

**Table 3.3:** *Meaning of the `ms` state variable.*

This dual use is a pragmatic implementation choice. It avoids extending the parcel header and therefore preserves compatibility with the rest of the OpenFOAM parcel pipeline. The state machine enforces mutual exclusion between the two meanings: before column breakup, `ms` is treated as a stripped-mass reservoir; after a state transition, specific negative or positive sentinel values define the breakup path. A future refactoring could replace this convention with a dedicated enumeration or explicit state field, but the present design is sufficient for the coupled breakup, drag and child-creation logic used in this work.

The variable `user` is used as a coarser persistent regime tag, as reported in Table 3.4. It is intentionally separated from `ms`: `ms` identifies the detailed breakup path, whereas `user` distinguishes the parcel family used by the drag model and by the child-signalling logic.

| Value | Meaning |
|---|---|
| user = 0 | Original core lineage. This includes the injected blob parcel and the retained core parcel during the first post-column PE sequence. The actual drag regime is then disambiguated using the column-breakup clock and the PE latch variables. |
| user = 1 | Wave-shed child generated by KH/Wave stripping. Before a PE latch, this parcel family is transported with spherical drag and is tested through the standard PE admission path. |
| user = 2 | Post-column or post-PE child generated by a Pilch–Erdman breakup event. Before any subsequent PE latch, this parcel family is transported as a non-deforming child and re-tested against the standard PE admission criteria. |

**Table 3.4:** *Meaning of the* user *parcel-regime flag.*

## 3.4 Breakup State Machine in `Madabhushi.C`

The breakup model is implemented as a state machine inside the `Madabhushi` model update. At each Lagrangian update, the model receives the current parcel diameter, breakup clock, statistical weight, material properties, carrier-gas state and relative velocity. The active path is then selected from the persistent parcel variables introduced in Section 3.3. This state-machine structure is necessary because the KH/Wave stripping model, the Madabhushi column-breakup latch, the Lambert ligament correction, the Pilch–Erdman breakup sequence and the drag-regime selection must remain mutually consistent but must not be applied simultaneously.

The implemented logic can be summarised through the following paths:

1. **Intact-core path.** Newly injected core parcels with user = 0 and no active PE latch are treated as liquid-column blobs. While $t_c < t_{cb}$, they undergo KH/Wave surface stripping. The parent diameter is relaxed toward the KH target child diameter and the removed mass is accumulated in ms.
2. **Column-breakup latch.** When $t_c \geq t_{cb}$, the intact-column representation is terminated. Any residual KH stripped mass that has not yet been emitted as a wave child is reabsorbed into the core parcel, the parent diameter is made consistent with the represented mass, and the parcel is latched into the first post-column Pilch–Erdman event with ms = −1.
3. **Sentinel-child admission.** Wave-shed children and post-PE children are not forced directly into a new PE event. They first enter a sentinel path, identified for example by ms = −10 or ms = −20, and are tested against the standard PE admission criteria. This avoids repeated fragmentation of already stable droplets.
4. **PE-active evolution.** Once a parcel is admitted to a PE event, the onset quantities $D_{p,0}$, $U_{\mathrm{rel},0}$, $t_{\mathrm{sec}}$ and the corresponding state marker are frozen in the reused parcel fields. The parcel then evolves through the deformation stage until the elapsed PE time reaches the breakup time $t_b$.
5. **Fragmentation and relatching.** When the PE breakup time is reached, the parent is fragmented into $n_f$ child parcels. The first child is returned through the standard

OpenFOAM breakup interface, whereas the additional children are stored in a pending-child buffer and later consumed by the modified parcel-creation routine. Newly generated fragments that remain unstable can be latched into a standard PE event; otherwise they remain in a sentinel child path and are transported with spherical drag until a later local condition admits a new PE cycle.

The separation between these paths is essential. KH/Wave stripping is restricted to the intact-core stage; the Lambert factor is restricted to the first post-column PE breakup of the original core parcel; and standard PE relatching is allowed only after a child fragment has passed both the critical Weber number and maximum-stable-diameter checks. In this way, the state machine prevents the different semi-empirical closures from acting on the same parcel in incompatible ways.

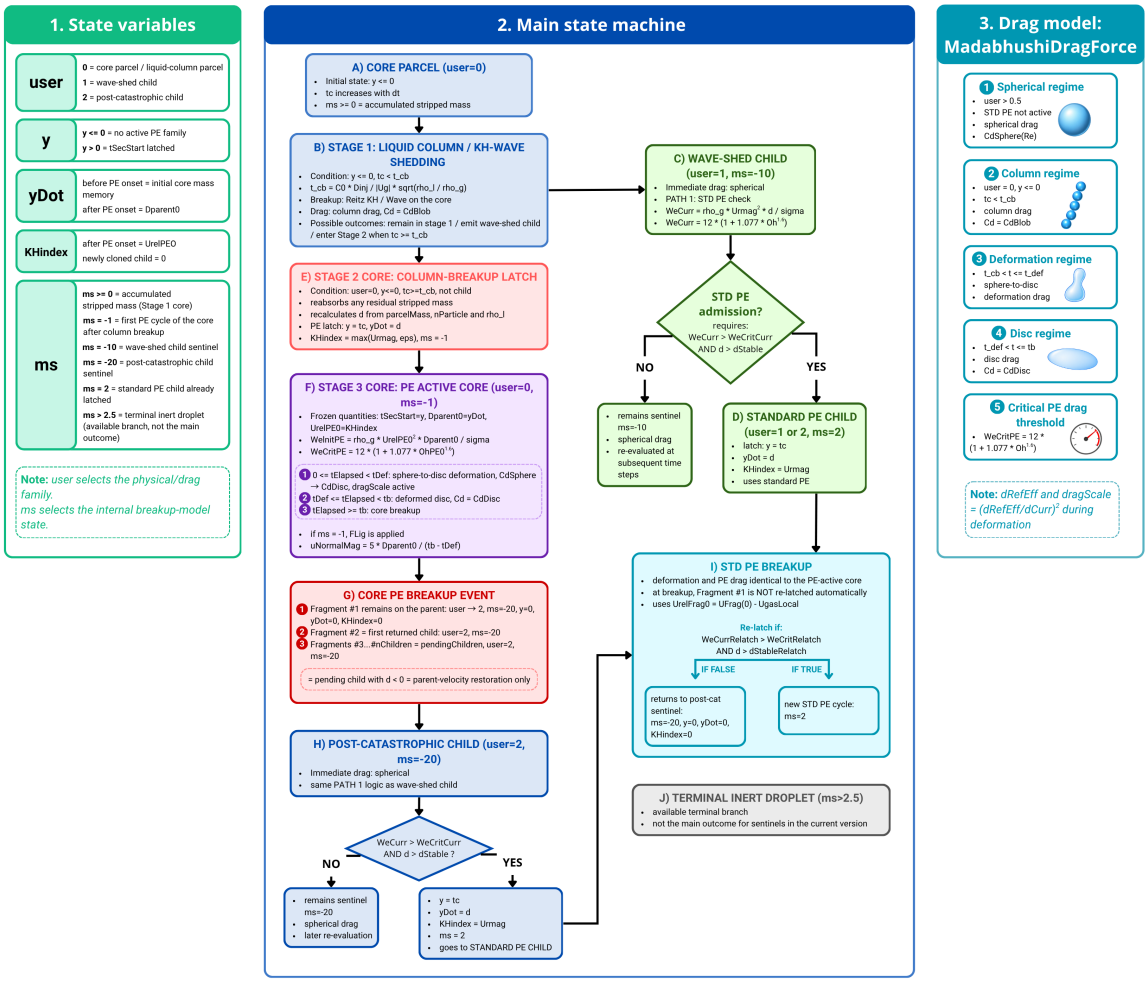



**Figure 3.1:** *State-machine overview of the coupled Madabhushi–Lambert breakup and drag implementation. The diagram follows the state convention of Table 3.3 and the regime-tag convention of Table 3.4.*

Figure 3.1 provides a code-level overview of this logic. The figure is not an additional physical model; it is a compact representation of how the correlations introduced in Chapter 2 are activated inside the parcel state machine.

**Implementation note.** The legend of Figure 3.1 includes an `ms = 3` marker labelled *stable droplet* that does not correspond to any value assigned by the implemented state

machine. Stable post-breakup fragments remain in the sentinel state `ms` $= -20$, and the guard `if (ms > 2.5) return false` in the breakup update serves as an inert terminal branch for any future extension, but is never reached in the current implementation because no execution path assigns `ms` $= 3$.

### 3.4.1 KH/Wave Stripping and Column-Breakup Latch

For intact-core parcels, the breakup model first evaluates the KH/Wave conditions using the local carrier-gas state interpolated at the parcel position. If the Wave model is active, the parent diameter is relaxed toward the target KH child diameter and the removed mass is accumulated in `ms`. When the accumulated stripped mass exceeds the shedding threshold, the breakup model requests the creation of a wave-shed child parcel and assigns the corresponding child markers through the child-signalling interface.

In the implementation, $m_{0,\mathrm{wave}}$ is stored in the `yDot` field on the first invocation of the breakup update for a core parcel (before any Pilch–Erdman onset). The same field is later repurposed to hold $D_{p,0}$ at the column-breakup latch, as documented in Table 3.2. The two roles are mutually exclusive: the latch overwrites `yDot` exactly once, after the residual stripped mass has been reabsorbed and the diameter has been made consistent with the updated parcel mass.

The column-breakup time is evaluated using the same local gas state employed in the formulation of Section 2.4:

$$t_{cb} = C_0 \frac{D_{\mathrm{inj}}}{|\mathbf{u}_g(\mathbf{x}_p, t)|} \sqrt{\frac{\rho_l}{\rho_g(\mathbf{x}_p, t)}}. \tag{3.6}$$

When the parcel clock reaches $t_{cb}$, the residual stripped mass is reabsorbed into the core parcel before the PE latch is applied. This operation is a numerical consistency repair: the stripped-mass reservoir represents physical liquid that has already been removed from the diameter relaxation, but that may not yet have reached the five-percent emission threshold required to create a separate wave-shed child. Reabsorbing this residual mass before the column-breakup latch ensures that the parent mass, diameter and statistical weight remain mutually consistent.

The first post-column PE event is then latched by storing the onset time, the parent diameter and the relative velocity in the reused parcel fields:

$$\begin{aligned} \mathtt{y} &= t_c, \\ \mathtt{yDot} &= D_{p,0}, \\ \mathtt{KHindex} &= U_{\mathrm{rel},0}, \\ \mathtt{ms} &= -1. \end{aligned} \tag{3.7}$$

No child parcel is created at the instant $t_c = t_{cb}$. Column breakup is therefore implemented as a state transition from the intact-column representation to a latched PE deformation event. Actual fragmentation occurs later, when the elapsed PE time reaches $t_b$.

### 3.4.2 PE Admission, Fragmentation and Lambert Flag

For wave-shed children, post-column fragments and post-PE fragments, the model evaluates the standard PE admission criteria before starting any new PE cycle. The current local Weber number is compared with the viscosity-corrected critical Weber number, and the current diameter is compared with the maximum stable diameter. Only fragments that are both aerodynamically unstable and larger than the stable limit are latched into a new PE event. This admission step is important because it prevents the state machine from repeatedly fragmenting droplets that should instead be transported as stable or marginally stable children.

If the parcel is admitted, the onset quantities are stored in the same fields used by the first post-column PE latch. During the PE-active interval, the latched values remain fixed and are used consistently by both the breakup model and the deformation-dependent drag force. When the breakup time $t_b$ is reached, the parent parcel is fragmented into $n_f$ children. The child diameters, represented masses and velocities are computed by the breakup model according to the Pilch–Erdman and Lambert relations described in Section 2.6 and Section 2.5.

The Lambert ligament correction is applied only when the current PE event is the first PE breakup of the original core parcel, identified by the marker `ms` $= -1$. After this event, the generated children follow the standard PE admission logic and do not retain the Lambert correction. This restriction is necessary because the Lambert factor represents the ligament-like nature of the first post-column fragments, not a general correction for all subsequent secondary-breakup events.

## 3.5 Breakup-Model Child Signalling

The standard OpenFOAM breakup interface is naturally oriented toward returning a single child parcel through the usual `addParcel` mechanism. The Madabhushi–Lambert PE breakup event, however, can generate multiple fragments. To support this without rewriting the OpenFOAM parcel-tracking infrastructure, the breakup model communicates child information to `SprayParcel.C` through an extended signalling interface.

The first child is passed through the standard breakup return values. In addition to its diameter and represented mass, the breakup model provides two state markers:

$$\begin{aligned} \texttt{childMsInit()} &\rightarrow \text{initial } \texttt{ms} \text{ of the first child,} \\ \texttt{childUserInit()} &\rightarrow \text{initial } \texttt{user} \text{ of the first child.} \end{aligned} \tag{3.8}$$

These markers ensure that the child does not inherit the parent breakup state after the clone operation.

Additional children are stored in a pending-child buffer inside the breakup model. Each pending child is represented by

$$\mathcal{C}_i = (d_i, m_i, \mathbf{u}_i), \qquad i = 2, \ldots, n_f, \tag{3.9}$$

where $d_i$ is the target child diameter, $m_i$ is the represented child mass and $\mathbf{u}_i$ is the assigned child velocity. The buffer is local to the breakup-model instance on the current

processor and is cleared after the pending children have been consumed by the parcel-creation routine.

This design separates breakup physics from parcel instantiation. The `Madabhushi` breakup model computes the child properties and state markers, while `SprayParcel.C` creates the actual OpenFOAM parcel objects, assigns their statistical weights and inserts them into the cloud. In parallel decomposed runs, the newly created children are initially placed in the same local cell as the parent; subsequent processor-boundary crossing is handled by the standard OpenFOAM parcel-tracking and migration machinery.

## 3.6 Child Creation and Mass Conservation in `SprayParcel.C`

The modified `SprayParcel.C` retains the standard OpenFOAM clone-based parcel creation strategy. A child parcel is first created as a copy of the parent, and its diameter, represented mass, velocity and state variables are then overwritten with the values computed by the breakup model. This approach is minimally invasive, but it requires careful resetting of all fields that would otherwise carry stale PE or KH information from the parent copy.

For the first child of a breakup event, the state assignment is

$$\begin{aligned} \texttt{child->ms()} &= \texttt{childMsInit()}, \\ \texttt{child->user()} &= \texttt{childUserInit()}. \end{aligned} \tag{3.10}$$

The PE-related state variables are then reset:

$$\begin{aligned} \texttt{child->tc()} &= 0, \\ \texttt{child->KHindex()} &= 0, \\ \texttt{child->y()} &= \texttt{y0()}, \\ \texttt{child->yDot()} &= \texttt{yDot0()}. \end{aligned} \tag{3.11}$$

The reset ensures that the newly created parcel enters the state machine as an independent fragment. It will be admitted to a standard PE event only if the local stability conditions are satisfied after creation.

Extra children are created from the pending-child buffer. Their state assignment is

$$\begin{aligned} \texttt{extraChild->ms()} &= -20, \\ \texttt{extraChild->user()} &= \texttt{pendingChildUserFlag()}. \end{aligned} \tag{3.12}$$

The marker $\texttt{ms} = -20$ identifies a post-PE child that must be tested by the standard PE-admission path before any further breakup. Its PE latch variables are also reset so that no active PE stage is inherited from the parent.

A key implementation correction is that the retained parent parcel is not reclassified before the child is cloned. Reclassifying the parent too early would cause the retained core parcel to enter a child branch, making the drag model select spherical drag instead

of blob drag while the parcel still belongs to the original core lineage. This was found to suppress the correct streamwise acceleration and crossflow deflection of the liquid column. In the final implementation, child-specific flags are assigned only to newly created child parcels, while any optional update of the retained parent is requested explicitly by the breakup model after the breakup decision has been completed.

Mass conservation is enforced through the parcel definition

$$m_p = N_p \rho_l \frac{\pi}{6} d_p^3, \tag{3.13}$$

where $N_p$ is the OpenFOAM statistical parcel weight `nParticle`. If a child with diameter $d_c$ and represented mass $m_c$ is created, its statistical weight is

$$N_{p,c} = \frac{m_c}{\rho_l (\pi/6) d_c^3}. \tag{3.14}$$

Numerical safeguards are applied in the code to avoid division by vanishing diameters or pathological parcel weights.

Before the original core parcel is latched into the first PE stage, residual stripped mass is reabsorbed and the core diameter is reconstructed from the updated represented mass:

$$d_p = \left( \frac{m_p}{N_p \rho_l \pi/6} \right)^{1/3}. \tag{3.15}$$

This operation is a numerical consistency repair of the parcel triple $(m_p, d_p, N_p)$, not an additional physical breakup model. Without it, the subsequent PE timescale, Weber number and child-size distribution would be computed from a diameter inconsistent with the liquid mass still represented by the core parcel.

## 3.7 State-Dependent Drag Force

The drag law is implemented in the runtime-selectable `MadabhushiDragForce` class, derived from `ParticleForce<CloudType>`. The full drag contribution is returned by `calcCoupled`, while `calcNonCoupled` returns zero. This prevents double counting and ensures that blob, spherical and deformation/disc drag are all selected inside a single state-dependent coupled-force branch.

At every force evaluation, the drag model reads the persistent parcel state:

$$\begin{aligned} \texttt{userFlag} &\leftarrow \texttt{p.user()}, \\ \texttt{msState} &\leftarrow \texttt{p.ms()}, \\ \texttt{tSecStart} &\leftarrow \texttt{p.y()}, \\ \texttt{Dparent0} &\leftarrow \texttt{p.yDot()}, \\ \texttt{UrelPE0} &\leftarrow \texttt{p.KHindex()}. \end{aligned} \tag{3.16}$$

These quantities are combined with the parcel age and the locally evaluated column-breakup time to determine whether the parcel should receive blob drag, spherical child drag or PE deformation/disc drag.

For PE-active deformed parcels, the semi-implicit coefficient returned to OpenFOAM is

$$S_p^{\text{def}} = m_p \frac{3}{4} C_{D,\text{eff}} \chi \frac{Re_p \mu_g}{\rho_l d_p^2}, \tag{3.17}$$

where $C_{D,\text{eff}}$ follows the deformation ramp defined in Equation (2.63) and $\chi$ is the projected-area correction introduced in Equation (2.66). The same local column-breakup time used by the breakup model is also evaluated by the drag force so that the intact core remains in the blob-drag regime only while the column-breakup condition has not yet been reached and no PE latch is active.

The branch selection is summarised in Table 3.5.

| **Parcel state** | **Typical condition** | **Drag law** |
|---|---|---|
| Intact core | user = 0, no PE latch, $t_c < t_{cb}$ | Blob drag with $C_{D,\text{blob}}$. |
| Wave-shed child | user = 1, no active PE latch | OpenFOAM spherical drag. |
| Post-column/post-PE child | user = 2, no active PE latch | OpenFOAM spherical drag. |
| PE-active core parcel | user = 0, active PE latch after column breakup | Deformation/disc drag with projected-area factor $\chi$. |
| PE-active child parcel | user = 1 or user = 2, standard PE latch active | Deformation/disc drag with projected-area factor $\chi$. |
| Stable or non-admitted child | Sentinel child not admitted to a new PE event | Spherical transport drag. |

**Table 3.5:** *State-dependent drag selection in* `MadabhushiDragForce`.

This synchronisation between breakup and drag is essential. If a child parcel were allowed to retain the core state, it would incorrectly receive blob drag. Conversely, if the retained core parcel were reclassified as a child too early, it would lose blob drag before the correct column-breakup or PE transition, altering the predicted penetration and crossflow deflection. The final implementation therefore lets the breakup model assign child-specific state markers explicitly, while the drag model interprets the same markers to select the appropriate force law.

## 3.8 DBM Source-Term Smoothing in `sprayFoamDBM.C`

The diffusion-based coarse-graining formulation introduced in Section 2.8 is implemented at solver level. The custom solver follows the standard `sprayFoam` time loop, but inserts the smoothing operation immediately after Lagrangian cloud evolution:

$$\texttt{parcels.evolve()} \longrightarrow \texttt{smoothSourceTerms.H} \longrightarrow \text{PIMPLE flow solution.} \tag{3.18}$$

This ordering is important. During `parcels.evolve()`, the Lagrangian tracking, breakup and drag models accumulate the raw PCM source fields `UTrans` and `UCoeff`. The smoothing routine then redistributes these fields before the Eulerian momentum

predictor uses them. If the smoothing were applied after the momentum equation, the gas phase would still receive the raw cell-centred PCM sources.

The implementation acts on both the explicit and implicit momentum source fields:

$$\begin{aligned} \texttt{UTrans} &\rightarrow \texttt{UTrans}_s, \\ \texttt{UCoeff} &\rightarrow \texttt{UCoeff}_s. \end{aligned} \tag{3.19}$$

For `UTrans`, the vector field is treated component-wise by the finite-volume discretisation. For `UCoeff`, the same screened diffusion operator is applied to the scalar coefficient field. In both cases, the per-cell cloud source is first converted into a volumetric density by division by the cell volume, smoothed on the Eulerian mesh, and then multiplied back by the cell volume before being written to the cloud source arrays.

The smoothing equation is discretised with zero-gradient boundary conditions, consistent with the conservative property discussed in Section 2.8. The bandwidth is read from the case dictionary through `smoothBandwidth`. The precise value is treated as a numerical parameter and is investigated in the DBM sensitivity study.

## 3.9 Time-Loop Integration and PIMPLE Coupling

The coupling between the Lagrangian cloud, the DBM smoothing and the Eulerian flow solver is performed once per Eulerian time step. The implemented sequence is summarised in Table 3.6.

| Step | Operation |
|---|---|
| 1 | Read time controls, update Courant number and time step. |
| 2 | Advance the Lagrangian cloud with `parcels.evolve()`. This includes tracking, interpolation of carrier properties, breakup, child creation, drag evaluation and accumulation of raw cloud source fields. |
| 3 | Apply DBM smoothing to `UTrans` and `UCoeff`. |
| 4 | Enter the Eulerian flow solution and solve density, momentum, species/energy and pressure equations according to the PIMPLE algorithm. |
| 5 | Write fields, parcels and optional debug outputs. |

**Table 3.6:** *Time-loop sequence of the coupled solver.*

The parcel motion itself may be sub-stepped internally by the OpenFOAM Lagrangian tracking procedure according to the usual parcel-tracking controls. The breakup update receives the sub-step duration $\Delta t$ used for the current parcel update. Therefore the KH/Wave relaxation, column-breakup clock and PE timing are evaluated on the actual Lagrangian update interval rather than only on the outer Eulerian write interval. In practice, the global time step is selected so that both the carrier-flow Courant number and the near-injector breakup timescales remain adequately resolved.

