

Validation Dossier: literature-equivalence evidence for the in-house microfluidics solver

Study S7 (arc MF), in-house microfluidics solver

Generated 2026-09-14 (S7.11/MF.11 close-out session)

Why this document exists

`PRE_REQUISET_ADVANCE_PLANNING/studies/s7_microfluidics/MF_VALIDATION_PROGRAM.md` (the standing validation program, defined 2026-08-24) commits to exactly this deliverable: a LaTeX rendering of the benchmark registry, organized so a reader can find, for any validated claim, what it is testing, its literature basis, whether an analytical solution exists, and the measured comparison against whatever reference is available. The markdown equivalent (`reports/mfV_validation_dossier.md`) already existed; this document adds the structured, per-case LaTeX form `MF_VALIDATION_PROGRAM.md`'s own Phase C names.

COMSOL-comparability argument, stated honestly

Verbatim from MF_VALIDATION_PROGRAM.md §5:

This is not a head-to-head comparison against COMSOL. The COMSOL artifacts this study was commissioned to replace do not exist on this machine (verified MF.0: `deep find` over `/data`, `/data2`, `/home/krithika`, `mlocate`, full git history — nothing found). What follows is equivalence evidence: for each benchmark, the certified reference interval, this solver's measured result, its ASME V&V 20-style comparison error, and — where a published COMSOL result on the same benchmark could be located — that result in the same row.

A validation dossier with no misses in it is not a validation dossier. This one keeps: the Ghia cavity $Re=400/1000$ mesh-independent mismatch (MF.1); DFG 2D-3's 13% $c_{L,max}$ deviation and $\Delta p(t=8)$ sign flip, traced to resolution sensitivity near a zero-crossing rather than a sign-convention bug (MF.9); the two independent two-way-coupling criteria that disagree by 5–6 orders of magnitude for intermediate particle concentrations (MF.9); and whatever A.5 (Dean-vortex resolution) and A.6 (Segré–Silberberg equilibrium) actually return, pass or fail. The six real, load-bearing bugs MF.8/MF.9 caught and fixed — including a driver script silently reporting a perturbation-forcing spike as a genuine periodic-state force coefficient — are evidence the validation process works, not an embarrassment to omit.

This session adds: DFG 3D-1Z's own genuine mesh-convergence finding (c_D converges into the certified interval, c_L does not yet, see §6.28), and the 3D ForwardDiff-MMS closure that an earlier session had marked done without actually deriving a case (see §4.4).

The five-tier benchmark ladder

Full detail in `MF_VALIDATION_PROGRAM.md` §1; summarized here. A result at a higher tier is worth more than the same-looking pass at a lower one — Tier 0 only proves the code solves its own discretized equations correctly, Tier 3 proves those equations are the right equations.

Tier	Name	What it establishes
0	Code verification	Discrete solution converges to the equation the code claims to solve; no physical-fidelity claim
1	Analytical solutions	Closed-form NS/Stokes solutions on simple geometries
2	Certified literature benchmarks	Community-converged reference intervals (DFG, Ghia, Gartling)
3	Physical-experiment validation	Real measured phenomena (Taylor–Aris dispersion, Segré–Silberberg pinch)
4	Solution verification	Richardson/GCI mesh-independence

Reading this dossier

Cases are grouped by `CFDCore.jl` module (the same sort key `render_markdown_table` uses), one `\section` per module and one `\subsection` per case. Each case has three parts: **Purpose** and **Literature** (hand-written; filled for every case as of 2026-09-20 — where no external reference applies the section says so explicitly rather than carrying an invented citation, and every bibliography entry was verified by web search before transcription), **Analytical solution** (structured: either the closed form itself or an explicit statement that none exists), and **Result** (regenerated live from the registry: tolerance, measured result, verdict, traceability, and any `honest_caveat`).