## 3.10 Case Dictionaries and Runtime Parameters

The implementation is controlled through the usual OpenFOAM runtime dictionaries. The exact case dictionaries used for each validation case are reported with the corresponding numerical setup, but the relevant model entries are summarised here to clarify reproducibility.

The spray-cloud dictionary selects the custom breakup model and exposes the model constants:

**Schematic breakup-model dictionary entries**

```
breakupModel Madabhushi;

MadabhushiCoeffs
{
  C0 ...;
  FLig ...;
  b0 ...;
  b1 ...;
  Dinj ...;
  debug ...;
}
```

The injector diameter `Dinj` is required by the breakup model even when the same value is also specified in the drag-force dictionary.

The drag force is selected in the cloud force sub-dictionary:

**Schematic drag-force dictionary entries**

```
particleForces
{
  MadabhushiDrag
  {
    type MadabhushiDragForce;
    CdBlob ...;
    C0 ...;
    Dinj ...;
  }
}
```

The DBM smoothing bandwidth is read at solver level, for example:

**Schematic DBM dictionary entry**

```
smoothBandwidth ...;
```

The coefficient $B_0$ is intentionally kept as a runtime parameter. The Lambert LJICF cases use values consistent with the Lambert reference setup, whereas the Zhang LJIGF/GLOP cases use a separately calibrated value because the confined gas-film

configuration modifies the near-field stripping behaviour.

## 3.11 Build and Execution Workflow

The custom parcel, breakup and drag components are compiled as user libraries through the standard OpenFOAM `wmake libso` workflow. The typical organisation is:

- `liblagrangianSpray.so`, containing the modified `SprayParcel.C` required for multi-child creation and state reset;
- `libMadabhushiSprayModels.so`, containing the custom breakup model `Madabhushi.C/H`;
- `libMadabhushiDragModels.so`, containing `MadabhushiDragForce.C/H`;
- the modified solver executable `sprayFoamDBM`;
- case-level dictionaries loading the user libraries and selecting the custom models at runtime.

A schematic build sequence is:

**Schematic build sequence**

```
wmake libso userLibs/lagrangianSpray
wmake libso userLibs/MadabhushiSprayModels
wmake libso userLibs/MadabhushiDragModels
wmake sprayFoamDBM
```

Two aspects of the build configuration deserve explicit mention. First, `SprayParcel::calcBreakup` is a class template whose instantiation occurs at the call site inside `sprayFoamDBM`; it therefore compiles into the solver binary rather than into `liblagrangianSpray.so`. Consequently, the `Make/options` file of the solver must list the local `lnInclude` path of the customised spray sources *before* the system `lagrangian/intermediate/lnInclude` path, so that the compiler resolves the modified template against the local version rather than the unmodified system header. Second, the `override` specifier cannot be used reliably in the derived breakup class when the base-class template is instantiated from a pre-compiled system shared library; the multi-child virtual methods in `Madabhushi.H` are therefore declared as `virtual` without `override`.

At runtime, the case loads the custom libraries through the usual `libs` entry in the control dictionary or in the cloud properties dictionary, depending on the local case organisation. The important reproducibility point is that `liblagrangianSpray.so`, the breakup model, the drag force and the solver must all be compiled against the same OpenFOAM version and selected consistently in the case dictionaries.

## 3.12 Debug Outputs and Numerical Consistency Checks

Several optional debug outputs were introduced during development. They are not part of the production post-processing workflow, but they were essential to verify that the state machine, child creation and drag selection behaved as intended.

The breakup model writes separate CSV-style logs for Wave stripping, column-breakup latch, PE evolution, PE breakup and standard PE admission. The drag model writes a separate log containing the selected drag branch, current state markers, effective diameter, Reynolds number, drag coefficient and area correction. The parcel-cloning routine also records parent and child states at creation events.

The main consistency checks performed during implementation were:

- **Mass conservation at breakup.** For every child-creation event, the sum of the represented child masses and the retained parent mass was checked against the pre-breakup represented mass, within round-off and parcel-weight tolerances.
- **Correct state reset of children.** Newly created children were checked to ensure that they did not inherit stale `y`, `yDot`, `KHindex` or `tc` values from the parent.
- **Lambert-factor localisation.** The ligament correction was verified to act only when `ms` $= -1$, i.e. during the first PE breakup of the original core parcel.
- **Drag-regime synchronisation.** Core parcels were checked to retain blob drag until column breakup, while child parcels were checked to receive spherical drag before any PE latch and deformation/disc drag only during PE-active evolution.
- **DBM conservation.** The volume integral of the smoothed `UTrans` and `UCoeff` fields was compared with the corresponding raw PCM integral before smoothing, consistently with the conservation property of Equation (2.70).
- **Analytical timing checks.** For selected parcels in simplified carrier-flow conditions, the computed column-breakup time was compared against the analytical expression of Equation (3.6).

These checks do not replace validation against experiments, which is presented in the following chapters. They ensure, however, that discrepancies observed in the validation cases can be attributed to modelling assumptions and numerical setup rather than to obvious inconsistencies in the parcel state machine.

## 3.13 Implementation Summary

The final implementation keeps the OpenFOAM spray framework intact while adding the state information required by the Madabhushi–Lambert model. The breakup library controls the physical state machine and fragment properties; `SprayParcel.C` performs conservative child creation and state reset; the drag library selects the force law from the current parcel state; and the solver applies DBM smoothing to the accumulated Eulerian source terms before the flow equations are advanced.

This separation of responsibilities improves robustness and maintainability. The physical modelling assumptions remain concentrated in the breakup and drag libraries, while the solver-level change is limited to source-term smoothing. At the same time, the persistent parcel flags provide enough information to keep breakup, child creation, drag and Eulerian coupling synchronised over the full parcel lifetime.

# 4 Lambert LJICF Validation: J10 Sensitivity and J20 Follow-up

## 4.1 Purpose of the Lambert Validation

The Lambert LJICF dataset, introduced in Lambert et al. [8], provides a well-characterised experimental and computational reference for Jet-A1 kerosene injected as a round liquid jet into a subsonic air crossflow. The injector diameter is $D_{\mathrm{inj}} = 457\,\mu$m and the momentum flux ratio spans two conditions: J10 ($J \approx 10$) and J20 ($J \approx 20$), where $J = \rho_l U_l^2/(\rho_g U_g^2)$ is the liquid-to-gas momentum flux ratio defined in Equation (2.3). At $J \approx 10$ the crossflow is energetic relative to the liquid jet, deflecting the column and triggering primary breakup within a moderate streamwise distance from the injector while a coherent ligament region persists close to the nozzle; at $J \approx 20$ the liquid momentum is comparatively stronger, so the column penetrates further into the crossflow and bends more gradually, sustaining a longer intact-column region before primary breakup.

Phase-Doppler measurements of the Sauter mean diameter $D_{32}$ and the streamwise spray velocity $U_x$ are available at two cross-sections downstream of the injector: $x/d = 30$ and $x/d = 60$, where $d = D_{\mathrm{inj}}$. These experimental data were originally reported by Gopala et al. [17] and Sekar et al. [18] and were subsequently adopted by Lambert et al. [8] to validate their Fluent SBES-DPM solution; the benchmark used in the present work therefore comprises both the experimental measurements and the Lambert numerical solution, with Lambert et al. [8] retained as the primary reference for the combined dataset. The nearer plane captures the spray during its initial dispersion phase, when the droplet size distribution still retains memory of the primary breakup; the farther plane reflects the more developed state following additional secondary breakup and inertial sorting. Together they constrain both the breakup-model calibration and the source-term diffusion sensitivity. Alongside the experimental data, Lambert et al. [8] also provide a reference Lagrangian Discrete Phase Model (DPM) solution obtained with the SBES turbulence model in ANSYS Fluent, which serves as a secondary validation target representing a high-fidelity RANS/LES-hybrid benchmark.

The objective of this chapter is to assess the predictive accuracy of the Madabhushi–Lambert primary breakup model implemented in the present OpenFOAM solver—described in Chapter 2 and detailed in Chapter 3—against this benchmark. The J10 configuration is used first for a systematic sensitivity study of the DBM bandwidth parameter $b$, which governs the spatial extent of the parcel source-term diffusion defined in Section 2.8. The J20 case then repeats the same bandwidth sweep at a higher momentum flux ratio, serving as a follow-up validation that tests whether the

compromise identified for J10 remains valid as $J$ increases; no physical parameter of the breakup model is re-calibrated between the two conditions. Both $D_{32}$ and $U_x$ profiles are examined: the former mainly reflects breakup, coalescence, transport, inertial segregation and the trajectory history of the parcels, whereas the latter reflects the parcel acceleration and the two-way momentum exchange with the carrier phase.

## 4.2 Numerical Setup and Post-processing Methodology

The computational domain reproduces the wind-tunnel geometry of Lambert et al. [8]: a rectangular test section with the crossflow entering from the inlet and the liquid injected through a wall-flush circular nozzle (Figure 4.1). Table 4.1 summarises the geometric, mesh, and operating parameters for both momentum flux conditions, and Figure 4.2 shows the structured hexahedral grid.

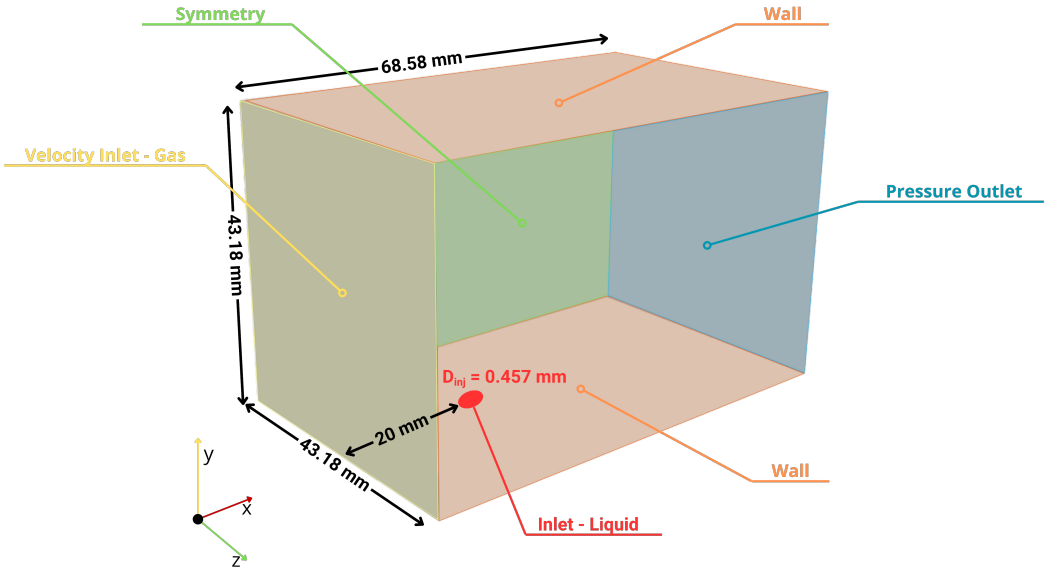
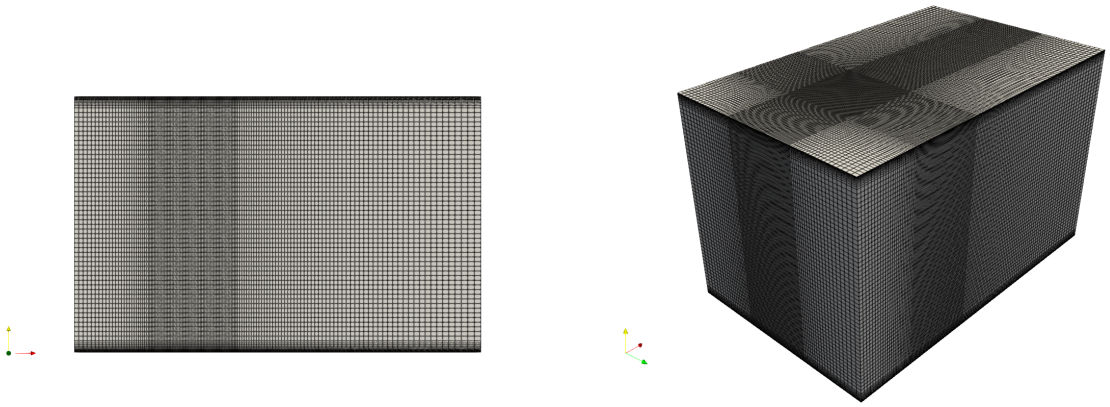



**Figure 4.1:** *Lambert LJICF computational domain and boundary conditions.*



**4.2a** *Front view* **4.2b** *Isometric view*

**Figure 4.2:** *Lambert LJICF mesh: structured hexahedral grid with near-nozzle in-plane refinement and wall-normal boundary-layer clustering.*

| Parameter | Symbol | J10 | J20 | Unit |
|---|---|---|---|---|
| *Case-specific* | | | | |
| Momentum flux ratio | $J$ | $\approx 10$ | $\approx 20$ | – |
| Injection velocity | $U_l$ | 30.4 | 42.9 | m/s |
| Liquid density | $\rho_l$ | $\approx 800$ | | kg/m$^3$ |
| Volumetric flow rate | $Q$ | $4.98 \times 10^{-6}$ | $7.04 \times 10^{-6}$ | m$^3$/s |
| *Domain and mesh* | | | | |
| Domain $L \times H \times W$ | – | $68.6 \times 43.2 \times 43.2$ | | mm$^3$ |
| Nozzle distance from inlet | – | 20 | | mm |
| Mesh ($145 \times 110 \times 100$) | – | $\approx 1.6 \times 10^6$ | | cells |
| Near-nozzle resolution (in-plane) | – | $\approx 0.25$ | | mm |
| First-cell height | – | $\approx 1.2$ | | $\mu$m |
| Wall-normal resolution | $y^+$ | $\approx 1$ | | – |
| *Carrier phase* | | | | |
| Turbulence model | – | $k$–$\omega$ SST (RAS) | | – |
| Crossflow velocity | $U_g$ | 112 | | m/s |
| Inlet turbulent kinetic energy | $k$ | 47 | | m$^2$/s$^2$ |
| Inlet specific dissipation rate | $\omega$ | 4140 | | s$^{-1}$ |
| Pressure | $p$ | $5.07 \times 10^5$ | | Pa |
| Temperature | $T$ | 300 | | K |
| Gas density | $\rho_g$ | $\approx 5.9$ | | kg/m$^3$ |
| *Dispersed-phase parameters* | | | | |
| Injector diameter | $D_{\text{inj}}$ | 457 | | $\mu$m |
| Column-breakup time coefficient | $C_0$ | 3.44 | | – |
| KH child-size coefficient | $b_0$ | 2.44 | | – |
| KH breakup-time coefficient | $b_1$ | 10 | | – |
| Ligament correction factor | $F_{\text{Lig}}$ | 0.4 | | – |
| Blob drag coefficient | $C_{D,\text{blob}}$ | 1.48 | | – |
| Disc drag coefficient | $C_{D,\text{disc}}$ | 1.2 | | – |
| *Numerical method* | | | | |
| Parcel interpolation | – | `cell` | | – |
| DBM bandwidths | $b$ | 0, $D_{\text{inj}}$, $3D_{\text{inj}}$, $6D_{\text{inj}}$ | | – |

**Table 4.1:** *Lambert LJICF numerical setup and dispersed-phase model parameters.*

The dispersed phase is governed by the Madabhushi–Lambert breakup model (Chapter 3), which handles KH/Wave surface stripping, column-breakup latching, the Lambert ligament correction and Pilch–Erdman secondary breakup, with inter-parcel collisions handled by the Nordin [12] trajectory model with coalescence enabled (Section 2.2). The DBM source-term treatment applies the screened-diffusion smoothing operator of Sun and Xiao [10] (Section 2.8) to the interphase momentum-exchange source terms, with the parameter $b$ interpreted as a Gaussian-equivalent bandwidth; the same four bandwidth values are applied to both the J10 and J20 conditions.

The measurement planes and the compared quantities follow the Lambert benchmark [8], whereas the numerical sampling is defined independently for the present solver. Wall-normal profiles of $D_{32}$ and $U_x$ are extracted along the $y$-direction at the spanwise symmetry plane $z = 0$ using 30 uniform bins over $y \in [0, 15]$ mm within a transverse window $|z| \leq 1$ mm, and bins containing fewer than 100 parcel crossing events are excluded. Within each retained bin $D_{32}$ is computed from the `nParticle`-weighted parcel population (Equation (2.6)) and $U_x$ is the corresponding number-weighted mean parcel velocity; no additional $|U_x|$ flux-area correction is applied. This numerical sampling definition should be kept in mind when comparing with the Phase-Doppler measurements, whose weighting is not necessarily identical and is not documented in detail for the present quantities.

## 4.3 J10: DBM Bandwidth Sensitivity on $D_{32}$ and $U_x$ Profiles

Figures 4.3 and 4.4 show the wall-normal profiles of $D_{32}$ and $U_x$ at $x/d = 30$ and $x/d = 60$ for the full DBM bandwidth sweep applied to the J10 case. Each panel displays the Fluent/Lambert SBES-DPM reference (solid blue), the experimental phase-Doppler data (orange squares), the PCM baseline ($b = 0$, green, `ce11` interpolation), and the three DBM variants $b = D_{\text{inj}}$ (red), $b = 3D_{\text{inj}}$ (purple), and $b = 6D_{\text{inj}}$ (brown). A curve is plotted only in those wall-normal bins that contain at least 100 parcel crossing events; where a curve terminates, the interruption therefore reflects an insufficient statistical sample in that band of $y$ rather than the physical absence of dispersed phase.

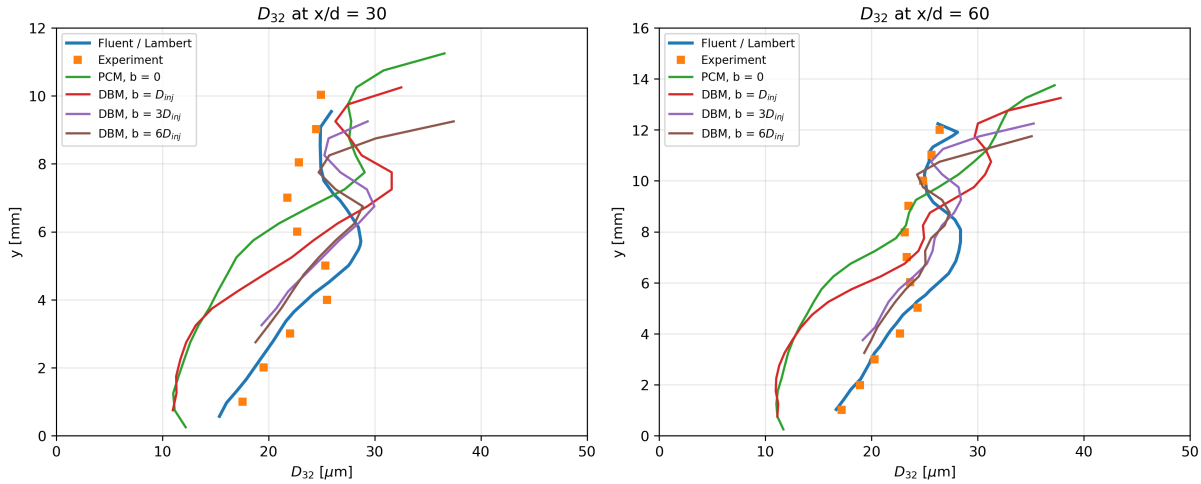



**Figure 4.3:** *$D_{32}$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J10: PCM and DBM $b \in \{D_{\text{inj}}, 3D_{\text{inj}}, 6D_{\text{inj}}\}$, with SBES and experiment.*

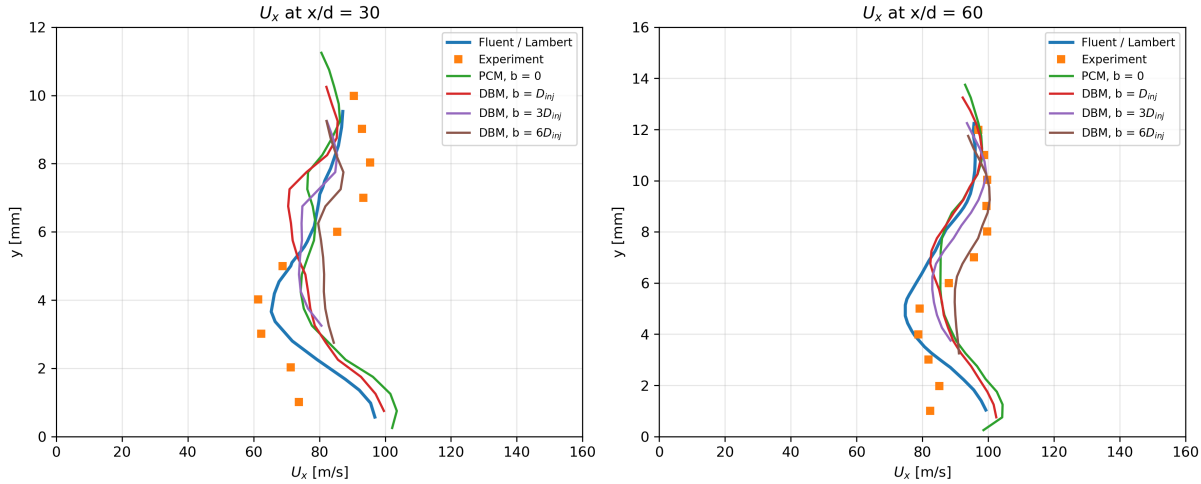



**Figure 4.4:** *$U_x$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J10; same curves as Figure 4.3.*

The most prominent feature of the $D_{32}$ comparison at $x/d = 30$ is an increase of the simulated Sauter mean diameter in the mid-to-upper portion of the spray. This plane lies approximately in the region where the first Pilch–Erdman breakup event is expected to occur, so it can intercept parent parcels close to completing their deformation and breakup, parcels still undergoing breakup, and only recently generated fragments; their size distribution has not yet reached the more developed state observed farther downstream. In the approximate wall-normal band $y \approx 6$–$9$ mm the simulated profiles display elevated $D_{32}$ values and a local bulge in the mid-to-upper portion of the spray. Several OpenFOAM configurations overestimate the experimental data over part of this band, most clearly around $y \approx 7$–$9$ mm, but the behaviour is not identical for every curve and at every point: in the lower part of this interval the PCM baseline can instead remain below the measurements. The PCM baseline (green) shows the largest penetration in the $y$ direction, extending the populated region to higher elevations than the DBM cases, so that part of its overestimation is displaced toward higher $y$ and toward the outer edge of the plume.