1 ChannelGeometry

1.1 Channel boundary-loop closure and enclosed-area preservation

Purpose

The channel wall polygon is the single geometric object every later stage (meshing, tagging, particle containment) is built on. This check would catch a closure error in `channel_boundary_loop` (a polygon that does not return to its start), a wrongly-ordered top/bottom wall traversal that flips the orientation, or an offset construction that silently changes the enclosed area. For a straight channel the area must equal length times width exactly; for a mildly curved centerline it must equal arc length times width up to the expected first-order curvature error, so a larger discrepancy indicates a real construction bug rather than curvature.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Yes: enclosed area = length*width for a straight channel (exact); arc-length*width is a first-order estimate for a mildly-curved centerline.

Result

Category closed_form

Tolerance / bar rtol=1e-10 (straight, exact rectangle); rtol=0.02 (mildly-curved, first-order discretization)

Measured result 1e-10 relative error (straight); ~2% relative error (curved serpentine)

Verdict **PASS**

Traceability ChannelGeometry | centerlines and boundary loops (MF.0) (ChannelGeometry)

2 ElectroThermal

2.1 MF.13: electric-potential (conduction) solver on the unstructured 2D/3D path – exact quadrupole field and grad of E squared, uniform-field conductor, weak Dirichlet pinning of shared nodes

Purpose

The dielectrophoretic, electrophoretic and thermophoretic force kernels of this module need an electric field $\mathbf{E} = -\nabla\phi$, and until this unit the field had to be prescribed by hand. This case checks the solver that supplies it: a steady conduction problem $\nabla \cdot (\sigma \nabla \phi) = 0$ with Dirichlet electrodes and insulating (natural) boundaries elsewhere, on the same unstructured 2D and 3D meshes as the flow solver. The test is against exact answers, the hyperbolic quadrupole $\phi = G(y - y_c)(z - z_c)$ (whose $\mathbf{E}$ and $\nabla|\mathbf{E}|^2$ a quadratic element reproduces exactly) and the uniform field of a conductor, so that any residual is a defect of the implementation and not of the discretisation. It also records a defect that this check exposed: nodes shared by a Dirichlet face and a natural face carry a single label and were left free, so the electrode reaction was missing at one node per shared edge; enforcing every Dirichlet boundary weakly as well removes it.

Literature

The governing equation is the quasi-static, real-conductivity limit of current conservation; the role of the resulting $\nabla|\mathbf{E}|^2$ in the dielectrophoretic force follows Jones [20], and the use of a conduction potential as the first stage of an electrokinetic calculation is standard in the microelectrode literature [29]. No published number is being reproduced: the reference is the closed-form quadrupole, chosen because it lies in the finite-element space.

Analytical solution

Hyperbolic quadrupole $\phi = G(y - y_c)(z - z_c)$: $\mathbf{E} = -\text{grad } \phi$ and $\text{grad}(\mathbf{E}^2) = 2 G^2 (y - y_c, z - z_c)$ in closed form; uniform-field conductor.

Result

Category closed_form

Tolerance / bar quadrupole $\mathbf{E}$ and $\text{grad}(\mathbf{E}^2)$ to 1e-6 relative; uniform conductor field to 1e-6

Measured result $\mathbf{E}$ and $\text{grad}(\mathbf{E}^2)$ to about 1e-9 relative after the fix (about 5e-5 before)

Verdict **PASS**

Traceability MF.13 ElectroThermal | potential solver: exact quadrupole field and $\text{grad}(E^2)$, uniform-field conductor, refusals (ElectroThermal)

Honest caveat. A steady, linear, real-conductivity conduction problem: no frequency dependence, no permittivity (so no displacement current, no double layer), one-way (nothing feeds back on sigma). For AC/DEP use the potential must be supplied in RMS-equivalent volts. The penalty for the weak Dirichlet term is $1e6$ sigma over the smallest bounding-box extent, which leaves a leakage of order $1e-6$ relative. Validated against closed forms only; no experiment and no COMSOL comparison exist.