Closer to the wall the behaviour differs between bandwidths. At low $y$ the PCM and $b = D_{\text{inj}}$ cases significantly underestimate both the experimental data and the Lambert numerical solution. The $b = 3D_{\text{inj}}$ and $b = 6D_{\text{inj}}$ curves, by contrast, do not reach the lowest band, because the wider spatial diffusion of their source terms produces a more compact spray structure and leaves fewer than 100 events in the bins with $y \lesssim 3$ mm. Their absence in this band must not be read as agreement with the experiment, nor as a physical absence of parcels; it rather signals a different spatial distribution of the parcels. Above the near-wall region, $b = 3D_{\text{inj}}$ and $b = 6D_{\text{inj}}$ generally provide the best overall agreement with both the experimental data and the Lambert solution, although neither is superior at every individual position.

The two larger bandwidths remain close to one another over much of the common interval, which suggests a response that tends to saturate as $b$ increases; the behaviour is nevertheless not monotonic point by point. At $x/d = 30$, $b = 3D_{\text{inj}}$ represents the best overall compromise, whereas at $x/d = 60$, $b = 6D_{\text{inj}}$ attains the lowest global error, as quantified in Table 4.2. At the farther plane the breakup is more developed and the profiles become more regular; increasing $b$ shifts the intermediate region toward larger diameters, correcting the PCM underestimation, but it can produce a local overestimation at the upper elevations. The associated improvement in the global metrics is therefore better described as a compensation between different ranges of $y$ than as a uniform, point-by-point convergence.

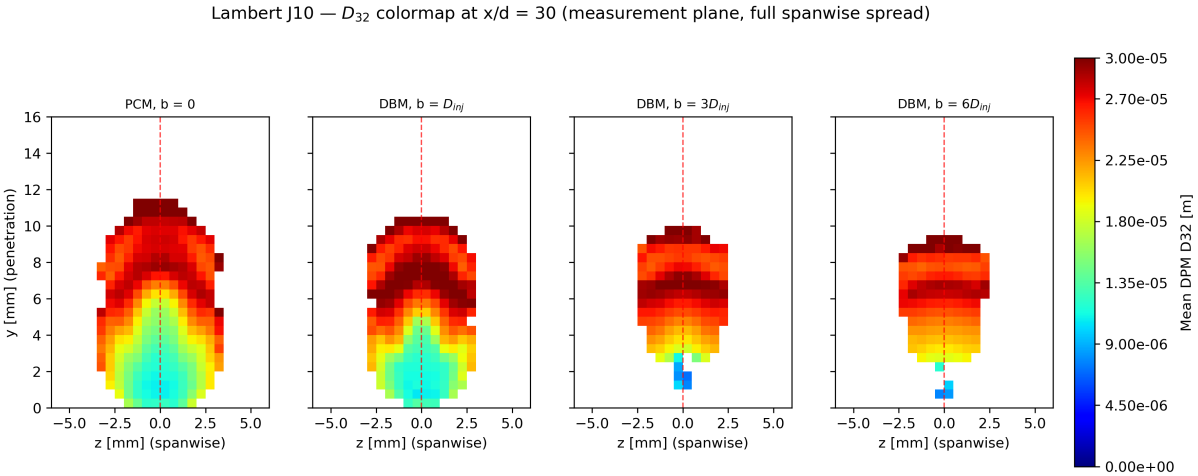
The localised $D_{32}$ bulge observed at the near plane attenuates by $x/d = 60$, where the distribution has had more time to evolve through secondary breakup and inertial selection. Its position is consistent with the passage of the core of the deflecting column and with the presence of parcels that have just completed, or are still completing, their first Pilch–Erdman cycle. In the SBES reference the resolved turbulent structures may promote an earlier deformation and fragmentation of the column, so that the $x/d = 30$ plane can sample an already more developed region in the Lambert solution, whereas in the SST-RAS OpenFOAM simulation part of the population may still be in deformation or active breakup. The central offset between the two solutions can thus be associated with a different breakup timing induced by the turbulence treatment and by the consequent evolution of the relative velocity. This interpretation should be regarded with caution: it is a physically plausible explanation rather than a conclusion demonstrated by the profiles alone, and it could be verified through statistics conditioned on the parcel state and on the location of the breakup events.

The two-dimensional distribution of $D_{32}$ over the measurement-plane cross-section is shown in panels (a) and (b) of Figure 4.5 as colour maps in the spanwise–wall-normal $(z, y)$ plane at $x/d = 30$ and $x/d = 60$ respectively. Each composite juxtaposes the four bandwidths—PCM ($b = 0$), $b = D_{\text{inj}}$, $b = 3D_{\text{inj}}$, and $b = 6D_{\text{inj}}$—on a common colour scale, with a vertical dashed line marking the $z = 0$ symmetry axis. In every panel the spray shows a cool near-wall core (cyan–green, small $D_{32}$) topped by a warm cap (orange–red, large $D_{32}$) in the upper, outer region. Increasing $b$ has two visible effects: the cross-section becomes progressively narrower in the spanwise direction, and the cool core contracts. The contraction is most pronounced at $x/d = 30$, where at $b = 6D_{\text{inj}}$ the cool core nearly disappears and the warm cap dominates a compact plume; at $x/d = 60$ the cool region is reduced but a residual lower-$D_{32}$ band still persists at intermediate $y$. The binned penetration height also decreases with bandwidth, clearly at $x/d = 30$ and to a lesser extent at $x/d = 60$. When $b$ grows well beyond the

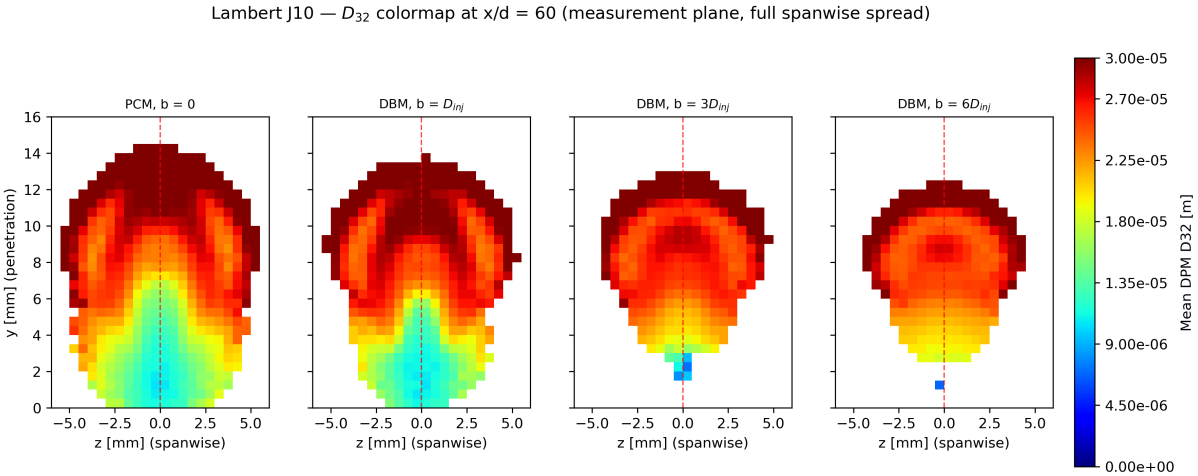
physical column scale ($b \gg D_{\mathrm{inj}}$), the screened-diffusion smoothing operator spreads the carrier momentum source terms over a region appreciably wider than the nozzle, accelerating the carrier in lateral zones and altering the local velocity field that the parcels experience. This redistribution modifies the carrier field, the parcel slip velocity and the acceleration history, and thereby changes the spatial distribution of the parcels that reach the sampling plane, thinning its near-wall, low-$y$ part; the two-dimensional maps consequently show a more compact spray with reduced vertical extent as $b$ increases. The maps display the resulting distribution but do not, on their own, allow a single mechanism responsible for the contraction to be isolated, indicating only that an over-large bandwidth distorts not only the size distribution but also the spatial structure of the spray.

Regarding the $U_x$ profiles (Figure 4.4), the effect of the bandwidth is less pronounced than on $D_{32}$, but it remains significant, particularly in the mid-to-upper band. At $x/d = 30$ the differences are limited in the upper portion of the profile and become appreciable in the central region: $b = 6D_{\mathrm{inj}}$ reproduces several experimental points well between about $y = 6$ and 9 mm, whereas PCM and $b = D_{\mathrm{inj}}$, which retain valid bins down to the lower elevations, markedly overestimate the measured velocity in the near-wall region and remain below the experimental data in part of the intermediate band, where the measurements rise from the local minimum toward the upper plume. The Lambert reference likewise overestimates the near-wall velocity, but it exhibits a more pronounced local minimum between about $y = 3$ and 5 mm. At $x/d = 60$ the agreement of the larger bandwidths improves further: $b = 6D_{\mathrm{inj}}$ follows the experimental data with high accuracy in the approximate band $y = 7$–$12$ mm and, at some points, lies closer to the experiment than the SBES solution, while $b = 3D_{\mathrm{inj}}$ also provides good agreement over its valid interval. Both curves are, however, truncated in the lower region for lack of sufficient parcel crossing events, so their near-wall behaviour cannot be assessed; PCM and $b = D_{\mathrm{inj}}$ are the only cases covering the lower band, where they qualitatively follow the Lambert profile but generally overestimate the experimental data at the lowest elevations.

The origin of the local $U_x$ minimum observed in the Lambert solution cannot be established unambiguously from the one-dimensional profiles alone. Possible contributions include a local deficit of the carrier-phase velocity in the wake of the column, the presence of larger and more inertial droplets, the correlation between diameter and velocity, turbulent dispersion, and the differing statistical weighting of the parcels. The combination observed close to the wall—an underestimation of $D_{32}$ together with an overestimation of $U_x$ in the cases that retain valid bins there—is qualitatively consistent with the effect of a jet wake that is not fully resolved by the SST-RAS closure, although this can be regarded only as one possible contribution rather than the sole cause. Overall, $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ reproduce $U_x$ very well in the mid-to-upper band, whereas the near-wall region can be assessed only for PCM and $b = D_{\mathrm{inj}}$.
The quantitative consequences of the different statistically valid wall-normal ranges are discussed in Section 4.4.




**4.5a** *$D_{32}$ colormap at $x/d = 30$.*




**4.5b** *$D_{32}$ colormap at $x/d = 60$.*

**Figure 4.5:** *Two-dimensional $D_{32}$ maps, J10: PCM, $b = D_{\mathrm{inj}}$, $b = 3D_{\mathrm{inj}}$, $b = 6D_{\mathrm{inj}}$ ($y$ wall-normal, $z$ spanwise, common colour scale).*

## 4.4 Quantitative Error Metrics and Trade-off Analysis

Table 4.2 reports normalised root-mean-square errors (NRMSE) for $D_{32}$ and $U_x$ at both measurement planes, computed against the experimental phase-Doppler data for the four bandwidth values. Each NRMSE is the root-mean-square simulated–measured deviation evaluated at the wall-normal sampling locations along the $y$ direction for which both an experimental measurement and a valid numerical prediction are available, normalised by the mean measured value and reported as a percentage. A numerical profile contributes only through the wall-normal bins that satisfy the minimum threshold of 100 parcel crossing events, so the interval of $y$ effectively spanned by the metric is set, case by case, by the extent of the statistically converged portion of each profile.

| Case | NRMSE $D_{32}$ $x/d = 30$ | NRMSE $D_{32}$ $x/d = 60$ | NRMSE $U_x$ $x/d = 30$ | NRMSE $U_x$ $x/d = 60$ |
|---|---|---|---|---|
| PCM ($b = 0$) | 29.5 | 28.2 | 19.7 | 12.1 |
| DBM $b = D_{\mathrm{inj}}$ | 30.3 | 26.8 | 20.2 | 11.6 |
| DBM $b = 3D_{\mathrm{inj}}$ | 19.1 | 13.9 | 13.9 | 6.5 |
| DBM $b = 6D_{\mathrm{inj}}$ | 21.3 | 10.1 | 17.2 | 6.3 |

**Table 4.2:** *J10 error metrics: NRMSE (in %) for $D_{32}$ and $U_x$ at $x/d = 30$ and $x/d = 60$, computed against the experimental phase-Doppler data.*

The metrics in Table 4.2 make the picture emerging from the profiles of Section 4.3 quantitative, provided they are read together with the differing statistical extent of the underlying curves. For $D_{32}$ at $x/d = 30$ the error is 29.5% for the PCM baseline and essentially unchanged, 30.3%, for $b = D_{\mathrm{inj}}$: a bandwidth comparable with the injector diameter does not improve the size prediction at the near plane. The minimum is reached at $b = 3D_{\mathrm{inj}}$ (19.1%), while $b = 6D_{\mathrm{inj}}$ (21.3%) remains better than the PCM but worse than $b = 3D_{\mathrm{inj}}$; the response is therefore not monotonic in $b$. At $x/d = 60$ the error decreases more regularly as the bandwidth grows— 28.2%, 26.8%, 13.9% and 10.1% for PCM, $b = D_{\mathrm{inj}}$, $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ respectively—with the two larger bandwidths producing the most substantial reductions and $b = 6D_{\mathrm{inj}}$ attaining the lowest value. As already noted for the profiles, this far-plane gain reflects a compensation between different ranges of $y$ rather than a uniform, point-by-point improvement.

The $U_x$ metrics follow a consistent pattern. At $x/d = 30$ the error is 19.7% for the PCM and 20.2% for $b = D_{\mathrm{inj}}$, again without improvement over the baseline; $b = 3D_{\mathrm{inj}}$ gives the minimum (13.9%) and a further increase to $b = 6D_{\mathrm{inj}}$ raises the error again (17.2%), consistent with a very wide diffusion over-modifying the coupled field in the near-wall region. At $x/d = 60$ the errors of the two larger bandwidths are very close, 6.5% for $b = 3D_{\mathrm{inj}}$ and 6.3% for $b = 6D_{\mathrm{inj}}$: $b = 6D_{\mathrm{inj}}$ holds the smallest value, but its margin over $b = 3D_{\mathrm{inj}}$ is marginal, in line with the tendency towards saturation already observed in the profiles.

These metrics must, however, be compared with care, because the four configurations are not evaluated over the same set of experimental positions. As shown in Section 4.3, the $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ profiles do not retain statistically converged bins in the near-wall band, so their global error is computed over a $y$ interval that excludes the

lowest elevations—precisely the region in which the PCM and $b = D_{\mathrm{inj}}$ cases display their largest deviations. A distinction must therefore be drawn between the number of statistically valid numerical bins and the number of experimental points actually entering the metric: the post-processing retains markedly fewer valid wall-normal bins for the wider bandwidths (of the order of 13–14 against 20–23 for the PCM at $x/d = 30$), while the experimental reference itself provides only 10 points at $x/d = 30$ and 12 at $x/d = 60$, and the NRMSE includes only those experimental points that fall within each case's valid span. Intersecting each case's valid span with the experimental positions, at $x/d = 30$ the PCM and $b = D_{\mathrm{inj}}$ cases retain all 10 available points, whereas $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ retain only 6 and 7 respectively; at $x/d = 60$ the corresponding counts are 12, 12, 9 and 8 out of the 12 available. These counts show that the wider bandwidths are evaluated over a reduced subset of the measurements, so that part of their apparent advantage could in principle stem from the excluded near-wall band. To assess the influence of this unequal coverage, the errors were also recomputed over the subset of experimental positions common to the statistically valid span of all four configurations. The resulting common-support metrics, reported in Table 4.4, preserve the original ranking for every plane and observable. The different profile extents therefore affect the absolute NRMSE values, but do not by themselves account for the advantage of the two larger bandwidths.

The two observables respond to the bandwidth variation through different but coupled mechanisms. The DBM does not modify the parcel diameters directly; it redistributes the interphase momentum source terms in space and thereby alters the carrier-phase field sampled by the parcels. The resulting changes in slip velocity, local Weber number, parcel trajectories, residence times and breakup history affect the predicted $D_{32}$ only indirectly. The $U_x$ profiles, by contrast, reflect more directly the Eulerian momentum balance and the two-way momentum exchange between the carrier and dispersed phases, although they are not independent of the underlying Lagrangian population. This distinction is consistent with the $U_x$ error improving by a smaller absolute margin than the $D_{32}$ error and with its rising again at $b = 6D_{\mathrm{inj}}$ at the near plane: the wider smoothing bandwidth modifies the local carrier-phase velocity deficit, alters the parcel acceleration history and reshapes the coupled wake structure, which can penalise the near-field velocity profile.

In summary, for J10 a bandwidth comparable with the injector diameter ($b = D_{\mathrm{inj}}$) stays close to the PCM baseline, $b = 3D_{\mathrm{inj}}$ provides the best near-field compromise for both observables, and $b = 6D_{\mathrm{inj}}$ attains the lowest far-field $D_{32}$ error while offering only a marginal far-field $U_x$ advantage. Increasing $b$ therefore does not produce a monotonic improvement across all quantities and planes, and the choice of bandwidth must weigh accuracy, statistical coverage and preservation of the spatial structure of the spray simultaneously. Whether the compromise identified here carries over to a higher momentum flux ratio is examined for the J20 condition in Section 4.7.

## 4.5 J10: Representative-Bandwidth Assessment

On the basis of the profiles of Section 4.3, the error metrics of Section 4.4 and the colour maps of Figure 4.5, the bandwidth $b = 3D_{\mathrm{inj}}$ is adopted as the preferred representative

bandwidth for the J10 condition. The choice rests on the combination of these three pieces of evidence rather than on the aggregate NRMSE alone.

The smallest bandwidth, $b = D_{\mathrm{inj}}$, offers only a limited benefit over the PCM baseline (Table 4.2) and is therefore not retained. The largest bandwidth, $b = 6D_{\mathrm{inj}}$, improves chiefly the far field, particularly $D_{32}$ at $x/d = 60$, but the colour maps show that it contracts the spray cross-section more markedly and depletes the near-wall parcel population more severely, and at the near plane it degrades the $U_x$ prediction. The intermediate value $b = 3D_{\mathrm{inj}}$ provides the best near-field balance for both observables, while modifying the spatial structure less severely than $b = 6D_{\mathrm{inj}}$ and providing a clearer benefit than $b = D_{\mathrm{inj}}$, and is adopted as the representative bandwidth for J10.

The two diffusive cases do not fully cover the near-wall region, but the common-support metrics of Table 4.4 do not overturn this interpretation, so the advantage of $b = 3D_{\mathrm{inj}}$ is not an artefact of its smaller valid $y$ interval. The selection is thus supported by the joint consistency of profiles, metrics and colour maps, and not by the aggregate error value in isolation. The smoothing length is kept of the order of the physical injection scale, in the spirit of Sun and Xiao [10]: the diffusion bandwidth should be large enough to reduce source localisation and local oscillations of the Lagrangian source terms, yet small enough not to erase physically relevant gradients of the breakup region. The specific value $b = 3D_{\mathrm{inj}}$ is not prescribed by that reference; it satisfies these two requirements for the present mesh and operating point, and is regarded as a case- and mesh-specific representative bandwidth rather than a universal optimum.

The residual high-$y$ overprediction of $D_{32}$, which persists with different intensity for all bandwidths, indicates that the discrepancy cannot be removed by the DBM parametrisation alone. It remains consistent with contributions from the SST-RAS representation of the wake and turbulent transport, the local slip-velocity history, parcel trajectories and sampling effects, none of which can be isolated from the available one-dimensional profiles. The J20 condition is subsequently used in Section 4.7 to assess whether the compromise identified at J10 remains valid when the liquid-to-gas momentum flux ratio is increased, without any further recalibration of the bandwidth.

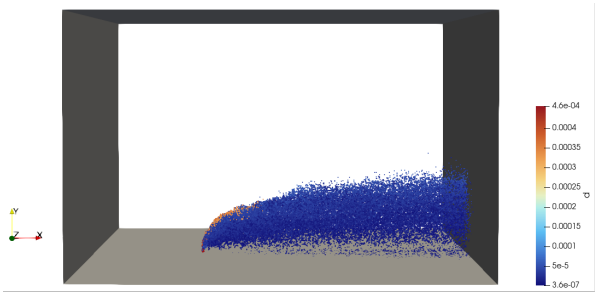
## 4.6 Spray and Carrier-Field Visualisation (J10)

Figures 4.6 and 4.7 provide three-dimensional views of the parcel diameter $d$ and streamwise velocity $U_x$ for the four source-projection bandwidths, whereas Figures 4.8 and 4.9 show the corresponding carrier pressure $p$ and streamwise velocity $U_x$ of the continuous phase on the $z = 0$ symmetry plane. All four figures are representative instantaneous snapshots rather than statistical samples: an exact synchronisation of the physical export time across the cases is not guaranteed, and the colour-bar limits, set individually for each export, are not perfectly uniform across the panels. The case-to-case comparison is therefore qualitative and rests on the shape, position and apparent extent of the flow features rather than on absolute colour values, in contrast to the statistically sampled profiles and maps discussed earlier.

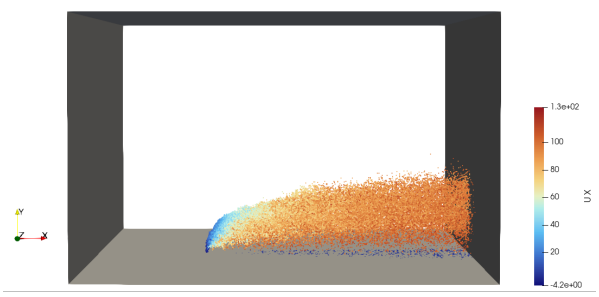
The diameter fields locate the largest parcels near the jet root and along the windward side of the deflected column, a population consistent with core parcels and recently

formed fragments that have not yet completed their aerodynamic breakup; farther downstream the cloud is progressively dominated by smaller parcels. The PCM and $b = D_{\mathrm{inj}}$ cases retain a broader population toward the lower wall-normal region, whereas $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ show a more compact and less near-wall-populated distribution, consistently with the cross-sectional $D_{32}$ maps; the spanwise contraction itself is resolved more clearly in Figure 4.5 than in these side views. The parcel-velocity fields show a progressive acceleration from the low-$U_x$ population near the injector toward the crossflow value downstream, with a slower branch persisting near the lower plume boundary; this wall-normal stratification is consistent with the profiles at $x/d = 30$ and $x/d = 60$. The two larger bandwidths again exhibit a more compact high-velocity region with fewer statistically populated near-wall trajectories, so the bandwidth effect on $U_x$ is weaker than on $D_{32}$ but cannot be regarded as negligible.

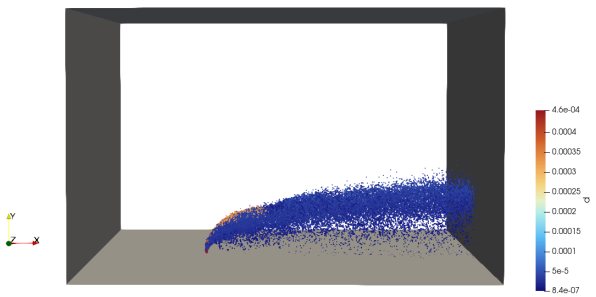
The carrier-pressure sections exhibit the expected windward stagnation-pressure rise immediately upstream of the jet base and a lower-pressure region on the leeward side, while the $U_x$ sections show an elongated velocity deficit that originates at the injector and extends downstream along the lower wall before recovering toward the undisturbed crossflow; this stagnation–wake topology is preserved for all bandwidths. As the bandwidth increases, the carrier $U_x$ field appears progressively less dominated by highly localised momentum-exchange contributions at the parcel locations: the velocity deficit remains present, but it becomes smoother and less cell-localised, consistently with the conservative redistribution of the interphase momentum source by the DBM operator. This redistribution spreads the source spatially without removing the total exchanged momentum, so the wider smoothing bandwidths reshape the local carrier field that sets parcel slip velocity, acceleration and breakup history rather than the global level of coupling. Because the colour-bar limits differ between panels, these observations remain qualitative and rely on the shape, position and apparent extent of the perturbed region, not on absolute colour values.



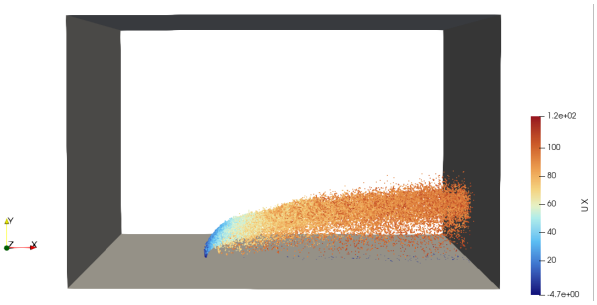
**4.6a** *PCM ($b = 0$): parcel $d$*



**4.6b** *PCM ($b = 0$): parcel $U_x$*

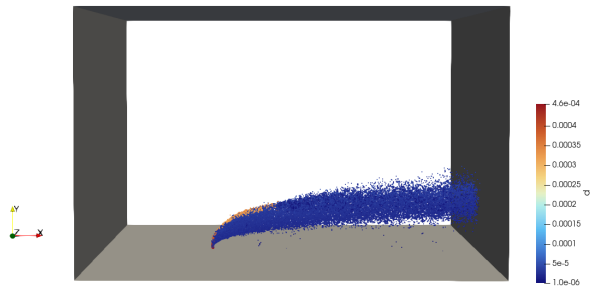


**4.6c** *$b = D_{\text{inj}}$: parcel $d$*

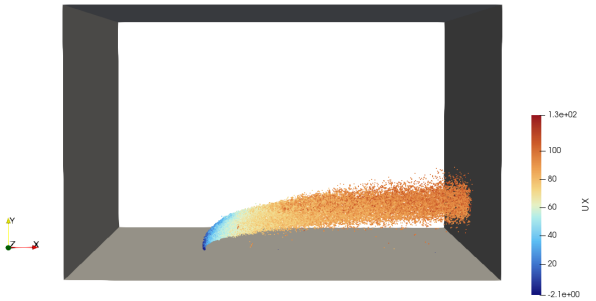


**4.6d** *$b = D_{\text{inj}}$: parcel $U_x$*

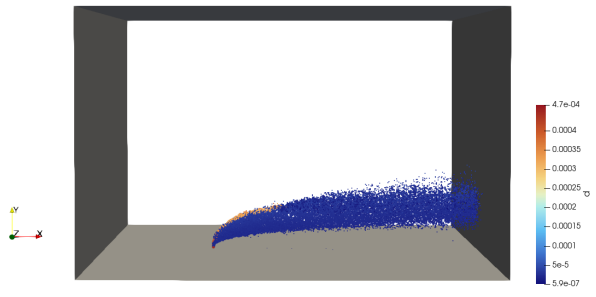
**Figure 4.6:** *J10 spray fields: parcel diameter $d$ (left, in m) and parcel streamwise velocity $U_x$ (right, in m/s), PCM and $b = D_{\text{inj}}$.*



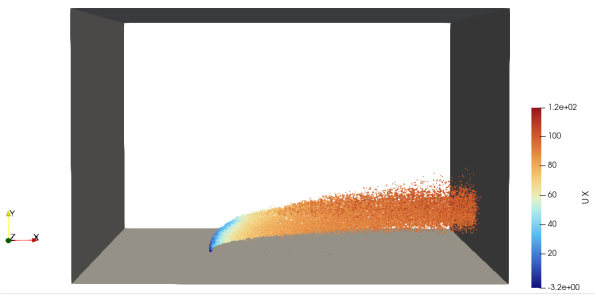
**4.7a** *$b = 3D_{\text{inj}}$: parcel $d$*



**4.7b** *$b = 3D_{\text{inj}}$: parcel $U_x$*

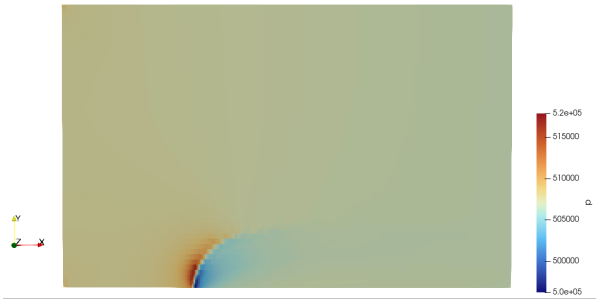


**4.7c** *$b = 6D_{\text{inj}}$: parcel $d$*

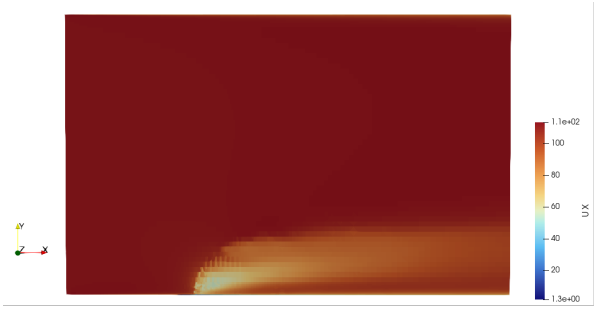


**4.7d** *$b = 6D_{\text{inj}}$: parcel $U_x$*

**Figure 4.7:** *J10 spray fields: parcel diameter $d$ (in m) and parcel streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom).*




**4.8a** *PCM ($b = 0$): carrier $p$*




**4.8b** *PCM ($b = 0$): carrier $U_x$*

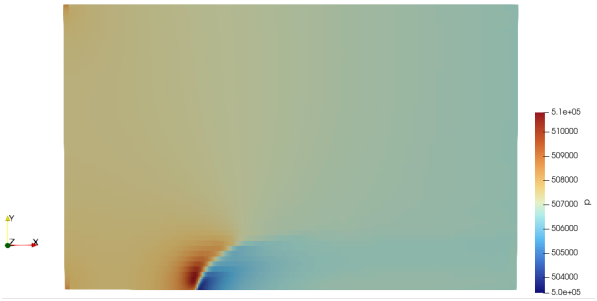



**4.8c** $b = D_{\text{inj}}$*: carrier $p$*




**4.8d** $b = D_{\text{inj}}$*: carrier $U_x$*

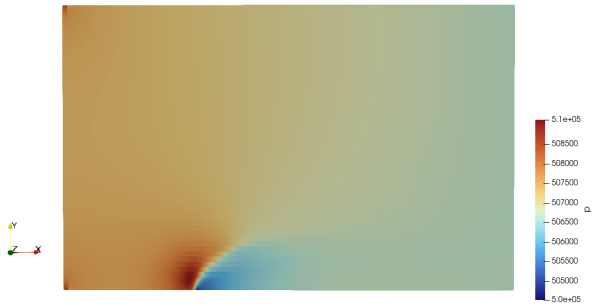
**Figure 4.8:** *J10 carrier fields on the $z = 0$ plane: pressure $p$ (left, in Pa) and carrier streamwise velocity $U_x$ (right, in m/s), PCM and $b = D_{\text{inj}}$.*




**4.9a** $b = 3D_{\text{inj}}$*: carrier $p$*




**4.9b** $b = 3D_{\text{inj}}$*: carrier $U_x$*




**4.9c** $b = 6D_{\text{inj}}$*: carrier $p$*




**4.9d** $b = 6D_{\text{inj}}$*: carrier $U_x$*

**Figure 4.9:** *J10 carrier fields on the $z = 0$ plane: pressure $p$ (in Pa) and carrier streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom).*

## 4.7 J20 Follow-up Validation

The J20 configuration reproduces the Lambert LJICF benchmark at the higher momentum flux ratio $J \approx 20$ and repeats the entire bandwidth sweep $b/D_{\mathrm{inj}} \in \{0, 1, 3, 6\}$ adopted for J10, without any re-calibration of the physical breakup parameters. It therefore serves as a follow-up validation that tests whether the compromise identified at the lower momentum flux ratio remains valid as $J$ increases. The numerical setup and the post-processing methodology are identical to those described in Section 4.2.

The stronger liquid momentum relative to the crossflow increases the wall-normal penetration of the jet, so the plume reaches higher elevations than in the J10 case, consistent with the larger penetration extent of the J20 colour maps (Figure 4.12). The portion of the spray evolution sampled by the near plane is therefore expected to be shifted toward larger $y$ relative to J10.

Figures 4.10 and 4.11 show the $D_{32}$ and $U_x$ profiles at both planes for J20. At $x/d = 30$ the PCM and $b = D_{\mathrm{inj}}$ cases markedly underestimate both the experimental data and the Lambert reference in the lower and intermediate band, approximately up to $y \approx 8$–$9$ mm, whereas in the upper region ($y \approx 10$–$15$ mm) the curves cross one another and display irregularities and local overestimations. The PCM baseline reaches the largest wall-normal penetration and prolongs the profile up to the highest elevations. The $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ cases generally provide the best agreement within their statistically valid range, but they do not reach the near-wall region, because the lower bins do not exceed the threshold of 100 parcel crossing events; their absence at low $y$ must not be read as agreement with the experiment.

At $x/d = 60$ the PCM and $b = D_{\mathrm{inj}}$ cases underestimate $D_{32}$ over much of the lower and intermediate band, whereas $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ correct this underestimation substantially; $b = 6D_{\mathrm{inj}}$ attains the lowest NRMSE but can locally overestimate the uppermost part of the profile, so the improvement is not uniform point by point, and the differing statistical coverage of the four cases must be kept in mind when comparing them. The crossings at $x/d = 30$ are compatible with parcels sampled at different stages of deformation and fragmentation. The wall-normal band in which these crossings and irregularities occur is shifted toward approximately $y = 9$–$13$ mm, consistently with the larger penetration at $J = 20$; this is compatible with a displacement of the sampled transition region, but the one-dimensional profiles do not directly localise the breakup events. The departure from the SBES reference may stem from the different representation of the turbulent structures, of the local relative velocity and of the breakup history in the present SST-RAS framework, and should not be assigned to a single mechanism from the available profiles alone. The SBES reference tracks the experimental data closely at both planes, whereas the present simulations reproduce the overall rise but not its detailed shape.

The $U_x$ profiles for J20 (Figure 4.11) follow the same pattern as J10: the bandwidth influence on the parcel streamwise velocity is less pronounced than that observed for $D_{32}$, although it remains measurable. At $x/d = 30$ the PCM and $b = D_{\mathrm{inj}}$ cases strongly overestimate the experimental velocity in the near-wall region; the Lambert reference also overestimates locally there, but it reproduces a more pronounced minimum. The $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ cases lack valid bins at the lowest elevations, and in the band $y \approx 4$–$7$ mm the OpenFOAM simulations generally remain above the measurements;

the agreement improves in the mid-to-upper band, where $b = 3D_{\text{inj}}$ attains the lowest overall NRMSE. At $x/d = 60$ the experimental points cover approximately the whole interval $y \approx 1$–$15$ mm: the PCM and $b = D_{\text{inj}}$ cases overestimate the velocity at the lower elevations, $b = 3D_{\text{inj}}$ follows the data closely in the central band, and $b = 6D_{\text{inj}}$ tends to retain higher values, most visibly around $y \approx 5$–$8$ mm, while the curves converge more in the upper region. As in J10, the local minimum of the SBES profile is not reproduced with the same intensity; because these curves represent the velocity of the parcels and not the carrier field directly, its origin cannot be isolated from the one-dimensional profiles and is likely affected by a combination of the carrier wake, the parcel inertia, the size–velocity correlation, the turbulent dispersion and the statistical weighting of the sampling. The corresponding NRMSE values are summarised in Table 4.3.

| Case | NRMSE $D_{32}$ $x/d = 30$ | NRMSE $D_{32}$ $x/d = 60$ | NRMSE $U_x$ $x/d = 30$ | NRMSE $U_x$ $x/d = 60$ |
|---|---|---|---|---|
| PCM ($b = 0$) | 27.9 | 31.9 | 18.1 | 9.5 |
| DBM $b = D_{\text{inj}}$ | 23.8 | 26.0 | 18.0 | 9.0 |
| DBM $b = 3D_{\text{inj}}$ | 11.4 | 9.9 | 14.5 | 5.2 |
| DBM $b = 6D_{\text{inj}}$ | 11.8 | 7.5 | 18.6 | 6.5 |

**Table 4.3:** *J20 error metrics: NRMSE (in %) for $D_{32}$ and $U_x$ at $x/d = 30$ and $x/d = 60$, computed against the experimental phase-Doppler data.*

The metrics in Table 4.3 reproduce the qualitative ordering already observed for J10. For $D_{32}$ the two larger bandwidths reduce the error strongly relative to both the PCM baseline and $b = D_{\text{inj}}$: at $x/d = 30$ they are very close, with $b = 3D_{\text{inj}}$ (11.4%) marginally ahead of $b = 6D_{\text{inj}}$ (11.8%), whereas at $x/d = 60$ the smallest $D_{32}$ error is attained by $b = 6D_{\text{inj}}$ (7.5%). For $U_x$ the minimum error is obtained with $b = 3D_{\text{inj}}$ at both planes, whereas $b = 6D_{\text{inj}}$ noticeably worsens the near-plane velocity prediction, consistent with a very wide diffusion over-modifying the coupled field close to the injector. As discussed for J10 in Section 4.4, these values must be read together with the differing statistical coverage of the profiles, since the wider bandwidths are evaluated over fewer experimental points: at $x/d = 30$ the PCM, $b = D_{\text{inj}}$, $b = 3D_{\text{inj}}$ and $b = 6D_{\text{inj}}$ metrics use 11, 11, 8 and 9 experimental points respectively, and at $x/d = 60$ they use 13, 13, 9 and 10 of the 14 available points. Because $D_{32}$ and $U_x$ are evaluated using the same parcel-crossing validity mask and the same experimental sampling positions in the Lambert benchmark [8], these point counts apply to both metrics. Overall, $b = 3D_{\text{inj}}$ is the most balanced compromise for J20—best on $U_x$ at both planes and very close to the minimum on $D_{32}$ at the near plane—rather than the minimum for every individual metric.

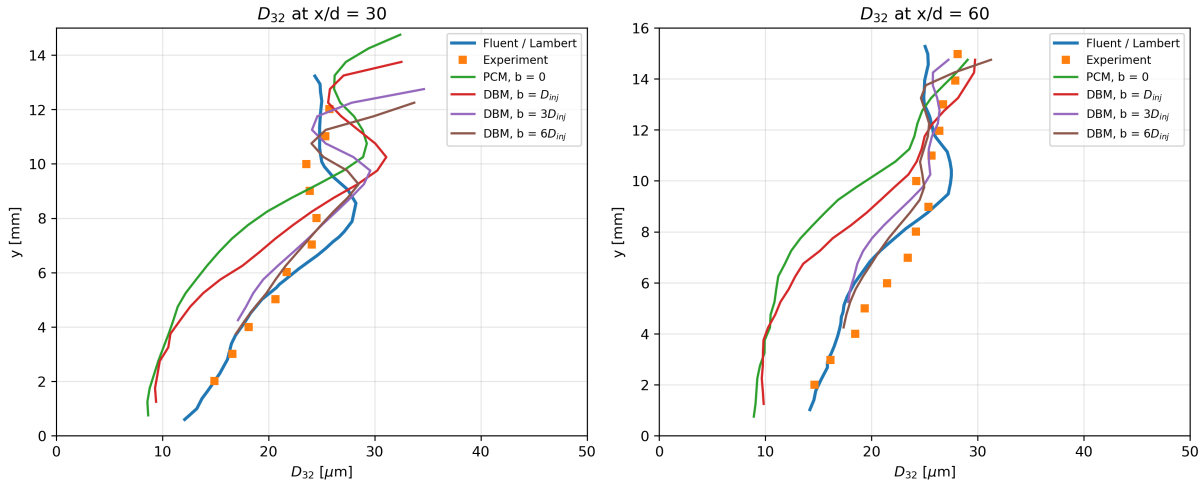



**Figure 4.10:** *$D_{32}$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J20: PCM and DBM $b \in \{D_{\text{inj}}, 3D_{\text{inj}}, 6D_{\text{inj}}\}$, with SBES and experiment.*

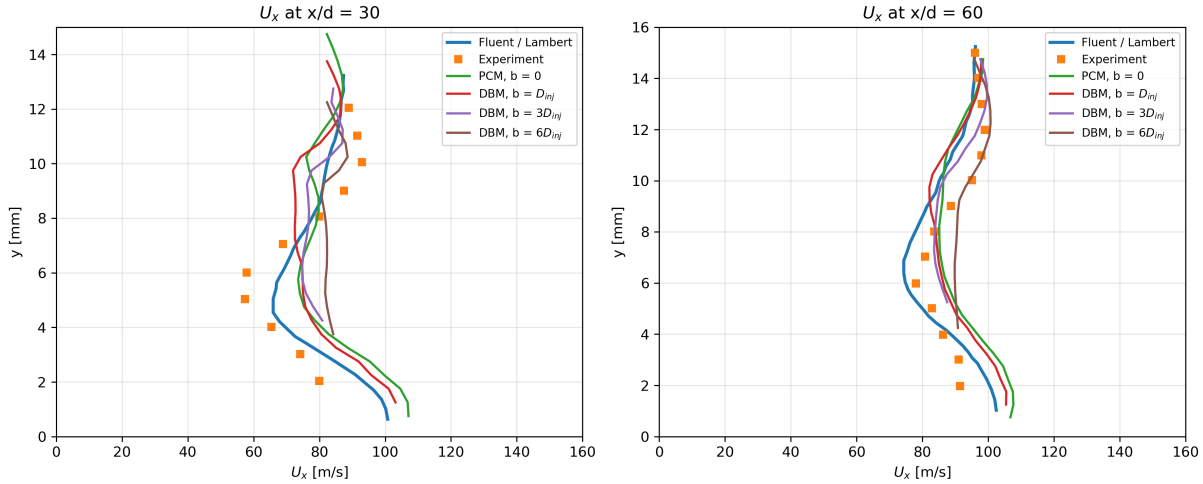
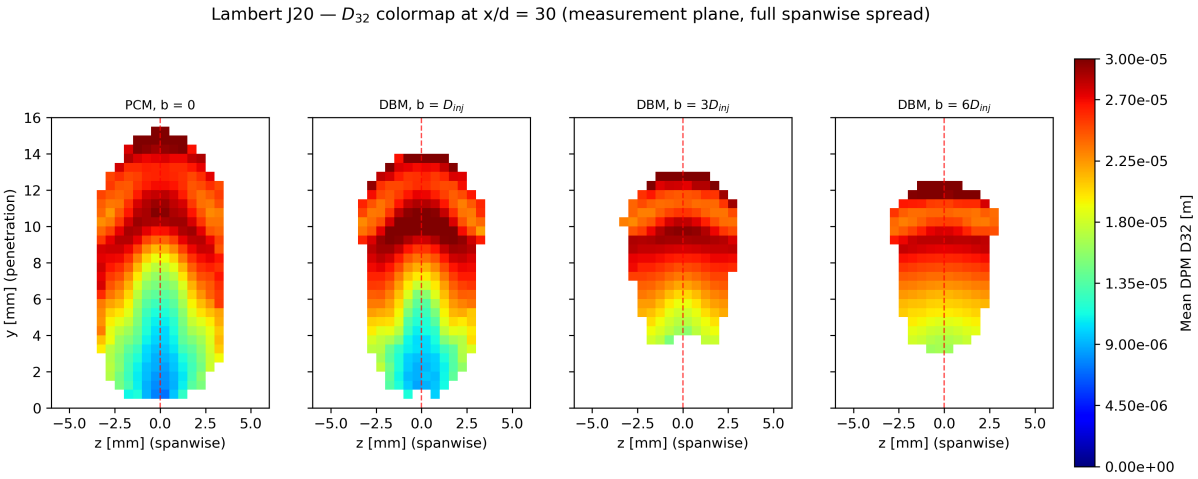
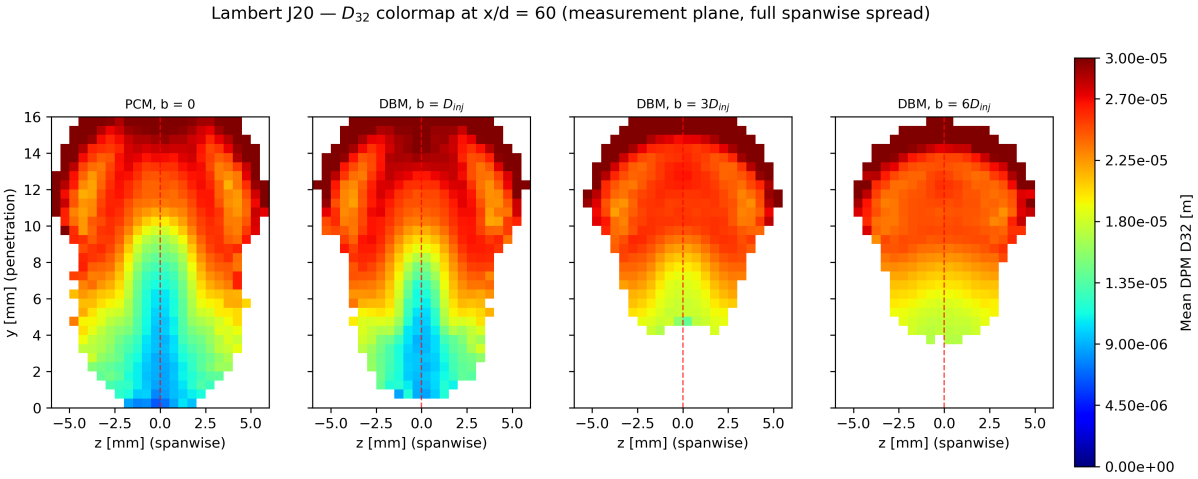



**Figure 4.11:** *$U_x$ wall-normal profiles at $x/d = 30$ (left) and $x/d = 60$ (right), J20; same curves as Figure 4.10.*

The two-dimensional $D_{32}$ colour maps at both planes are shown in Figure 4.12. The same overall structure observed for J10 is reproduced: a cool low-$D_{32}$ region beneath a warm high-$y$ cap. For the PCM and $b = D_{\text{inj}}$ cases this cool region extends down to elevations close to the wall, whereas for $b = 3D_{\text{inj}}$ and $b = 6D_{\text{inj}}$ the map begins only at $y \approx 3$–$4$ mm, because the lower bins do not contain a sufficient parcel population to be retained. Increasing $b$ narrows the section in the spanwise direction, reduces the valid population in the lower region, and makes the high-$D_{32}$ region occupy a larger fraction of the residual section; the binned penetration height decreases clearly at $x/d = 30$ and only marginally at $x/d = 60$, where the four panels reach comparable wall-normal extents. Consistent with the higher momentum flux ratio, the J20 plume penetrates to larger wall-normal distances than the J10 plume at the same plane. The contraction is the combined result of the modified carrier field and parcel statistics and cannot be attributed to a single direction of parcel deflection. The vertical dashed line in each panel marks the $z = 0$ symmetry axis.