2.2 MF.13: Joule heating – steady conduction with source $\sigma(\text{grad } \phi)^2$ [squared magnitude] reproduces the parabolic temperature rise $\sigma E^2 y(H-y)/(2k)$ and the total power

Purpose

Joule heating is the coupling COMSOL’s microfluidics offerings pair with current conservation: the dissipated power $Q = \sigma|\nabla\phi|^2$ is a heat source for conduction. This case checks the one-way version built here, a temperature solved from the potential of the previous case, against the exact temperature rise of a uniform current in a slab with isothermal walls, $T - T_w = \sigma E^2 y(H-y)/(2k)$, at three heights, and against the energy balance (integrated power equal to $\sigma E^2 LH$). It is deliberately a check of the implementation only: the model has no temperature-dependent conductivity, no fluid convection of heat and no electrothermal body force.

Literature

Heat generation by an applied field and the temperature gradients it produces are described in the review of forces in microelectrode structures by Ramos and co-workers [29], which also explains why the electrothermal body force, absent here, is what makes the effect matter for fluid motion. The verification solution is the elementary parabolic profile of a heated slab and needs no citation beyond it.

Analytical solution

Uniform current density in a slab with isothermal walls: $T(y) = T_w + \sigma E^2 y(H-y)/(2k)$.

Result

Category closed_form

Tolerance / **bar** temperature profile to $1e-4$ relative; integrated power to $1e-6$ relative

Measured result profile and power match; centreline rise matches $\sigma E^2 H^2/(8k)$ to about $4e-11$ relative after the weak-Dirichlet fix ($1.4e-3$ before)

Verdict **PASS**

Traceability MF.13 ElectroThermal | Joule heating: parabolic temperature rise, energy balance, refusals (ElectroThermal)

Honest caveat. ONE-WAY and steady: sigma, k and the flow do not respond to temperature, there is no electrothermal body force (the buoyancy and permittivity-gradient forces that drive electrothermal flow), no convective heat transport by the fluid and no temperature-dependent conductivity. Closed-form verification only; no COMSOL or experimental comparison.

3 FieldSampler

3.1 Curved-duct material derivative (Du/Dt) vs analytic centripetal acceleration, and streamline radius-holding

Purpose

In a straight channel the material derivative is identically zero, so its computation had never been tested where it matters. In the curved duct it must equal the centripetal acceleration u_θ^2/r at mid-radius. This check would catch a wrong sign, a missing convective contribution in the sampler's Du/Dt , and a streamline-radius drift, all of which feed the pressure-gradient and added-mass forces on particles.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Yes: the analytic centripetal-acceleration Du/Dt for fully-developed curved flow.

Result

Category closed_form

Tolerance / bar Du/Dt relative error target: small (no pre-registered bar, reported as measured); streamline radius drift $\leq 5e-3 \cdot (Ro - Ri)$

Measured result Du/Dt relative error 0.095% at mid-radius; streamline radius drift well within $5e-3 \cdot (Ro - Ri)$

Verdict **PASS**

Traceability MF.4 | curved-duct case: Du/Dt , streamline, and equilibrium-position ensemble (items 6,15,16) (FieldSampler)

3.2 Field sampler exactness, robustness, and wall distance vs analytic Poiseuille

Purpose

The particle tracer never touches the finite-element space directly; it samples velocity, pressure and their derivatives through an interpolation layer. Because the meshes are straight-sided, a Taylor-Hood P_2 field is exactly a ten-term quadratic per cell, so the sampler should reproduce the solved field to round-off. This check catches errors in cell location, coordinate mapping or reconstruction, and confirms the wall-distance routine against the analytic Poiseuille geometry.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Yes: the plane-Poiseuille closed-form velocity/pressure field.

Result

Category `closed_form`

Tolerance / bar 0 locator failures on random interior points; velocity/pressure error μ i field scale; gradient `rtol=1e-6`; `Du/Dt atol=1e-8` (exactly zero expected, fully-developed flow); wall distance `atol=1e-9`

Measured result 0/500 (build-time exactness test) and 0/2000 (MF.2's own wider sweep) locate failures; velocity error μ 1e-5 (peak scale 2.5e-4 m/s); pressure error μ 0.01 (scale 1.0); gradient matches analytic to $\sim$ 1e-6-1e-10 relative; wall distance exact to μ 1e-9

Verdict **PASS**

Traceability MF.2 `FieldSampler` | `exactness and robustness (FieldSampler)`

3.3 FieldSampler 3D: exact P2/P1 reconstruction on tets, robust locator, wall distance, channel scale

Purpose

The 3D sampler must reproduce the solved field exactly at interior points, because on straight-sided tetrahedra a P_2 field is exactly a ten-term quadratic per cell. The check confirms this against the Hagen–Poiseuille pipe field, together with the robustness of the point locator, the wall-distance routine and the channel-scale calculation. A failure would corrupt every 3D trajectory.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Yes, indirectly: cross-checked via the exact Hagen-Poiseuille closed-form pipe-flow field (sampled off the FEM solution, not compared to the closed form directly by this entry).

Result

Category `cross_check`

Tolerance / bar sampler exactness `rtol=1e-8`; `locate_3d exterior=0` (exact); `wall_distance_3d/channel_scale_3d rtol=0.05-0.1` (coarse mesh faceting, not a bug – same discipline as the 2D exactness bar)

Measured result all checks passed on the calibration pipe case

Verdict **PASS**

Traceability MF.10 FieldSampler | 3D exactness and robustness (FieldSampler)

3.4 MF.V B.2: multi-wall lift superposition produces a genuine off-axis (cornerward) force a single-wall model cannot, on a cheap synthetic square duct

Purpose

A single-nearest-wall lift can only produce a continuous ring in a square duct, because a particle near a corner never feels its second, equally close wall. This case builds and mechanism-tests a multi-wall superposition that gives corner points a genuine off-axis force. On a synthetic square duct it reduces exactly to one wall at a face centre and to two equidistant walls at a corner. Its later run against the real ensemble (recorded in the wiring and clean-field cases) found discrete attractors at the corners, not the face centres.

Literature

The target it was built to address is the four face-centre equilibria reported by Di Carlo et al. [11]; the per-wall term is Ho-Leal's [18].

Analytical solution

No closed form; a geometric/mechanism self-consistency check (exact single-wall reduction, qualitative corner off-axis force), not a literature-benchmark comparison.

Result

Category cross_check

Tolerance / bar exact reduction to inertial_lift_3d at a face centre (not approximate – one wall found, one term); qualitative mechanism check at a corner (a real second nonzero force component appears), not a quantitative literature target

Measured result face centre: 1 wall found, $F_{mw} = F_{single}$ exactly; corner: 2 walls found (equidistant, near-orthogonal normals), multiwall force has strictly more nonzero components than the single-wall force at the identical point

Verdict **FINDING**

Traceability MF.V FieldSampler | wall_distances_multi_3d: square-duct corner vs face-centre (B.2) (FieldSampler)

Honest caveat. This entry confirms the MECHANISM on a small, cheap, non-production synthetic square duct. The follow-up it originally left open has since been RUN (S7.12/MF.12 Step 4, 2026-09-20): :ho_leal_3d_multiwall is now wired into ForceModel3D/trace_particles_3d! (see :mf12_multiwall_tracer_wiring) and was run against the real Di Carlo ensemble on a noise-free field (see :mf12_di_carlo_gate_clean_field). Result: it produces DISCRETE attractors, but at the four CORNER points of the ring, NOT at Di Carlo et al. (2007)'s face-centre positions – so it does not resolve the specificity miss; it replaces one wrong answer with another. (Older text in this study calls the literature target '4 discrete corners'; the target is face-centre. Corner equilibria are reported for other Re/size regimes, which this result neither confirms nor refutes.)

4 FlowBenchmarks2D

4.1 Exact-solution reference formulas verified via ForwardDiff before any FEM comparison relies on them

Purpose

Every downstream comparison trusts an exact solution to be exact. This check applies automatic differentiation (ForwardDiff) to substitute the Kovasznay solution and the hand-derived curved-annulus azimuthal profile back into their own governing equations before any finite-element comparison uses them. It caught a real error during development (an initial particular-solution constant in the curved-annulus profile), which a finite-element comparison alone would have mis-attributed to the solver.