**4.12a** *$D_{32}$ colormap at $x/d = 30$.*

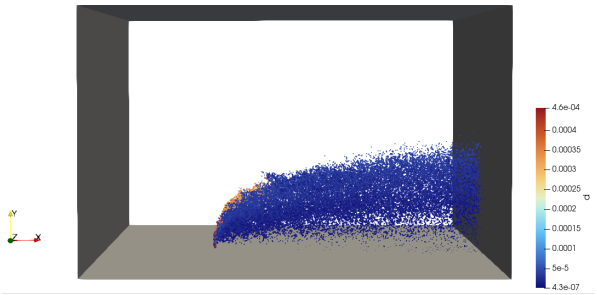



**4.12b** *$D_{32}$ colormap at $x/d = 60$.*

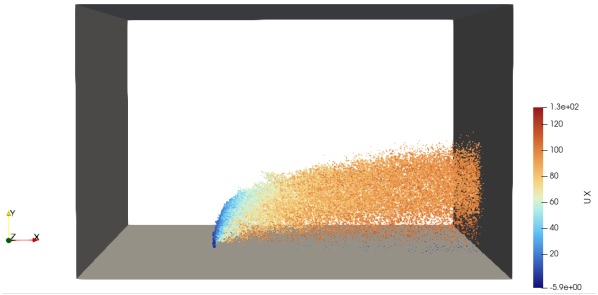
**Figure 4.12:** *Two-dimensional $D_{32}$ maps, J20: PCM, $b = D_{\mathrm{inj}}$, $b = 3D_{\mathrm{inj}}$, $b = 6D_{\mathrm{inj}}$ ($y$ wall-normal, $z$ spanwise, same colour scale as Figure 4.5).*

The three-dimensional parcel views (Figures 4.13 and 4.14) show larger-diameter parcels near the jet root and along the windward envelope, with smaller parcels dominating downstream; as the parcel state is not displayed, they locate the larger-diameter region, not the instant of breakup. Parcels accelerate from low $U_x$ at the root, retaining a wall-normal velocity stratification. With non-common colour bars the bandwidth differences are only qualitative.

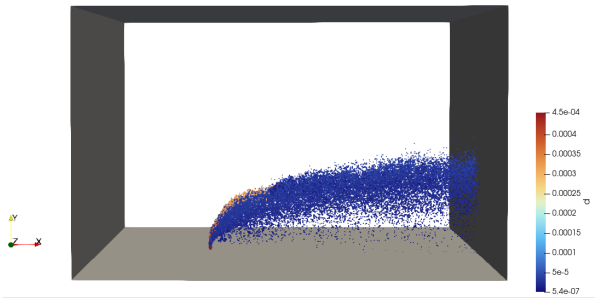
The two-dimensional carrier sections on the $z = 0$ plane (Figures 4.15 and 4.16) show a windward pressure rise, a leeward pressure reduction, a $U_x$ wake deficit and a gradual recovery toward the undisturbed crossflow. As the colour bars are not uniform across cases, only the shape, position and extent of the perturbed region are comparable. As for J10, the carrier $U_x$ perturbation becomes less cell-localised and more spatially regularised as $b$ increases, consistent with the conservative redistribution of the interphase momentum source—a different spatial distribution of the same exchanged momentum, not a weaker global coupling. Negative $U_x$ colour-bar values are read only as a near-zero or locally reversed velocity near the jet root, not an extended recirculation region, since no dedicated field-level recirculation check was performed.



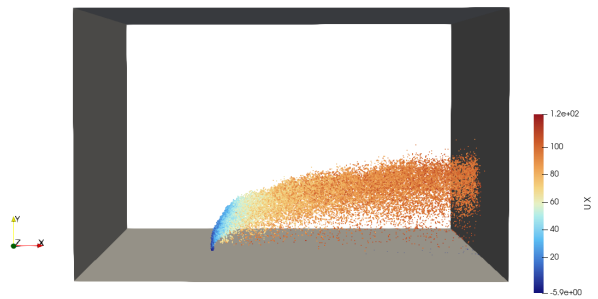
**4.13a** *PCM ($b = 0$): parcel $d$*



**4.13b** *PCM ($b = 0$): parcel $U_x$*

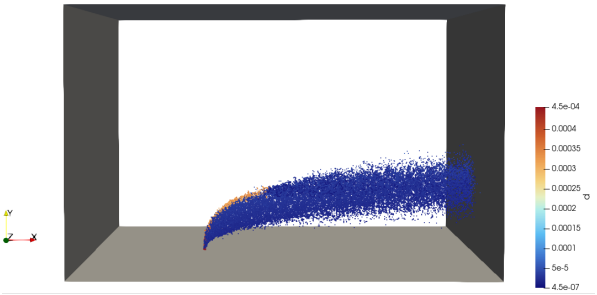


**4.13c** *$b = D_{\text{inj}}$: parcel $d$*

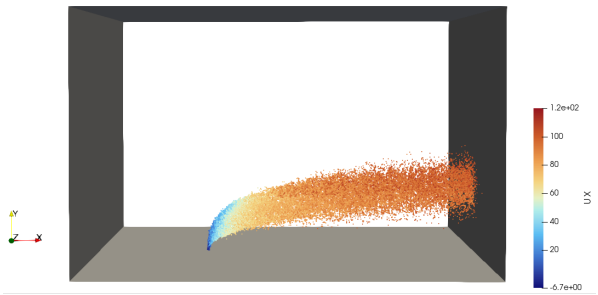


**4.13d** *$b = D_{\text{inj}}$: parcel $U_x$*

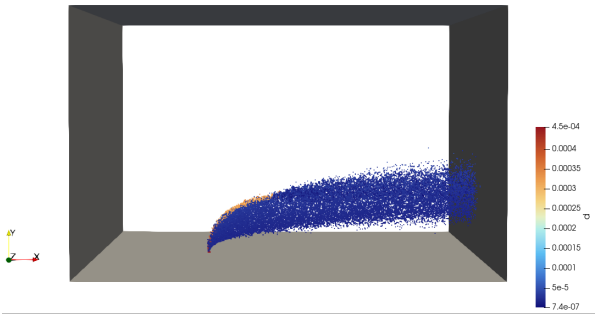
**Figure 4.13:** *J20 spray fields: parcel diameter $d$ (in m) and parcel streamwise velocity $U_x$ (in m/s), PCM and $b = D_{\text{inj}}$.*



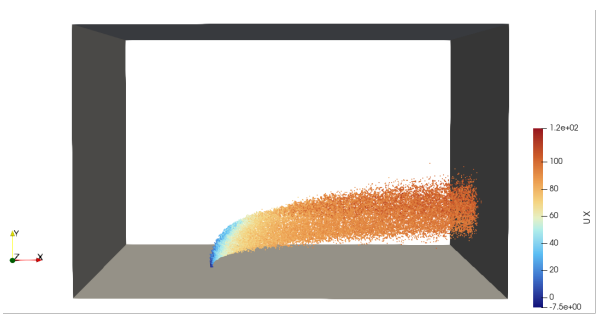
**4.14a** *$b = 3D_{\text{inj}}$: parcel $d$*



**4.14b** *$b = 3D_{\text{inj}}$: parcel $U_x$*

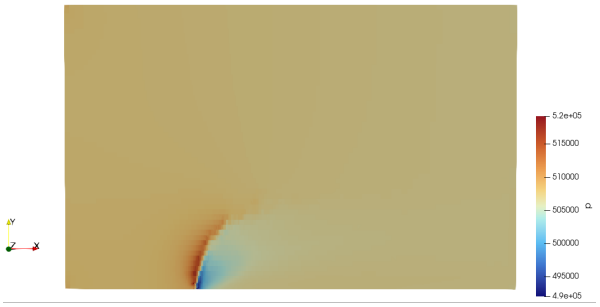


**4.14c** *$b = 6D_{\text{inj}}$: parcel $d$*

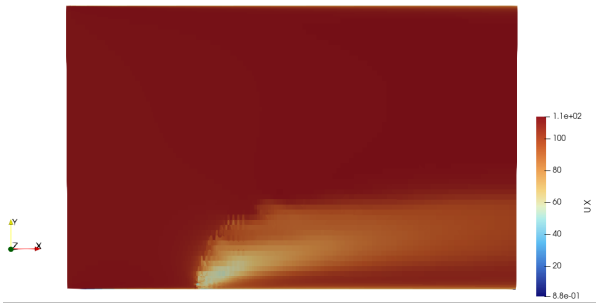


**4.14d** *$b = 6D_{\text{inj}}$: parcel $U_x$*

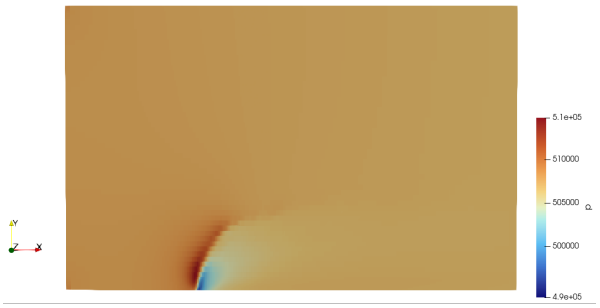
**Figure 4.14:** *J20 spray fields: parcel diameter $d$ (in m) and parcel streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom).*



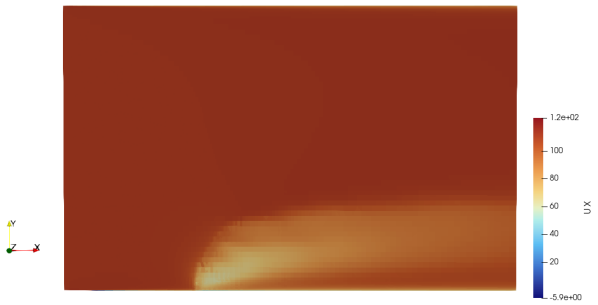

**4.15a** *PCM ($b = 0$): carrier $p$*




**4.15b** *PCM ($b = 0$): carrier $U_x$*

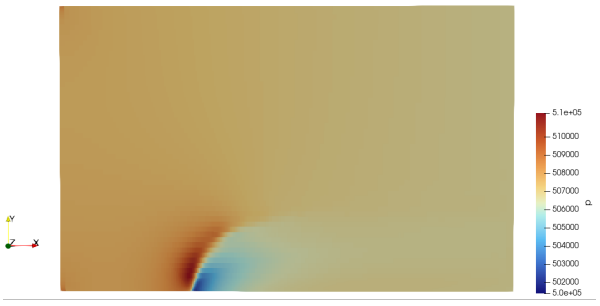



**4.15c** *$b = D_{\text{inj}}$: carrier $p$*




**4.15d** *$b = D_{\text{inj}}$: carrier $U_x$*

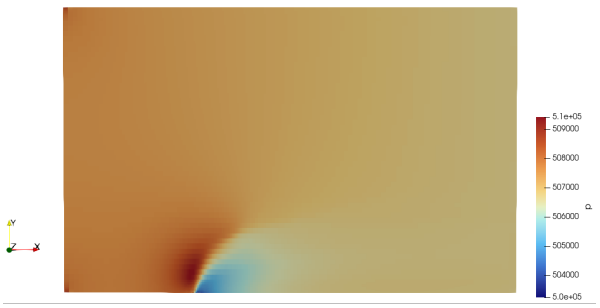
**Figure 4.15:** *J20 carrier fields on the $z = 0$ plane: pressure $p$ (in Pa) and carrier streamwise velocity $U_x$ (in m/s), PCM and $b = D_{\text{inj}}$.*




**4.16a** *$b = 3D_{\text{inj}}$: carrier $p$*




**4.16b** *$b = 3D_{\text{inj}}$: carrier $U_x$*




**4.16c** *$b = 6D_{\text{inj}}$: carrier $p$*




**4.16d** *$b = 6D_{\text{inj}}$: carrier $U_x$*

**Figure 4.16:** *J20 carrier fields on the $z = 0$ plane: pressure $p$ (in Pa) and carrier streamwise velocity $U_x$ (in m/s), $b = 3D_{\text{inj}}$ (top) and $b = 6D_{\text{inj}}$ (bottom).*

## 4.8 Summary of the Lambert Validation

The Lambert LJICF validation campaigns show that the Madabhushi–Lambert primary breakup model, as implemented in the present OpenFOAM solver, reproduces the qualitative structure of the spray at both measurement planes and for both momentum flux ratios, including the increase in wall-normal penetration from J10 to J20. The PCM baseline ($b = 0$) and $b = D_{\mathrm{inj}}$ retain the largest errors, mainly on $D_{32}$, with near-plane $D_{32}$ errors ranging approximately from 24% to 30% across the two campaigns (Tables 4.2 and 4.3). The two larger bandwidths $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ reduce the droplet-size error significantly, most clearly at the far plane, but the effect of the bandwidth is neither monotonic in $b$ nor uniform across planes and quantities, and the reductions are accompanied by a compression of the spray cross-section and a thinning of its near-wall base, as shown by the colour maps and discussed in Sections 4.3 and 4.5.

At the near plane the simulated $D_{32}$ shows, in J10, a localised elevation in the $y \approx 6$–$9$ mm band and, in J20, a more disordered crossing pattern shifted to $y \approx 9$–$13$ mm. These near-field irregularities are compatible with differences in the position of the sampling planes relative to the breakup region, in the local relative velocity, in the deformation and fragmentation history of the parcels, and in the representation of the turbulence and column wake in the present SST-RAS framework. They cannot be attributed with certainty to a single breakup-timing mechanism from the one-dimensional profiles alone.

The $U_x$ profiles are generally less sensitive to the bandwidth than $D_{32}$ in both J10 and J20. They are not, however, independent of it: the parcel streamwise velocity still reflects the redistribution of the momentum source term, the carrier wake and the parcel acceleration history. The local velocity minima seen in the SBES reference are not reproduced with the same intensity by the SST-RAS simulations, a difference that is consistent across all bandwidth values but cannot be ascribed to a single cause from the available profiles.

On balance, $b = 3D_{\mathrm{inj}}$ is retained as the representative bandwidth for both momentum flux ratios. It gives the lowest near-field NRMSE for both $D_{32}$ and $U_x$ in J10, the lowest $U_x$ error at both planes in J20, and remains within a few percentage points of the far-field $D_{32}$ minimum, while contracting the spray structure less than $b = 6D_{\mathrm{inj}}$; it therefore offers the best compromise between accuracy, statistical coverage and preservation of the spatial structure of the spray, and is not claimed to be a universal optimum. A robustness check that recomputes the NRMSE over the experimental points common to the valid span of all four configurations (Table 4.4) does not overturn the interpretation above and indicates that the improved performance of the larger bandwidths is not merely an artefact of reduced coverage.

| Case | NRMSE $D_{32}$ $x/d = 30$ | NRMSE $D_{32}$ $x/d = 60$ | NRMSE $U_x$ $x/d = 30$ | NRMSE $U_x$ $x/d = 60$ |
|---|---|---|---|---|
| *J10* | | | | |
| PCM ($b = 0$) | 27.6 | 25.8 | 14.5 | 8.9 |
| DBM $b = D_{\mathrm{inj}}$ | 27.2 | 24.0 | 17.7 | 9.5 |
| DBM $b = 3D_{\mathrm{inj}}$ | 19.1 | 11.5 | 13.9 | 6.9 |
| DBM $b = 6D_{\mathrm{inj}}$ | 22.3 | 10.1 | 14.5 | 6.3 |
| *J20* | | | | |
| PCM ($b = 0$) | 24.6 | 28.7 | 15.0 | 7.2 |
| DBM $b = D_{\mathrm{inj}}$ | 19.8 | 20.9 | 16.8 | 7.9 |
| DBM $b = 3D_{\mathrm{inj}}$ | 11.4 | 9.9 | 14.5 | 5.2 |
| DBM $b = 6D_{\mathrm{inj}}$ | 12.1 | 7.5 | 17.6 | 6.2 |

**Table 4.4:** *Robustness check: NRMSE (in %) for $D_{32}$ and $U_x$ recomputed over the experimental points common to the valid span of all four configurations at each plane, for J10 and J20.*

The residual high-$y$ $D_{32}$ overestimation persists, to varying degrees, for every bandwidth, which indicates that the DBM momentum-source diffusion alone cannot eliminate it; turbulence and wake representation, breakup history, parcel trajectories and statistical sampling may all contribute. The recourse to Sun and Xiao [10] is limited to the general principle of avoiding both spatially concentrated sources and excessive over-smoothing; the specific value $b = 3D_{\mathrm{inj}}$ is not prescribed by that reference. The transferability of this choice and of the overall modelling strategy to the Zhang LJIGF/GLOP configuration—which involves a different fluid, a different injector geometry, and a higher range of gas-film Weber numbers—is examined in Chapter 5.

# 5 Zhang LJIGF/GLOP Validation and DBM Necessity

## 5.1 Motivation: Differences Between Lambert LJICF and Zhang LJIGF

The Lambert benchmark of Chapter 4 validated the Madabhushi–Lambert model for a liquid jet in an *open* gaseous crossflow (LJICF): a slender Jet-A1 column ($D_{\mathrm{inj}} = 0.457$ mm) in an unbounded air stream. The Zhang configuration is a liquid jet in a *confined* gas film (LJIGF), realised in a gas–liquid orifice-plate (GLOP) atomiser: a water jet ($D_{\mathrm{inj}} = 1.6$ mm) injected normal to a gas stream bounded between two walls [3], closer to the pintle–ogive injectors of throttleable liquid-propellant engines than to Lambert's open crossflow. Three differences make this configuration a more demanding stress test of the DBM source coupling.

The first is geometric: the Zhang orifice is 3.5× wider (1.6 mm versus 0.457 mm), so the near-injector momentum source is more concentrated and, combined with the confinement, represents a larger fraction of the local carrier momentum; the two-way coupling is correspondingly stiffer. The second difference concerns both the working liquid and the drag calibration: the Zhang configuration uses water rather than Jet-A1, and the blob drag coefficient is not simply inherited from the Lambert case but recalibrated for the present case study to $C_{D,\mathrm{blob}} = 2.16$ (Section 2.7), a value larger than the $C_D = 1.696$ reported by Wu et al. [2] for the column-like regime, while the disc drag coefficient is kept at $C_{D,\mathrm{disc}} = 1.2$.

The third difference is numerical, and is the central theme of this chapter. The stiffer coupling makes the bare per-cell source projection unstable: in the finely resolved near-injector cells the concentrated parcel sources can drain the gas momentum and drive the local density towards zero, so that the unsmoothed configuration either fails outright with a floating-point exception or yields a badly distorted penetration. Where in the open Lambert crossflow the DBM was a beneficial refinement, in the confined GLOP geometry it becomes a *necessary* ingredient, both for numerical survival and for a physically meaningful spray. The interplay between the carrier-to-parcel interpolation scheme and the DBM smoothing that together make the cases runnable is examined in detail in Section 5.3.

The validation focuses on the intermediate base case $We_g^{\mathrm{film}} \approx 500$ (Zhang Sim. 2, throttle level 30%) and the higher-$We_g^{\mathrm{film}}$ follow-up $We_g^{\mathrm{film}} \approx 3500$ (Sim. 3, 70%). The

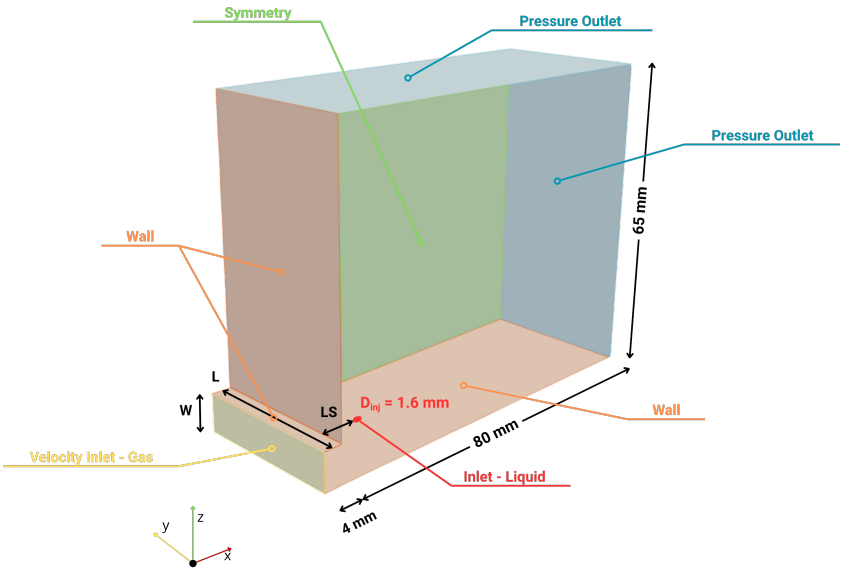



**Figure 5.1:** *Zhang LJIGF/GLOP computational domain and boundary conditions.*

most energetic case $We_g^{\text{film}} \approx 8500$ (Sim. 4, 100%) did not run to completion; the instability was not isolated, and the absence of DBM smoothing for the energy and mass coupling terms—which, unlike the momentum source, are left at their per-cell values (Section 5.3)—is one of several possible contributors discussed as future work in Section 6.3. The governing equations are given in Chapter 2 and the implementation in Chapter 3.