Literature

The Kovasznay solution is the classical exact Navier–Stokes solution behind a two-dimensional grid [22]. The verify-the-reference-first discipline follows the method-of-manufactured-solutions literature [32]. The curved-annulus profile is derived in this study and has no published counterpart to cite.

Analytical solution

Yes: Kovasznay’s exact Navier-Stokes solution (Kovasznay 1948) and a hand-derived curved-annulus azimuthal Stokes closed form, both independently re-verified via ForwardDiff before use.

Result

Category closed_form

Tolerance / **bar** ODE/PDE residual near machine precision (target $\sim 1\text{e-}13$ to $1\text{e-}10$)

Measured result curved-annulus ODE residual $\sim 1\text{e-}13$ after the sign fix (was off by ~ 2000 before it)

Verdict **PASS**

Traceability MF.1 self-checks | exact-solution formulas verified via ForwardDiff before use (FlowBenchmarks2D)

4.2 Backward-facing step reattachment-length trend vs Gartling (1990)

Purpose

The backward-facing step is the first benchmark in the ladder with a separation bubble, so it tests whether the solver reproduces a qualitative flow feature (reattachment length increasing with Reynolds number) rather than a smooth profile. A wrong convective term, a wrong outflow condition, or an under-resolved step corner would each degrade or reverse that trend. The check was run at the cheaper $\text{Re} = 200/400$ pair rather than the reference $\text{Re} = 800$ point, and it uncovered a real bug during development.

Literature

Gartling [15] tabulates the reference reattachment length (6.10 step heights at $Re = 800$). The source substituted here for Armaly’s original experimental table is documented in the reference’s own docstring, so the comparison is against the numerical benchmark, not the experiment.

Analytical solution

No closed form; validated against Gartling (1990)’s numerically-converged backward-facing-step reattachment-length trend (a numerical, not analytical, reference).

Result

Category `certified_literature`

Tolerance / **bar** monotonic increase with Re (trend check, not a quantitative match to 6.10)

Measured result `x_r/S`: 6.05 ($Re=200$) - `i` 9.74 ($Re=400$)

Verdict **PASS**

Traceability MF.1 FlowSolver2D | backward-facing step: reattachment length trend (Gartling 1990 geometry) (FlowBenchmarks2D)

Honest caveat. Not attempted at the reference $Re=800$: a first attempt ran for over two hours of CPU time without finishing and was aborted as disproportionate for one verification item in a time-boxed session. The cheaper $Re=200/400$ pair confirms the qualitative trend and both runs converge cleanly, but this is a trend check, not a quantitative match to Gartling’s 6.10 – disclosed, not silently substituted.

4.3 3D Tier 1 analytical references: rectangular-duct Fourier series + Hagen-Poiseuille, verified before use

Purpose

This is the FIRST 3D check run at all in this study (MF.10’s own gate before any 3D benchmark was trusted) – it establishes that `solve_flow_3d`’s Taylor-Hood assembly, boundary-condition application, and linear solve are correct in three dimensions on the simplest possible geometries, before any 3D result involving a curved/bluff-body domain (DFG 3D-1Z, the Dean-duct sweep) is attempted. Run in the order MF.1 established for the 2D solver: verify the reference FORMULAS themselves first (via `ForwardDiff`, independent of the FEM solve), then compare the FEM solution against them. Would catch: a wrong basis-function count/ordering in 3D, an incorrectly extended vector Laplacian, or a 2D-only piece of code silently truncating a 3D field.