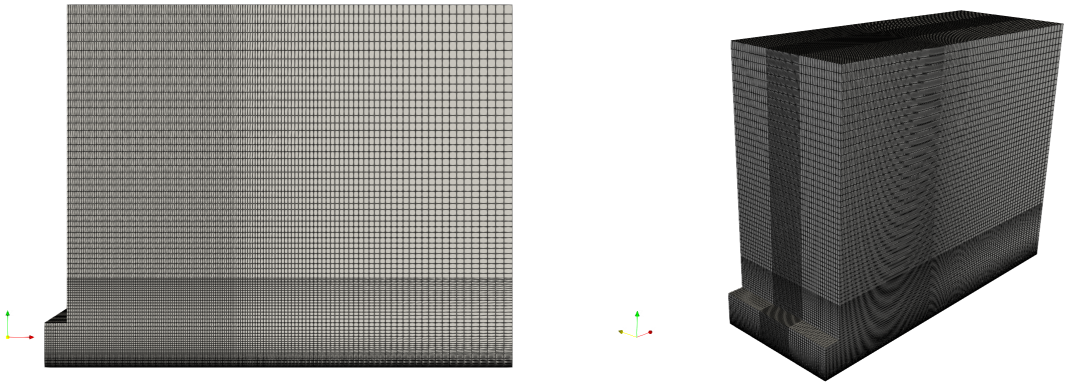
## 5.2 Numerical Setup and Reference Quantities

The structured hexahedral mesh reproduces the confined gas-film geometry of Zhang et al. [3] (Figure 5.1); the operating, geometric, mesh and dispersed-phase model parameters for both gas-film Weber numbers are collected in Table 5.1, and Figure 5.2 shows the grid. The gas enters through a confined inlet channel bounded between the bottom wall and a film top wall, and downstream of a backward-facing step the upper domain opens to its full height. The liquid orifice is centred on the bottom wall and injects normal to the gas film (in $+z$), with the streamwise $x$, wall-normal penetration $z$ and spanwise $y$ directions oriented as in Figure 5.1 and the spanwise extent closed by two symmetry planes. The mesh is strongly graded towards the bottom wall and the orifice to resolve the near-injector shear layer and coarsens towards the outlets.

The two cases differ primarily in the gas-film Weber number, the gas-film velocity and the liquid injection velocity, whereas the injector diameter, the liquid properties, the breakup-model coefficients and the drag coefficients are kept unchanged. The carrier is nitrogen under atmospheric conditions, with a uniform velocity imposed at the gas inlet, no-slip walls and pressure outlets; water is injected as blobs normal to the gas film, and the dispersed phase uses the Madabhushi–Lambert breakup model (Chapter 3), which handles KH/Wave surface stripping and Pilch–Erdman secondary breakup, with inter-parcel collisions and stochastic turbulent dispersion. The DBM bandwidth is swept from the bare PCM baseline ($b = 0$) up to $b = 6D_{\text{inj}}$, and the carrier-to-parcel

| Parameter | Symbol | $We_g^{\text{film}} \approx 500$ | $We_g^{\text{film}} \approx 3500$ | Unit |
|---|---|---|---|---|
| *Case-specific* | | | | |
| Zhang simulation (throttle) | – | Sim. 2 (30%) | Sim. 3 (70%) | – |
| Gas-film Weber number | $We_g^{\text{film}}$ | 508 | 3459 | – |
| Gas-film velocity | $U_g$ | 62.5 | 163.1 | m/s |
| Injection velocity | $U_l$ | 5.32 | 14.2 | m/s |
| *Domain and mesh* | | | | |
| Domain $L_x \times L_z \times L_y$ | – | $84 \times 65 \times 32$ | | $\text{mm}^3$ |
| Gas-film inlet height | – | 8 | | mm |
| Orifice position | $x$ | 15 | | mm |
| Mesh size | – | $\approx 1.3 \times 10^6$ | | cells |
| Near-wall spacing | $\Delta z_1$ | $\approx 4$ | | $\mu$m |
| *Carrier phase* | | | | |
| Turbulence model | – | $k$–$\omega$ SST (RAS) | | – |
| Carrier gas | – | $N_2$ (Sutherland) | | – |
| Pressure | $p$ | 101 325 | | Pa |
| Temperature | $T$ | 288 | | K |
| Gas density | $\rho_g$ | $\approx 1.19$ | | $\text{kg/m}^3$ |
| *Dispersed-phase and model parameters* | | | | |
| Working liquid | – | water | | – |
| Injector diameter | $D_{\text{inj}}$ | 1.6 | | mm |
| Liquid density | $\rho_l$ | $\approx 1000$ | | $\text{kg/m}^3$ |
| Column-breakup coefficient | $C_0$ | 0.46 | | – |
| KH child-size coefficient | $b_0$ | 2.44 | | – |
| KH breakup-time coefficient | $b_1$ | 1.73 | | – |
| Ligament correction factor | $F_{\text{Lig}}$ | 1 | | – |
| Blob drag coefficient | $C_{D,\text{blob}}$ | 2.16 | | – |
| Disc drag coefficient | $C_{D,\text{disc}}$ | 1.2 | | – |
| *Numerical method* | | | | |
| Parcel interpolation (PCM) | – | `cellPoint` | | – |
| Parcel interpolation (DBM) | – | `cell` | | – |
| DBM bandwidths | $b$ | 0, 1.6, 4.8, 9.6 | | mm |

**Table 5.1:** *Zhang LJIGF/GLOP numerical setup and dispersed-phase model parameters.*



**5.2a** *Front view (x–z plane)* **5.2b** *Isometric view*

**Figure 5.2:** *Zhang GLOP mesh.*

interpolation scheme differs between PCM and DBM for the stability reasons discussed in Section 5.3.

Post-processing mirrors the sectional binning of the Lambert analysis (Section 4.2), applied here to instantaneous parcel samples rather than to crossing events. Parcels are gathered from the saved Lagrangian fields over $t \in [0.02, 0.03]$ s in a slab of half-thickness 0.5 mm centred on the sampling plane $\Delta x = 60$ mm from the orifice ($x = 75$ mm) and binned along the penetration coordinate $z$ into 13 equal bins spanning 0–65 mm. Each bin is weighted by the parcel `nParticle` number, so the *equivalent droplet count* is the sum of `nParticle` over the parcels in the bin; bins whose equivalent droplet count falls below 25 are discarded. The reported diameters are the `nParticle`-weighted arithmetic mean $D_{10} = \sum n_p d / \sum n_p$ and Sauter mean $D_{32} = \sum n_p d^3 / \sum n_p d^2$, and $U_{\text{mean}}$ is the `nParticle`-weighted mean parcel speed $|U_p|$ (velocity magnitude, not a single component), consistent with the $|U_p|$ maps of Section 5.5. The reference profiles are the VOF-to-DPM/LES results of Zhang et al. [3] (Fig. 24), not direct experimental measurements.

## 5.3 Carrier-to-Parcel Interpolation, PCM Stability and DBM Necessity

The confined GLOP geometry exposes a stability limit of the bare per-cell source projection that never appeared in the open Lambert crossflow. Two distinct numerical choices must be kept separate here: the carrier-to-parcel *interpolation scheme* (`cell` or `cellPoint`), which sets how the carrier fields are sampled at the parcel position, and the parcel-to-carrier *source projection* (PCM or DBM), which sets how the parcel reaction is returned to the gas. The bare PCM projection deposits the entire momentum source of a parcel into the single cell that contains its centroid, regardless of the interpolation scheme. In the strongly refined near-orifice cells, where many parcels are concentrated in the injection region, the unsmoothed ($b = 0$) deposition produces an artificial momentum drain that can empty a gas cell of its momentum and drive the local density towards zero; the resulting floating-point overflow (SIGFPE) aborts the run before a statistically meaningful state is reached. In Lambert the open crossflow kept the local carrier momentum large compared with the parcel sources, so the `cell` + PCM

combination ran without difficulty; here the larger orifice and the confined gas film make the same combination prohibitively stiff.

To obtain a PCM baseline at all, the $b = 0$ case is therefore run with the `cellPoint` interpolation scheme. This does not redistribute the parcel source term—the deposition remains per-cell—but it lets each parcel sample a vertex-interpolated, smoother carrier velocity rather than the raw host-cell value, which weakens the locking feedback enough for the unsmoothed case to run to completion. The three DBM cases ($b > 0$) instead retain the standard `cell` interpolation used throughout Lambert and rely on the screened-diffusion smoothing operator of Sun and Xiao [10] to spread each deposited momentum source over a neighbourhood of cells, so the per-cell peak is regularised by the smoothing itself and the deposition is stable and consistent across all three bandwidths. The bare-PCM curve in the comparisons that follow is thus a `cellPoint`, $b = 0$ result, whereas the DBM curves use `cell` interpolation with screened-diffusion source smoothing. The PCM-versus-DBM comparison therefore varies both the bandwidth and the interpolation scheme; the three $b > 0$ cases ($D_{\mathrm{inj}}$, $3D_{\mathrm{inj}}$, $6D_{\mathrm{inj}}$), which share the `cell` scheme, isolate the effect of the bandwidth alone.

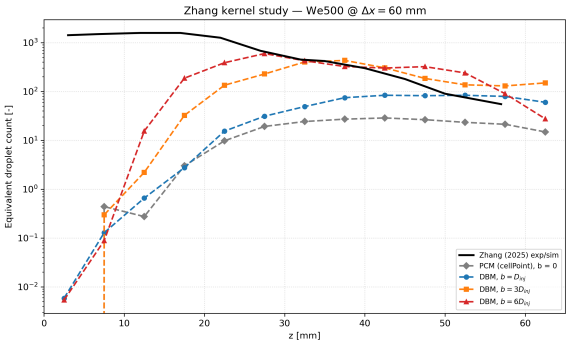
The `cellPoint` stabilisation is, however, only a numerical expedient: it keeps the run alive but does not make the bare PCM physically adequate. With $b = 0$ the source peak still lives on the grid scale, so the near-injector parcels sample a carrier velocity distorted by the very momentum they have just deposited; this feedback biases the parcel trajectories and yields a jet penetration that does not reflect the physics of the confined jet, independently of any error in the predicted drop sizes (Section 5.4). These two observations—the `cell` singularity at $b = 0$ and the residual grid-scale bias of the stabilised `cellPoint` PCM—establish the DBM as a *necessary* ingredient for the GLOP geometry rather than the optional refinement it was in Lambert.

## 5.4 DBM Bandwidth Sensitivity at Intermediate Gas-film Weber Number ($We_g^{\mathrm{film}} \approx 500$)

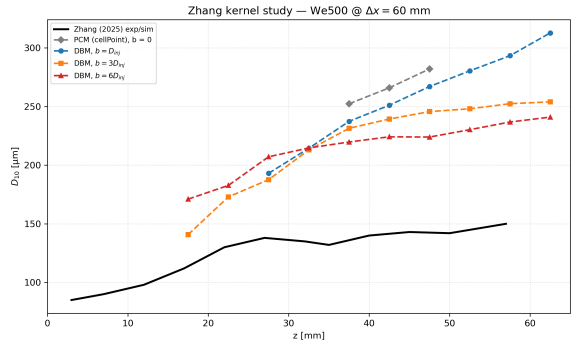
Figure 5.3 compares the four DBM bandwidth cases against the Zhang reference at the intermediate base condition $We_g^{\mathrm{film}} \approx 500$, for the equivalent droplet count, the arithmetic and Sauter mean diameters $D_{10}$ and $D_{32}$, and the mean droplet velocity $U_{\mathrm{mean}}$, all sampled along the penetration coordinate $z$ at $\Delta x = 60$ mm.

The bare PCM case (Figure 5.3(a), grey) under-populates the spray by one to two orders of magnitude and concentrates the few surviving droplets at large penetration ($z \gtrsim 35$ mm), far from the near-base peak of the Zhang reference ($z \lesssim 25$ mm); the near-base bins fall below the 25 equivalent-droplet threshold and are discarded, leaving valid diameter and velocity data only at high $z$. This is a spatial-distribution failure, not merely a size error: without smoothing, the near-injector parcels follow unphysical trajectories, consistent with the regularising role of the DBM anticipated in Section 5.3.

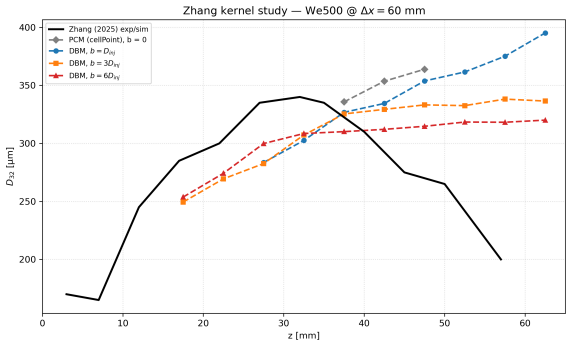
Increasing $b$ raises the equivalent droplet count and shifts it towards the jet base, bringing $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ close to the reference over part of the mid-field ($z \approx 30$–$40$ mm); none of the cases, however, reproduces the high near-base population of the Zhang reference at $z \lesssim 20$ mm (Figure 5.3(a)). The Sauter diameter follows the same



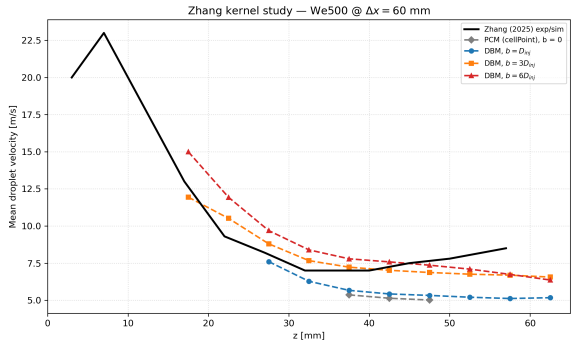

**5.3a** *Equivalent droplet count*




**5.3b** $D_{10}$




**5.3c** $D_{32}$




**5.3d** $U_{\text{mean}}$

**Figure 5.3:** $We_g^{\text{film}} \approx 500$*: DBM bandwidth sensitivity at* $\Delta x = 60$ *mm.*

bandwidth ordering (Figure 5.3(c)): larger bandwidths lower $D_{32}$ towards the reference peak ($\approx 340$ $\mu$m at $z \approx 30$ mm) and flatten the profile, but all cases over-predict $D_{32}$ downstream, where the Zhang VOF-to-DPM/LES reference decreases whereas the OpenFOAM profiles remain nearly flat or continue to increase. This behaviour is consistent with an excessive opening of the spray cone after the first Pilch–Erdman breakup event of the liquid core, which redistributes relatively large fragments over the downstream bins and delays the decay of the Sauter mean diameter. $D_{10}$ (Figure 5.3(b)) is over-predicted by roughly a factor of two by every case; the larger bandwidths reduce the downstream growth but do not remove the discrepancy.

Unlike the weak bandwidth dependence of $U_x$ in Lambert, $U_{\text{mean}}$ here is markedly sensitive to $b$ (Figure 5.3(d)), with ordering $b = 6D_{\text{inj}} > 3D_{\text{inj}} > D_{\text{inj}} > \text{PCM}$. The trend is consistent with the larger smoothing bandwidth redistributing the interphase momentum source over a broader neighbourhood, thereby reducing the cell-localised carrier response sampled by the parcels near the jet base, so that their deceleration is less abrupt; this interpretation should be regarded as a coupling-level explanation rather than a direct measurement of the local source field.

Two discrepancies persist even at the largest bandwidths: the $D_{10}/D_{32}$ over-prediction and the missing downstream $D_{32}$ roll-off. Both are more plausibly attributed to physical-model limitations (secondary atomisation, blob-drag calibration, RANS closure) than to the source projection, whose regularising effect is indicated by the improvement of the spatial count distribution with $b$.

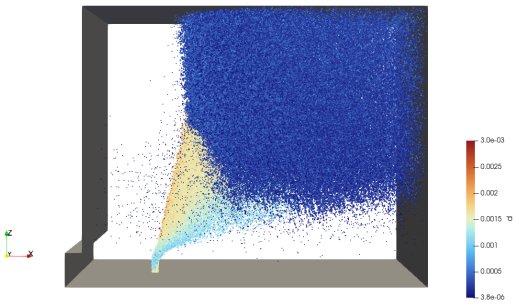
## 5.5 Spray and Carrier-Field Visualisation at $We_g^{\text{film}} \approx 500$

The field visualisations on the $y = 0$ plane complement the sectional profiles and are read qualitatively; the colour scales are not necessarily identical across panels, so the maps are compared by spatial structure rather than by colour value. The spray maps (Figures 5.4 and 5.5) show the spatial counterpart of the count redistribution of Section 5.4: bare PCM produces a diffuse cloud of fine droplets (deep blue) that fills the full 65 mm chamber height with no coherent jet structure—the over-dispersion of Section 5.3, the few surviving parcels being scattered towards the upper wall rather than concentrated near the base. Increasing $b$ progressively collapses this cloud onto an arched bent jet; by $b = 6D_{\text{inj}}$ the spray follows a well-defined trajectory with a compact near-orifice column of large blobs (warm colours, close to the injection diameter $D_0 = D_{\text{inj}} = 1.6$ mm) bending over before dispersing into a focused fine-droplet plume.

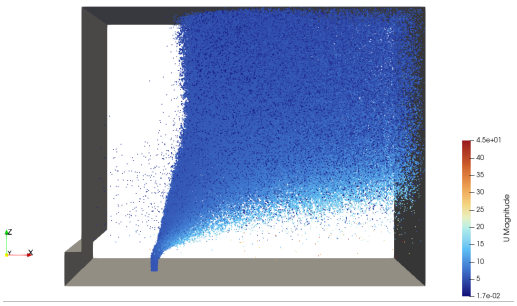
The diameter maps are qualitative snapshots of where larger and smaller parcels are located rather than direct evidence of individual breakup events: the largest parcels (close to the injection diameter) appear concentrated near the liquid column at the orifice, while finer parcels—consistent with the KH/Wave stripping of Chapter 3—are carried over the crest and downstream, so the parcel size tends to decrease along the trajectory. The companion parcel velocity-magnitude $|U_p|$ snapshots are consistent with the expected momentum exchange: parcels appear slow near the orifice (close to the liquid speed $U_l \approx 5.3$ m/s, deep blue) and faster downstream as they are entrained towards the gas-film speed ($\approx 62$ m/s), the warm downstream plume suggesting where the light children have nearly equilibrated with the carrier. These snapshots locate regions containing larger or faster parcels but do not by themselves identify the atomisation sequence.

The carrier fields (Figures 5.6 and 5.7) show the confined gas film as a high-speed band ($\approx 62$ m/s) hugging the bottom wall, with an oblique low-velocity streak rising from the orifice and leaning downstream—the momentum deficit imprinted by the liquid column, whose inclination follows the bent-jet trajectory as the gas surrenders momentum to the accelerating droplets. The pressure field shows a windward overpressure / leeward suction dipole at the jet base, the classic bluff-body signature of the column obstructing the film, superimposed on an otherwise uniform $\approx 1$ atm field. With bare PCM both gradients are confined to one or two cells—the near-singular source concentration that makes the ce11 projection diverge in Section 5.3—whereas a finite bandwidth spreads the accumulated momentum source over a screened-diffusion neighbourhood of Gaussian-equivalent bandwidth $b$ (Equation (2.73)), so that by $b = 6D_{\text{inj}}$ the stagnation dipole becomes a broad, mild depression and the deficit streak widens into a gentle wake. As the bandwidth increases, the carrier velocity-magnitude ($|U|$) perturbation appears less cell-localised and more spatially regularised, consistent with the conservative redistribution of the interphase momentum source—a different spatial distribution of the same exchanged momentum, not a weaker global coupling.

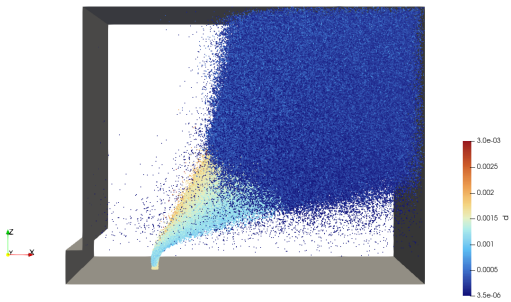
A second feature is worth emphasising: the carrier *topology*—film, deficit streak and pressure dipole—is qualitatively preserved across all four bandwidths, in contrast to the spray morphology, which changes from a domain-filling cloud to a focused jet. The smoothing operator does not create or remove flow features but only sets the smoothness with which the strong near-injector two-way coupling is resolved.