Literature

Hagen-Poiseuille: the exact parabolic solution for pressure-driven flow in a circular pipe, attributed to Hagen (1839) and Poiseuille (1840/1846) – textbook material, e.g. White, *Viscous Fluid Flow*. The rectangular-duct Fourier-series solution (the classical separation-of-variables result for pressure-driven flow in a rectangular cross-section) is the same standard reference. Both are the right comparison for THIS check specifically because they are closed-form (no independent numerical

solve to disagree with, isolating the FEM discretization’s own correctness) and geometrically trivial (a pipe, a rectangle), so a mismatch here can only mean a genuine 3D-assembly defect, not a resolution or geometry-approximation effect.

Analytical solution

Yes: the classical White/Berker rectangular-duct Fourier-series solution and the exact Hagen-Poiseuille parabolic profile.

Result

Category `closed_form`

Tolerance / **bar** PDE residual $\leq 0.1\%$ of `dpx` at the default `nmax=200`; wall BCs exact to $1e-12$

Measured result both self-checks pass; `solve_flow_3d` subsequently matches `rectangular_duct_velocity` to $\leq 2\%$ (coarse mesh vs. non-polynomial series) and `hagen_poiseuille_velocity` to $\leq 1e-8$ (exactly P2-representable quadratic)

Verdict **PASS**

Traceability MF.10 FlowBenchmarks2D | 3D Tier 1 exact-solution self-checks (validation program §1) (FlowBenchmarks2D)

4.4 3D ForwardDiff-manufactured-solution order of accuracy (validation program Tier 0, closed this session)

Purpose

The validation program’s Tier 0 row for 3D order of accuracy had been marked “covered by reuse” without a genuine 3D derivation. Closing it needed a new construction: a vector potential $\mathbf{u} = \nabla \times \Psi$, the 3D analogue of the 2D stream function, because a single scalar has no divergence-free 3D equivalent. Measured convergence rates of about 2.9–3.0 in velocity and 3.8 in pressure, with mass conservation at round-off, establish that the 3D assembly is consistent.


Literature

The order-of-accuracy procedure with manufactured solutions follows Roache [32]; the vector-potential construction is derived in this study.

Analytical solution

Yes: a genuine new 3D ForwardDiff-manufactured exact solution, built from a divergence-free vector-potential ($\mathbf{u}=\text{curl}(\Psi)$) construction (this session – the 2D module’s single-scalar stream-function trick has no direct 3D analogue).

Result



Genuine new 3D derivation this session: a divergence-free vector-potential ($u=\text{curl}(\Psi)$) construction, the 3D analogue of the 2D module's single-scalar stream-function trick, which has no direct 3D equivalent. $n=(4,6,8)$ on the unit cube, kept coarse per this study's own A.0 direct-LU cost-scaling finding.

Category closed_form

Tolerance / bar P2 velocity order ≥ 2.5 , P1 pressure order ≥ 1.5 (same bars as the already-validated 2D MMS case)

Measured result 3 refinement levels ($n=4,6,8$ on the unit cube, kept coarse per this study's own A.0 direct-LU cost-scaling finding): $\text{rate}_u=[2.92, 2.97]$, $\text{rate}_p=[3.83, 3.82]$, both clear of their bars; mass_conservation_residual at machine precision ($\approx 3e-18$) at every level; div-free check $\approx 5e-16$ at 3 spot points before the FEM solve was trusted

Verdict PASS

Traceability MF.V FlowBenchmarks2D | 3D ForwardDiff-MMS: measured order of accuracy (validation program Tier 0) (FlowBenchmarks2D)

5 FlowDiagnostics

5.1 Streamline tracing, equal-flux seeding, directional separation-point detection, and case_diagnostics

Purpose

These diagnostics turn a solved field into engineering outputs. Streamline drift is measured rather than assumed negligible; equal-flux seeding must place streamtubes at their target flux fractions; and directional separation-point detection must recover the known backward-facing-step reattachment point, cross-checked exactly against an independent inline probe on the same field. A failure would indicate a wrong integrator, a wrong flux weighting or a sign error in the separation criterion.

Literature

The separation cross-check uses the backward-facing-step benchmark geometry and reattachment reference of Gartling [15].

Analytical solution

No closed form; separation_points is cross-checked against this study's own independently-coded Gartling BFS probe, a numerical (not analytical) cross-check.

Result

Category cross_check

Tolerance / bar streamline drift: corrected to match the measured RK4-in-arclength truncation, not an unverified a priori bound; equal-flux seeding: rtol=0.02; separation_points vs MF.1 probe: rtol=1e-6; case_diagnostics: atol=1e-9

Measured result straight-channel streamline drift $\sim 2.1\text{e-}8$ m absolute ($\sim 0.1\%$ of H at $ds=H/20$); equal-flux seeding matches target to rtol=0.02; separation_points vs MF.1's Gartling probe agree to rtol=1e-6; case_diagnostics exact to 1e-9

Verdict **PASS**

Traceability MF.4 FlowDiagnostics | streamlines and equal-flux seeding (straight channel) (FlowDiagnostics)

Honest caveat. The originally-planned streamline-drift acceptance bar ($1\text{e-}9H$, an a priori assumption written before any real measurement) was wrong by many orders of magnitude – the actual mechanism is RK4-in-arclength truncation error ($\sim 2\text{e-}8$ m, $\sim 0.1\%$ of H at $ds=H/20$), not a claim to machine precision (the field's own representational exactness, shown by the FieldSampler entries above, is close to machine precision – the streamline's drift scales with the arclength step instead). The bound was corrected to match the measurement, not the other way around. Separately: separation_points' first implementation returned the FIRST sign change of any kind, which for the Gartling BFS is the step-corner SEPARATION point (~ 4.07 length units), not the far-downstream REATTACHMENT point (~ 9.70) MF.1's own probe measures – both routes agreed on the field to $\sim 1\text{e-}13$ relative but were finding different, both genuinely real, crossings; fixed by adding a 'direction' keyword.

5.2 Wall shear stress: closed form plus two independent cross-checks, and vorticity

Purpose

Wall shear stress must be computed from the sampler's exact per-cell polynomial, never from a boundary-triangulation gradient of a C^0 field, whose gradient is discontinuous across element edges (measured more than twenty times wrong in this codebase's own history). This check validates the method against the plane-Poiseuille closed form and against two independent cross-checks, and also checks vorticity, so a regression to the discontinuous gradient would be caught.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Yes for wall shear stress: the plane-Poiseuille closed-form value; independently cross-checked via a control-volume momentum balance (which has no closed form of its own).

Result

Category `cross_check`

Tolerance / bar WSS vs closed form: relative error over the mid-80% of channel length (outlet-adjacent do-nothing-BC deviation excluded, a real documented artifact); vorticity vs closed form: `rtol=1e-4`, centerline exactly zero to `1e-6`

Measured result WSS vs closed form: $6.99\text{e-}7$ max relative error; vs momentum balance: 0.044% ; vs `reaction_force`: $\sim 4.2\%$; vorticity: `rtol=1e-4` match, centerline zero to `1e-6`

Verdict **PASS**

Traceability MF.4 FlowDiagnostics | wall shear stress, vorticity, momentum balance (FlowDiagnostics)

Honest caveat. Three independent WSS routes agree to within a few percent of each other, not to a suspiciously tight bound. `wall_shear_force` agrees with the control-volume momentum balance to 0.044% , but with Gridap’s own `reaction_force` identity (applied to a boundary-layer coordinate indicator) only to $\sim 4.2\%$. Most plausibly this is because a nodal indicator function on a continuous P2 space is a coarser, quadrature-limited approximation to a true boundary delta function than either of the other two routes – `reaction_force` itself was already proven exact via Stokes-drag linearity in MF.1’s confined-cylinder test, so the gap is attributed to the indicator approach, not the underlying identity. All three agreeing to within a few percent IS the validation result, not a tighter number to chase.

5.3 Four-vortex transition search: natural continuation + deliberate perturbed-branch check, curved square duct

Purpose

This case documents a miss. The literature predicts a transition from two to four Dean vortices at a particular Dean number, and two independent searches (natural continuation over the whole sweep, and a deliberately perturbed initial state) both failed to find it, returning the same unique steady state to machine precision at $De = 150$. The result is reported as a miss so that a possible resolution or mesh limitation is not hidden.