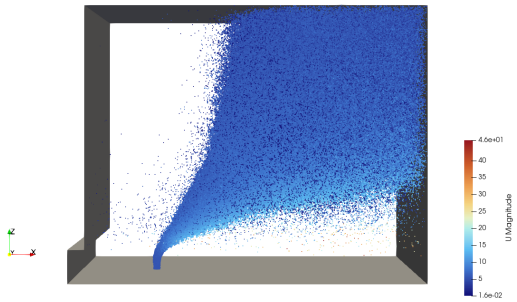
**5.4a** *PCM ($b = 0$): parcel $d$*



**5.4b** *PCM ($b = 0$): parcel $|U_p|$*

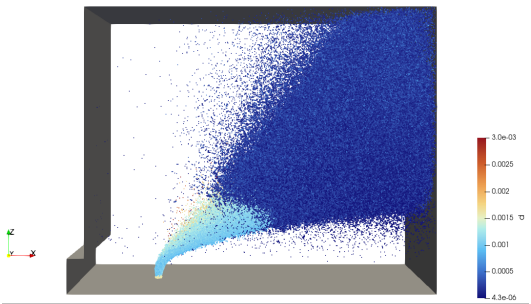


**5.4c** $b = D_{\text{inj}}$*: parcel* $d$

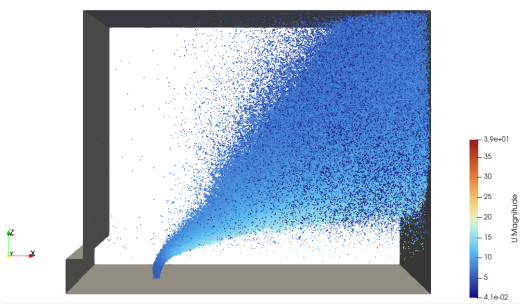


**5.4d** $b = D_{\text{inj}}$*: parcel* $|U_p|$

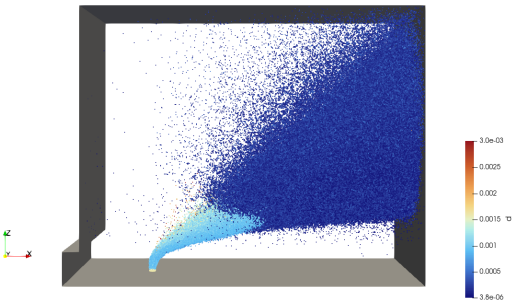
**Figure 5.4:** $We_g^{\text{film}} \approx 500$ *spray fields: parcel* $d$ *and* $|U_p|$*, PCM (top) and* $b = D_{\text{inj}}$ *(bottom).*



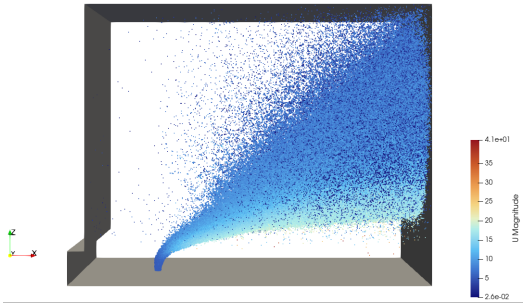
**5.5a** $b = 3D_{\text{inj}}$*: parcel* $d$



**5.5b** $b = 3D_{\text{inj}}$*: parcel* $|U_p|$



**5.5c** $b = 6D_{\text{inj}}$*: parcel* $d$

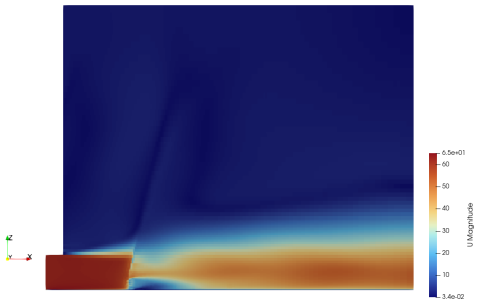


**5.5d** $b = 6D_{\text{inj}}$*: parcel* $|U_p|$

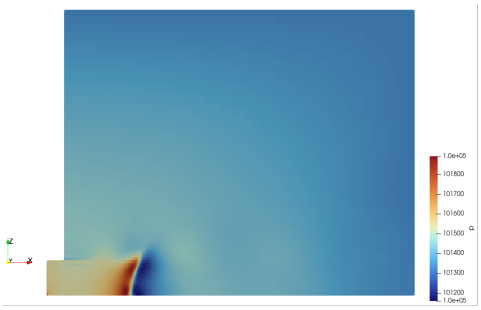
**Figure 5.5:** $We_g^{\text{film}} \approx 500$ *spray fields: parcel* $d$ *and* $|U_p|$*,* $b = 3D_{\text{inj}}$ *(top) and* $b = 6D_{\text{inj}}$ *(bottom).*



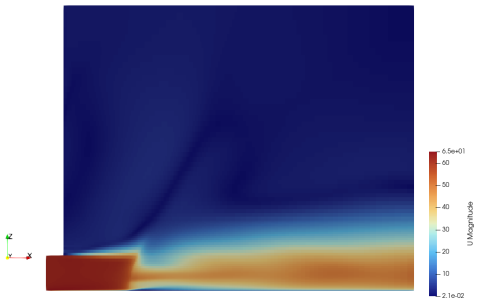
**5.6a** *PCM ($b = 0$): carrier $p$*



**5.6b** *PCM ($b = 0$): carrier $|U|$*

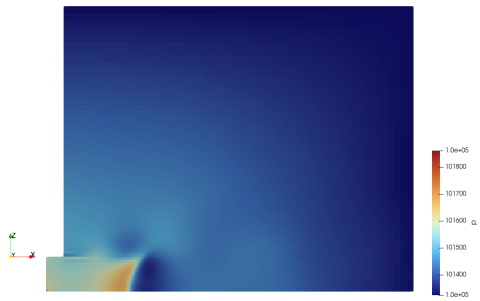


**5.6c** $b = D_{\text{inj}}$*: carrier* $p$

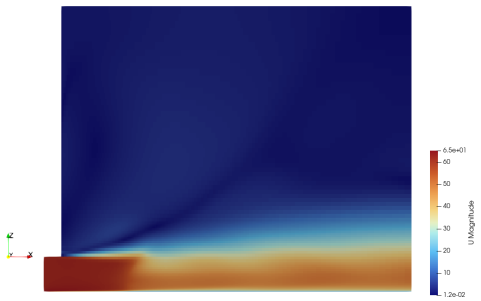


**5.6d** $b = D_{\text{inj}}$*: carrier* $|U|$

**Figure 5.6:** $We_g^{\text{film}} \approx 500$ *carrier fields on the* $y = 0$ *plane: pressure* $p$ *(left) and velocity magnitude* $|U|$ *(right), PCM and* $b = D_{\text{inj}}$.



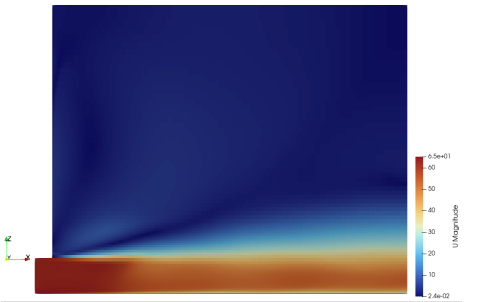
**5.7a** $b = 3D_{\text{inj}}$*: carrier* $p$



**5.7b** $b = 3D_{\text{inj}}$*: carrier* $|U|$



**5.7c** $b = 6D_{\text{inj}}$*: carrier* $p$



**5.7d** $b = 6D_{\text{inj}}$*: carrier* $|U|$

**Figure 5.7:** $We_g^{\text{film}} \approx 500$ *carrier fields on the* $y = 0$ *plane:* $p$ *and* $|U|$, $b = 3D_{\text{inj}}$ *(top) and* $b = 6D_{\text{inj}}$ *(bottom).*

## 5.6 Selected DBM Bandwidth and Physical Interpretation

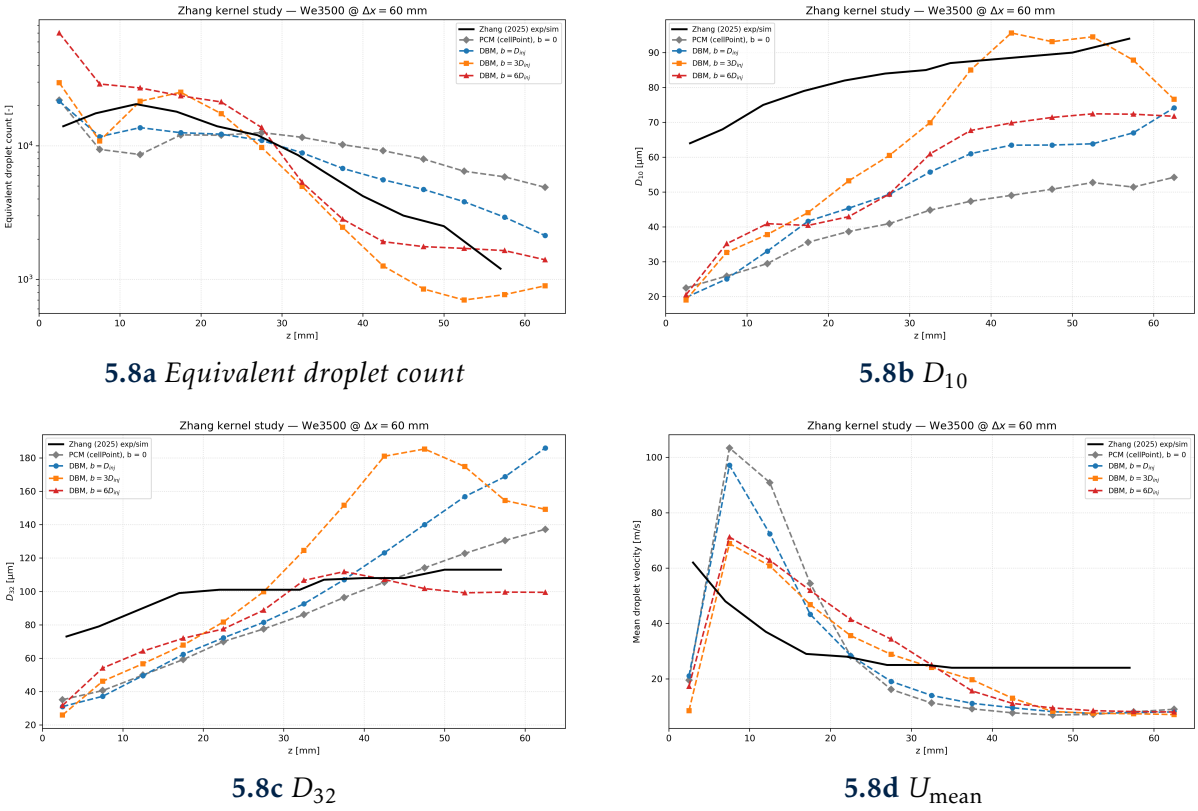
On the basis of Figure 5.3 and the field visualisations of Section 5.5, $b = 3D_{\text{inj}}$ is retained as the representative bandwidth for the Zhang cases at $We_g^{\text{film}} \approx 500$. It recovers most of the count redistribution towards the jet base together with the $D_{32}$ and $U_{\text{mean}}$ improvement, while applying a less aggressive smoothing than $b = 6D_{\text{inj}}$, which edges it only marginally on velocity and $D_{10}$ at the price of a more aggressive smoothing of the coupled near-injector region. This choice is consistent with the representative bandwidth $b = 3D_{\text{inj}}$ adopted for the Lambert J10 case (Section 4.5): a smoothing bandwidth of the order of a few injector diameters provides a reasonable compromise for both injector geometries, even though the larger Zhang orifice and the confined gas film concentrate the Lagrangian sources more strongly and thus correspond to a wider absolute bandwidth. This is a pragmatic engineering compromise rather than a universal optimum: as the higher-Weber follow-up below shows, the preferred bandwidth depends on the gas-film Weber number, on the metric considered and on the region of the profile, and a larger $b$ can be better for some metrics.

## 5.7 Higher Gas-film Weber Number: $We_g^{\text{film}} \approx 3500$

The same methodology (Section 5.4) is applied to the higher-$We_g^{\text{film}}$ follow-up, Zhang Sim. 3 at $We_g^{\text{film}} \approx 3500$ (throttle level 70%). Figure 5.8 reports the four profiles at $\Delta x = 60$ mm.

The higher Weber number produces a finer, more populous spray: the reference Sauter diameter drops to a plateau of $\approx 100$ $\mu$m (vs. $\approx 340$ $\mu$m at $We_g^{\text{film}} \approx 500$, consistent with SMD $\propto (We_g^{\text{film}})^{-0.6}$) and counts rise to $\mathcal{O}(10^4)$. The near-base population deficit that dominated the bare-PCM case at $We_g^{\text{film}} \approx 500$ largely disappears, and near the base all four cases reproduce the count within a factor of two to three (Figure 5.8(a)). Downstream the bandwidth controls the decay rate: bare PCM keeps too high a population in the far field, $b = D_{\text{inj}}$ reduces this only partly, and the larger bandwidths decay faster; $b = 3D_{\text{inj}}$ decays the fastest and tends to under-predict the far-field count, whereas $b = 6D_{\text{inj}}$ stays closer to the reference there.

The bandwidth ordering is broadly consistent with the intermediate case, but the size metrics now favour a larger bandwidth. For $D_{32}$ (Figure 5.8(c)), $b = 6D_{\text{inj}}$ is the best-performing tested bandwidth at $We_g^{\text{film}} \approx 3500$: it follows the Zhang reference over a substantial portion of the profile and captures the order of magnitude of the $\approx 100$ $\mu$m Sauter-diameter plateau more convincingly than the other cases. By contrast, $b = 3D_{\text{inj}}$ remains the representative global compromise, but it clearly over-predicts $D_{32}$ in the mid-to-high penetration region, while bare PCM and $b = D_{\text{inj}}$ miss the trend. $D_{10}$ (Figure 5.8(b)) is now *under*-predicted at low-to-mid penetration (opposite to $We_g^{\text{film}} \approx 500$); $b = 3D_{\text{inj}}$ can locally over-shoot at large $z$, while $b = 6D_{\text{inj}}$ is more regular. The preferred bandwidth for the size distribution can therefore exceed the value chosen as a representative compromise at $We_g^{\text{film}} \approx 500$ and in Lambert: in terms




**5.8a** *Equivalent droplet count*

**5.8b** $D_{10}$

**5.8c** $D_{32}$

**5.8d** $U_{\text{mean}}$

**Figure 5.8:** $We_g^{\text{film}} \approx 3500$*: DBM bandwidth sensitivity at* $\Delta x = 60$ *mm. Reference data from Zhang et al. [3], Fig. 24.*

of both far-field count decay and $D_{32}$, $b = 6D_{\text{inj}}$ is the most convincing of the tested bandwidths for this higher-$We_g^{\text{film}}$ condition, even though it is not adopted as a universal optimum.

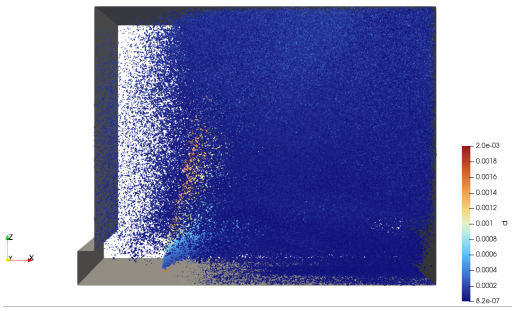
The mean velocity is reproduced substantially worse than at $We_g^{\text{film}} \approx 500$ (Figure 5.8(d)): all cases develop a spurious near-base peak ($\approx 100$ m/s for PCM vs. $\approx 50$ m/s in the Zhang reference) and then fall below the reference far-field plateau ($\approx 24$ m/s). A larger $b$ lowers the spurious peak but does not restore the far-field level, so this metric is best read as a limitation of the present model rather than a confirmation of the bandwidth effect. A plausible contributing factor is that only the momentum source is smoothed by the DBM operator while the mass and energy coupling terms are left at their per-cell values (Section 5.3); the same incomplete regularisation may also contribute to the $We_g^{\text{film}} \approx 8500$ non-completion (Section 6.3), although that mechanism has not been isolated.

## 5.8 Spray and Carrier-Field Visualisation at $We_g^{\text{film}} \approx 3500$

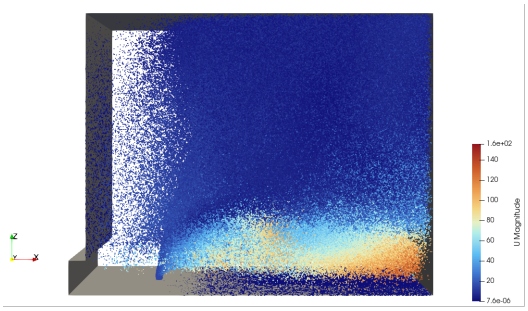
Read qualitatively as at $We_g^{\text{film}} \approx 500$ and with the same colour-scale caveat, the spray maps (Figures 5.9 and 5.10) show the same overall structure but a markedly finer and more deflected spray, driven by the stronger gas momentum ($U_g = 163$ m/s). In the parcel diameter maps the largest parcels remain concentrated in the initial jet/column region near the orifice, while the downstream cloud is dominated by much smaller

parcels; the spatial distribution of the high-$d$ regions changes strongly with $b$, and $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$ produce a parcel layout appreciably different from PCM and $b = D_{\mathrm{inj}}$. The parcel velocity-magnitude maps show that high-$|U_p|$ parcels remain mainly associated with the gas-film-driven lower and downstream part of the cloud, whereas the upper spray contains a large population of slower parcels; increasing $b$ modifies the extent and position of these high-speed regions but does not remove the discrepancy observed in the $U_{\mathrm{mean}}$ profiles (Figure 5.8(d)). These snapshots locate regions containing larger or faster parcels and are consistent with the spatial trends inferred from the profiles, but the parcel state and breakup events are not directly visualised. Even at the larger bandwidths the near-orifice column lies lower and bends more sharply downstream, while the bare-PCM over-dispersion is less catastrophic than at $We_g^{\mathrm{film}} \approx 500$ because the larger parcel population keeps every bin above the statistical threshold (Section 5.7).

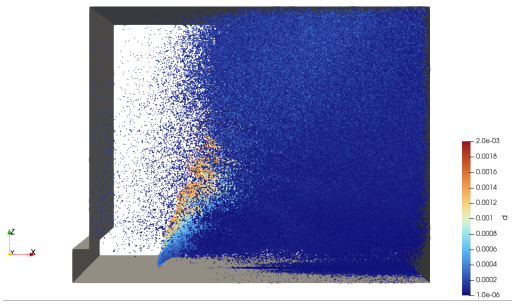
The carrier fields (Figures 5.11 and 5.12) are far more strongly perturbed than in the intermediate-$We_g^{\mathrm{film}}$ case. The $|U|$ map shows the high-speed film ($\approx 163$ m/s) along the bottom wall, a pronounced momentum-deficit region where the jet blocks the gas, and downstream re-acceleration, while the pressure field develops swings an order of magnitude larger than at $We_g^{\mathrm{film}} \approx 500$, reflecting the much larger dynamic pressure $\frac{1}{2}\rho_g U_g^2$. Increasing the smoothing bandwidth still redistributes the coupled momentum exchange, without weakening the global coupling, but the resulting pressure and velocity-magnitude fields are not a simple monotonic smoothing of the PCM solution: the position, extent and topology of the perturbed regions change substantially with $b$, indicating that the local carrier topology is highly sensitive to the source-projection length scale rather than being progressively blurred. The fields should therefore be read qualitatively, as an indicator of this strong sensitivity, and not as evidence of a single mechanism. Because the colour-bar limits are not common to all panels, only the topology, position and apparent extent of the perturbed regions are compared qualitatively. The regularisation also remains incomplete—only the momentum source is smoothed, not the energy or mass coupling terms—which is the field-level counterpart of the spurious velocity peak in Figure 5.8(d).



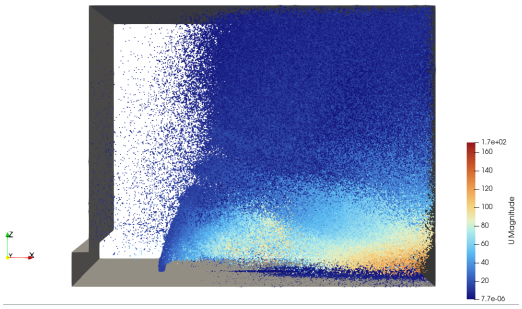
**5.9a** *PCM ($b = 0$): parcel $d$*



**5.9b** *PCM ($b = 0$): parcel $|U_p|$*

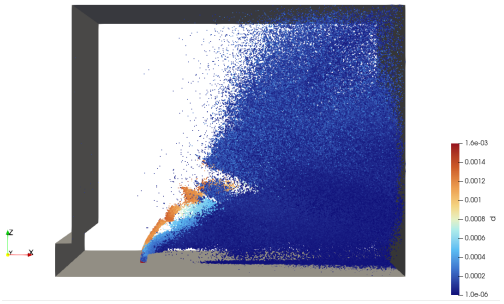


**5.9c** $b = D_{\text{inj}}$*: parcel* $d$

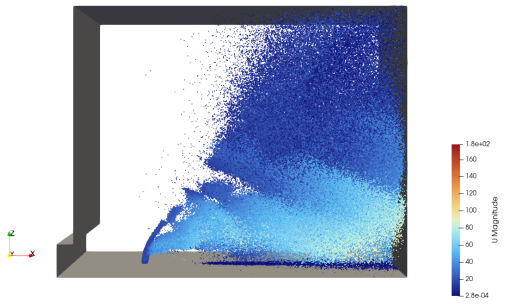


**5.9d** $b = D_{\text{inj}}$*: parcel* $|U_p|$

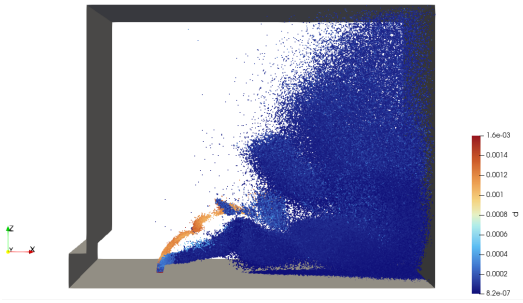
**Figure 5.9:** $We_g^{\text{film}} \approx 3500$ *spray fields on the* $y = 0$ *symmetry plane: parcel diameter* $d$ *(left) and velocity magnitude* $|U_p|$ *(right), PCM and* $b = D_{\text{inj}}$.



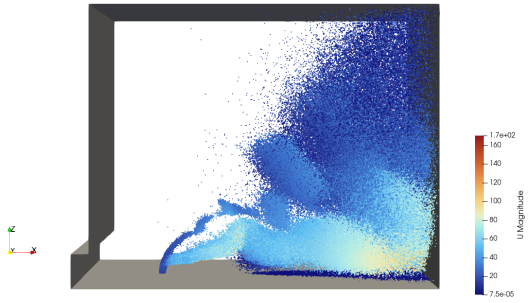
**5.10a** $b = 3D_{\text{inj}}$*: parcel* $d$



**5.10b** $b = 3D_{\text{inj}}$*: parcel* $|U_p|$

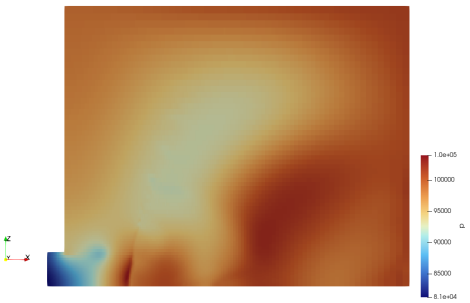


**5.10c** $b = 6D_{\text{inj}}$*: parcel* $d$

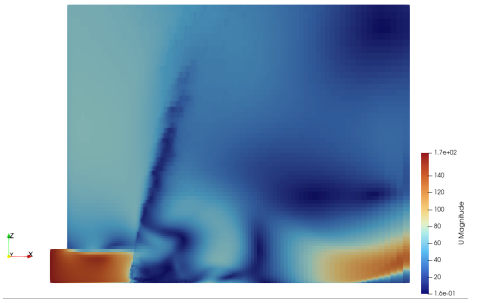


**5.10d** $b = 6D_{\text{inj}}$*: parcel* $|U_p|$

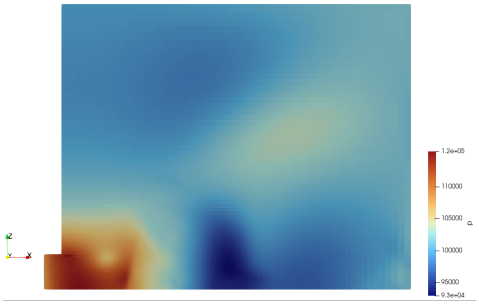
**Figure 5.10:** $We_g^{\text{film}} \approx 3500$ *spray fields: parcel* $d$ *and* $|U_p|$, $b = 3D_{\text{inj}}$ *(top) and* $b = 6D_{\text{inj}}$ *(bottom).*



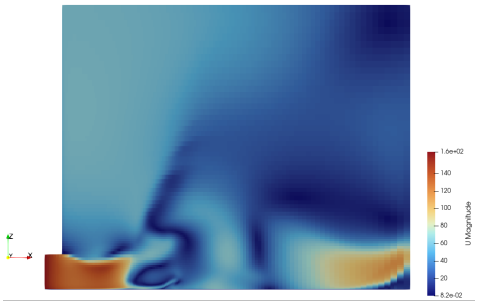
**5.11a** *PCM ($b = 0$): carrier $p$*



**5.11b** *PCM ($b = 0$): carrier $|U|$*

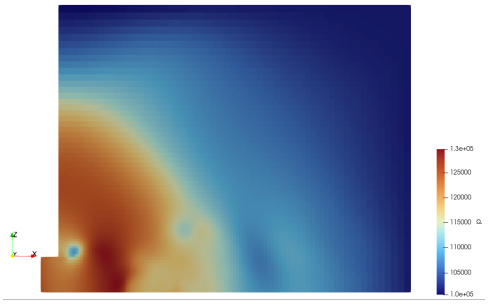


**5.11c** $b = D_{\text{inj}}$*: carrier* $p$

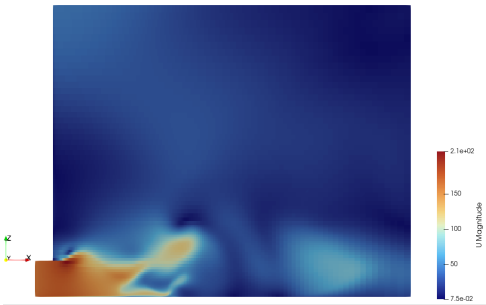


**5.11d** $b = D_{\text{inj}}$*: carrier* $|U|$

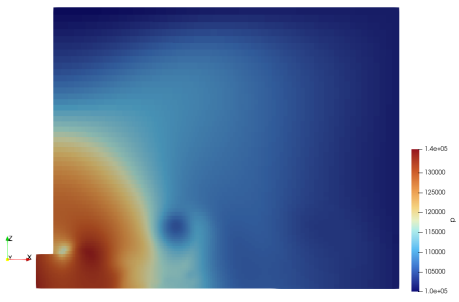
**Figure 5.11:** $We_g^{\text{film}} \approx 3500$ *carrier fields on the* $y = 0$ *plane: pressure* $p$ *(left) and velocity magnitude* $|U|$ *(right), PCM and* $b = D_{\text{inj}}$.



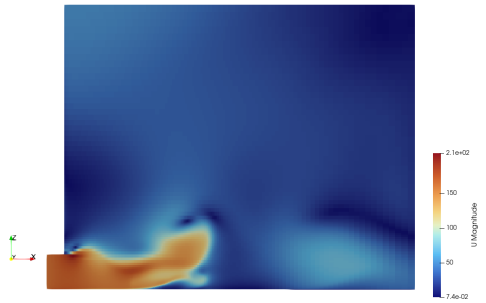
**5.12a** $b = 3D_{\text{inj}}$*: carrier* $p$



**5.12b** $b = 3D_{\text{inj}}$*: carrier* $|U|$



**5.12c** $b = 6D_{\text{inj}}$*: carrier* $p$



**5.12d** $b = 6D_{\text{inj}}$*: carrier* $|U|$

**Figure 5.12:** $We_g^{\text{film}} \approx 3500$ *carrier fields on the* $y = 0$ *plane:* $p$ *and* $|U|$, $b = 3D_{\text{inj}}$ *(top) and* $b = 6D_{\text{inj}}$ *(bottom).*

## 5.9 Summary of LJIGF/GLOP Findings

Across both gas-film conditions the DBM markedly improves agreement with the Zhang reference. At $We_g^{\text{film}} \approx 500$ increasing $b$ shifts the parcel population from a sparse, downstream-displaced cloud towards the near-base concentration of the Zhang reference and brings $D_{32}$ and $U_{\text{mean}}$ substantially closer to it; at $We_g^{\text{film}} \approx 3500$ the finer spray is reproduced within a factor of two to three in count, with $D_{32}$ best captured at the larger bandwidth $b = 6D_{\text{inj}}$. $b = 3D_{\text{inj}}$ is retained as the representative bandwidth, as a pragmatic compromise rather than a universal optimum—the preferred bandwidth depends on the gas-film Weber number, the metric and the region of the profile—and it should be borne in mind that the bare-PCM baseline differs from the DBM cases in the carrier-to-parcel interpolation scheme as well as in $b$ (Section 5.3), so that only the three $b > 0$ cases isolate the bandwidth effect cleanly. Residual discrepancies (diameter over/under-prediction, missing $D_{32}$ roll-off, poor velocity at $We_g^{\text{film}} \approx 3500$) are attributed primarily to physical-model limitations rather than to the source projection.

A particularly relevant positive outcome of the Zhang campaign is the $D_{32}$ prediction at $We_g^{\text{film}} \approx 3500$: despite the remaining discrepancies in count, $D_{10}$ and $U_{\text{mean}}$, the $b = 6D_{\text{inj}}$ case provides an encouraging agreement with the Zhang VOF-to-DPM/LES reference over a substantial part of the profile. This indicates that, once the momentum source is sufficiently regularised, the present breakup and drag framework can recover the correct order of magnitude of the Sauter mean diameter in the higher-$We_g^{\text{film}}$ regime, even though the velocity field and the full droplet-population distribution remain imperfect. This is a technically meaningful result for a case study that is numerically and physically far more severe than Lambert, and it does not alter the choice of $b = 3D_{\text{inj}}$ as the representative compromise across observables.

The campaign establishes that, unlike in the open Lambert crossflow, the DBM is a *necessary* ingredient for the confined GLOP geometry. The per-cell PCM (`ce11`) projection triggers a SIGFPE (Section 5.3); the stabilised `ce11Point` bare-PCM baseline scatters parcels into a diffuse, domain-filling cloud (Sections 5.5 and 5.8) that the DBM smoothing collapses onto a coherent bent jet, restoring a physically meaningful spatial distribution.

The field visualisations add a consistent mechanistic picture of why the bandwidth matters so much more here than in the open Lambert crossflow. The DBM acts on the Eulerian side, broadening the stagnation dipole and the oblique momentum-deficit wake of the carrier; at $We_g^{\text{film}} \approx 500$ the carrier *topology* is qualitatively preserved across the bandwidths, whereas at $We_g^{\text{film}} \approx 3500$ the pressure and velocity-magnitude fields become markedly sensitive to the smoothing length scale and do not reduce to a simple monotonic regularisation. Because in the confined GLOP geometry the Lagrangian source occupies a far larger fraction of the local carrier momentum than in Lambert (Section 5.1), the regularisation is decisive rather than a mere refinement, which is also why a moderate $b = 3D_{\text{inj}}$ already suffices while larger bandwidths give only diminishing returns.

The most energetic case $We_g^{\text{film}} \approx 8500$ could not be completed, and the instability mechanism was not isolated. Plausible contributors include the stiffer two-way coupling, time-step restrictions, the growth of the Lagrangian population and the localisation of

the source terms, together with the un-smoothed (per-cell) mass and energy coupling terms; the deteriorating velocity prediction already visible at $We_g^{\mathrm{film}} \approx 3500$ is consistent with the last of these but does not by itself identify the cause. Extending the screened-diffusion smoothing to the remaining Lagrangian source terms is a candidate remedy to be verified, discussed as future work in Section 6.3.

# 6 Conclusions

This thesis addressed the numerical modelling of liquid-jet primary atomisation in a gaseous crossflow, in both the open LJICF configuration and the confined LJIGF/GLOP variant, within a two-way-coupled Eulerian–Lagrangian framework in OpenFOAM. The chosen approach combines the Madabhushi–Lambert column-breakup and KH/Wave stripping model with a deformation-dependent drag law and a diffusion-based coarse-graining (DBM) regularisation of the Lagrangian momentum source. The model was assembled in Chapters 2 and 3 and assessed against two validation campaigns in Chapters 4 and 5.

## 6.1 Summary of Results and Model Advantages

On the implementation side, the work delivered a self-contained solver, `sprayFoamDBM`, in which the Madabhushi–Lambert model is realised through an extension of `BreakupModel.H` that supports multi-fragment child emission, a compact state machine encoded in the `ms` parcel variable that selects the active breakup stage—KH/Wave column stripping followed by Pilch–Erdman secondary breakup—, a state-dependent drag law that switches between blob, sphere and deforming-disc regimes, and a screened-diffusion DBM operator applied to the accumulated momentum source before the carrier momentum equation is solved (Chapter 3).

In the open Lambert LJICF benchmark (Chapter 4), the model reproduced the measured Sauter diameter and streamwise velocity profiles, with $b = 3D_{\mathrm{inj}}$ adopted as the representative bandwidth. The bare PCM projection already gave normalised $D_{32}$ errors of 28–32% and $U_x$ errors of 10–20%; both generally decreased with increasing bandwidth, with the largest reductions—roughly halving the $D_{32}$ error—at $b = 3D_{\mathrm{inj}}$ and $b = 6D_{\mathrm{inj}}$, consistent with the regularising role of the DBM. The residual mismatch concentrated at the near plane ($x/d = 30$) was traced to a breakup-timing offset between the SST-RAS closure and the reference SBES data, rather than to a deficiency of the breakup model itself.

In the confined Zhang LJIGF/GLOP configuration (Chapter 5), the DBM proved to be a necessary rather than an optional ingredient. The per-cell PCM projection (`cell`) triggered a floating-point exception, while the stabilised `cellPoint` variant ran but produced a spatially distorted, domain-filling droplet distribution; only with $b = 3D_{\mathrm{inj}}$ did the predicted spray collapse onto the coherent bent jet expected on physical grounds. The representative bandwidth $b = 3D_{\mathrm{inj}}$ is thus common to both campaigns; in the confined Zhang geometry, however, the DBM is not merely a refinement but a neces-

sary ingredient for a physically meaningful spray, reflecting the stronger near-injector coupling, where the source occupies a larger fraction of the local carrier momentum.

A central advantage of the approach is that the DBM regularisation introduces no calibration constant beyond the single bandwidth $b$, whose value is tied to the injector length scale (of the order of a few $D_{\mathrm{inj}}$) rather than tuned to match the data, in keeping with the coarse-graining interpretation of Sun and Xiao [10]. This distinguishes the method from purely empirical source-smoothing corrections. The validation supports the usefulness of the DBM framework within the tested Lambert and Zhang conditions, but it also shows that the preferred smoothing bandwidth depends on the observable, the operating regime and the region of the spray being compared: $b = 3D_{\mathrm{inj}}$ is retained as a representative compromise for the present solver and mesh rather than as a universal optimum, and a larger value such as $b = 6D_{\mathrm{inj}}$ can be preferable for specific observables, notably $D_{32}$ at $We_g^{\mathrm{film}} \approx 3500$.

## 6.2 Limitations of the Current Implementation

The most visible limitation stems from the SST-RAS closure. Without resolved turbulent eddies, the onset of the first Pilch–Erdman breakup of the intact core is delayed, producing the near-plane ($x/d = 30$) timing offset seen in both Lambert cases. A scale-resolving closure such as LES or SBES would capture the eddies that trigger earlier fragmentation, consistent with the closer agreement of the reference SBES curve in the Lambert comparison. More broadly, both reference datasets were produced with scale-resolving turbulence treatments—the Lambert benchmark with an SBES-DPM (Lagrangian) solution that does not resolve the liquid–gas interface, and the Zhang benchmark with an interface-resolving, LES-based VOF–LPT solver—whereas the present model relies on an SST-RAS carrier; the breakup-timing offset is therefore a structural consequence common to both campaigns rather than a deficiency of the breakup model itself.

The DBM is presently applied only to the momentum source, while the mass and energy coupling terms are left at their per-cell values; this appears adequate at $We_g^{\mathrm{film}} \approx 500$ and $We_g^{\mathrm{film}} \approx 3500$ but may become a limitation at higher Weber numbers. The most energetic case, $We_g^{\mathrm{film}} \approx 8500$, could not be run to completion, and its failure could not be assigned to a single mechanism. Possible contributors include the stronger two-way coupling, the increased Lagrangian population, time-step restrictions, the localisation of the source terms, local thermodynamic excursions and, if active, the un-smoothed scalar or energy source terms. A dedicated numerical study would be required to isolate their relative importance; the growing discrepancy in the mean droplet velocity already visible at $We_g^{\mathrm{film}} \approx 3500$ is consistent with, but does not by itself establish, the last of these.

Three further limitations concern the methodology rather than the physics. First, the representative bandwidth is not a fixed property of the model: the comparisons show that the preferred $b$ depends on the mesh resolution, the operating regime and the observable, so it must be assessed case by case rather than transferred blindly. Second, in the Zhang campaign the bare-PCM baseline differs from the DBM cases not only

in $b$ but also in the carrier-to-parcel interpolation scheme (`cellPoint` versus `cell`), so that the PCM-versus-DBM comparison conflates the two effects, whereas only the $b = D_{\mathrm{inj}}$, $3D_{\mathrm{inj}}$, $6D_{\mathrm{inj}}$ series isolates the bandwidth. Third, the validation statistics are snapshot-based and rely on sectional binning with a minimum-count threshold, so sparsely populated bins—particularly near the jet base of the most under-resolved cases—are discarded and the comparison is restricted to the better-sampled regions.

The blob drag coefficient is calibrated per configuration ($C_{D,\mathrm{blob}} = 1.48$ for the Lambert Jet-A1 jet and 2.16 for the Zhang water jet, with $C_{D,\mathrm{disc}} = 1.2$ retained throughout), so that a priori knowledge of the working fluid and geometry is required and the model cannot be transferred to a new configuration without recalibration. In the same spirit, the secondary breakup follows the Pilch–Erdman correlations [9], which are valid for the bag and multimode regimes; the stripping and catastrophic regimes encountered at very high Weber numbers are not represented.

Finally, the dispersed phase is treated as inert: the current implementation includes neither heat exchange nor phase change, restricting its applicability to cold atomisation. The absence of a resolved turbulent-dispersion mechanism further limits the fidelity of the predictions in the dilute spray region far from the injector, where turbulent fluctuations control the lateral spreading of the droplet cloud.

## 6.3 Future Developments

A natural extension is to generalise the screened-diffusion treatment to any additional Lagrangian source terms that are active in evaporating or reacting configurations—the energy and species sources (`heTrans`, `YiTrans`)—completing a framework that is already in place for the momentum source. This is motivated by the Zhang campaign, in particular the $We_g^{\mathrm{film}} \approx 8500$ non-completion and the deteriorating velocity prediction at $We_g^{\mathrm{film}} \approx 3500$; it should, however, be undertaken after quantifying whether these terms contribute significantly to the numerical instability observed in the most severe cases, and it requires modifying the time loop in `sprayFoamDBM.C`.

Closely tied to this is the introduction of a heat-exchange and evaporation (phase-change) model for the dispersed phase, which the present implementation treats as inert. The same Lagrangian energy and species source terms whose DBM regularisation is discussed above (`heTrans`, `YiTrans`) would carry the interphase heat and mass transfer, so that lifting the inert assumption extends the framework from cold atomisation to evaporating and ultimately reacting configurations.

A further priority is a systematic sensitivity study in the bandwidth, the mesh resolution and the time step: repeating the Lambert J10 and Zhang $We_g^{\mathrm{film}} \approx 500$ cases on refined and coarsened grids would clarify how far the representative bandwidth is a physical rather than a numerical length scale. The present results support a bandwidth of order a few injector diameters, with $b = 3D_{\mathrm{inj}}$ retained as the representative compromise; the appropriate value should, however, be assessed with respect to the case, mesh resolution, observable and operating regime rather than treated as a universal optimum, and larger values such as $b = 6D_{\mathrm{inj}}$ may be preferable for specific observables, notably $D_{32}$ at $We_g^{\mathrm{film}} \approx 3500$. Common-support error metrics, statistics conditioned on the

parcel state and a more direct localisation of the breakup events would further sharpen the comparison. Looking further ahead, a locally adaptive bandwidth proportional to the local cell size would extend the method to strongly non-uniform meshes, and an automated bandwidth selection could be envisaged, but as a case-dependent rather than a universal rule.

An extension to supersonic and compressible crossflow—relevant to advanced propulsion and hypersonic-vehicle applications—is a further natural direction. Because the column-breakup time scale is already formulated in terms of the local, instantaneous gas density and slip velocity, and the underlying `sprayFoamDBM` solver is fully compressible (it advances the continuity, energy and pressure equations of the carrier phase), the breakup model is in principle operative in the compressible and supersonic regime without modification to the breakup formulation itself. What such a regime does require is a dedicated validation effort: shock-capturing numerical schemes for the compressible carrier, together with a recalibration of the Kelvin–Helmholtz instability growth and of the drag laws at high Mach number.

# References


[1] G. M. Faeth, L.-P. Hsiang, and P.-K. Wu. "Structure and breakup properties of sprays". In: *International Journal of Multiphase Flow* 21(Suppl.) (1995), pp. 99–127.

[2] P.-K. Wu, K. A. Kirkendall, R. P. Fuller, and A. S. Nejad. "Breakup processes of liquid jets in subsonic crossflows". In: *Journal of Propulsion and Power* 13(1) (1997), pp. 64–73. DOI: 10.2514/2.5151.

[3] Meng Zhang, Shuai Zhou, Lekun Song, Xu Zhang, Jie-Yao Lyu, and Fangmian Dong. "Atomization dynamics of gas-liquid orifice type pintle injectors at different throttling levels". In: *Energy* 335 (2025), p. 138192. DOI: 10.1016/j.energy.2025.138192.

[4] X. Jiang, G. A. Siamas, K. Jagus, and T. G. Karayiannis. "Physical modelling and advanced simulations of gas–liquid two-phase jet flows in atomization and sprays". In: *Progress in Energy and Combustion Science* 36(2) (2010), pp. 131–167. DOI: 10.1016/j.pecs.2009.09.002.

[5] OpenCFD Ltd. *OpenFOAM: The Open Source CFD Toolbox — User Guide and C++ Source Guide, Version v2412*. ESI Group. 2024. URL: https://www.openfoam.com/documentation/.

[6] R. K. Madabhushi. "A model for numerical simulation of breakup of a liquid jet in crossflow". In: *Atomization and Sprays* 13(4–5) (2003), pp. 413–424.

[7] R. D. Reitz. "Modeling Atomization Processes in High-Pressure Vaporizing Sprays". In: *Atomisation and Spray Technology* 3 (1987), pp. 309–337.

[8] Markus Lambert, Thomas Esch, Markus Braun, and Hossam Elasrag. "Enhancement of the Madabhushi Liquid Jet in Crossflow Breakup Model by a Ligament Breakup Mechanism". In: *ILASS – Europe 2019, 29th Conference on Liquid Atomization and Spray Systems*. Paris, France, Sept. 2019.

[9] M. Pilch and C. A. Erdman. "Use of breakup time data and velocity history data to predict the maximum size of stable fragments for acceleration-induced breakup of a liquid drop". In: *International Journal of Multiphase Flow* 13(6) (1987), pp. 741–757.

[10] R. Sun and H. Xiao. "Diffusion-based coarse graining in hybrid continuum–discrete solvers: Theoretical formulation and a priori tests". In: *International Jour-*

*nal of Multiphase Flow* 77 (2015), pp. 142–157. DOI: 10.1016/j.ijmultiphaseflow.2015.08.014.

[11] F. R. Menter. "Two-equation eddy-viscosity turbulence models for engineering applications". In: *AIAA Journal* 32(8) (1994), pp. 1598–1605. DOI: 10.2514/3.12149.

[12] P. A. N. Nordin. "Complex Chemistry Modeling of Diesel Spray Combustion". PhD thesis. Chalmers University of Technology, 1998.

[13] A. Putnam. "Integrable Form of Droplet Drag Coefficient". In: *ARS Journal* 31(10) (1961), pp. 1467–1468. DOI: 10.2514/8.5825.

[14] L. Schiller and A. Naumann. "Über die grundlegenden Berechnungen bei der Schwerkraftaufbereitung". In: *Zeitschrift des Vereines Deutscher Ingenieure* 77 (1933), pp. 318–320.

[15] A. B. Liu, D. Mather, and R. D. Reitz. "Modeling the Effects of Drop Drag and Breakup on Fuel Sprays". In: *SAE Technical Paper 930072*. 1993. DOI: 10.4271/930072.

[16] Lorenzo Proli. *Diffusion-Based Spray Solver for OpenFOAM*. Version 1.1.1. 2026. DOI: 10.5281/zenodo.20666066. URL: https://doi.org/10.5281/zenodo.20666066.

[17] Y. Gopala, P. Zhang, O. Bibik, E. Lubarsky, and B. T. Zinn. "Liquid Jet in Crossflow – Trajectory Correlations Based on the Column Breakup Point". In: *48th AIAA Aerospace Sciences Meeting*. AIAA Paper 2010-214. Orlando, Florida, USA, 2010.

[18] J. Sekar, A. Rao, S. Pillutla, A. Danis, and S. Y. Hsieh. "Liquid Jet in Cross Flow Modeling". In: *Proceedings of ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*. Düsseldorf, Germany, 2014.