Literature

The four-vortex transition is discussed in the curved-duct literature reviewed by Berger, Talbot and Yao [2]; the two-vortex baseline is Dean’s [9]. The miss concerns this solver’s agreement, not the reference.

Analytical solution

No closed form; the four-vortex transition Dean number is an empirically-observed literature threshold (the Berger-Talbot-Yao 1983 comparator), not derived from a closed-form solution.

Result

Category certified.literature

Tolerance / bar outer-wall probe: expects a transition from 1 to 3 crossings at the literature’s critical De (definition/aspect-ratio-sensitive; not independently re-verified this session, see caveat); branch check: a relative difference at or below numerical Newton tolerance ($\sim 1e-10$) after a large deliberate perturbation would indicate a shared branch

Measured result outer-wall probe: exactly 1 crossing at all 13 De points $De=2$ to $De=260$; branch check: relative difference $2.1e-16$ (machine precision) after a 44%-of-u.s.mean perturbation

Verdict **MISS (documented)**

Traceability MF.10 A.5 FlowDiagnostics | Dean cross-section diagnostics (FlowDiagnostics)

Honest caveat. This gate clause MISSES outright – no transition found by either route, across a range that spans well past the $De \sim 150$ this study anticipated as the likely regime. A REAL, separate, unresolved finding from the same campaign: retrying the non-converged $De=190$ point with 8 continuation stages (up from 2) still failed to converge on its final 3 stages ($\rho=750/875/1000$ of 1000) even at 60 Newton iterations each, despite the first 5 stages (through $\rho=625$) converging cleanly in 5-8 iterations – something in the $De \sim 150-190$ range makes Newton’s method struggle sharply harder than either side of it. This is reported as OPEN, NOT as evidence for or against the four-vortex transition: it is equally consistent with approaching a genuine bifurcation/turning point in the solution branch and with a mesh-resolution or continuation-staging limitation, and this session did not isolate which. The literature’s own critical Dean number for a curved SQUARE duct is itself definition- and aspect-ratio-sensitive across sources; this session did not independently re-verify a specific numeric target the way A.6 did for Segre-Silberberg’s $0.6R$, so this miss is reported against the qualitative signature (an outer-wall-localized extra vortex pair) rather than a specific unverified De value.

5.4 Dean secondary-flow intensity scaling vs. De (Ookawara 2004) and vortex topology, curved square duct

Purpose

Curved channels drive a secondary Dean flow, and separation in them depends on its strength. This check sweeps the Dean number from about 2 to 260 with real Navier–Stokes solves and measures how the secondary-flow peak scales. At low Dean number the local exponent matches Dean’s De^2 perturbation result, and it declines toward the asymptote reported for higher Dean numbers, which would expose a wrong curvature term or an under-resolved secondary flow.

Literature

Dean [9] gives the De^2 scaling at low Dean number; Berger, Talbot and Yao [2] review curved-pipe flow including the higher-Dean-number behaviour; Ookawara et al. [27] report the asymptotic exponent for a curved microchannel.

Analytical solution

Partial: Dean (1928)’s own low-De asymptotic De^2 scaling is a closed-form perturbation result; the full intensity curve is compared against Ookawara (2004)’s empirically-fitted correlation, which is not a closed form.

Result

Category `certified_literature`

Tolerance / bar exponent: qualitative match to the De^2 -1.63 trend, no formal tolerance (a shape claim, not a point value); magnitude: within a factor of ~ 2.5 of Ookawara’s correlation, no stated reference uncertainty to compare a tighter bound against

Measured result exponent 1.99-2.01 (De 2-10) declining to 1.10 (De=127-150); `u_sec_peak/dean_velocity_lit` ratio 1.2-2.5 across all 10 converged points; `sym_residual` $3.2e-4$ to $7.7e-3$

Verdict **FINDING**

Traceability MF.10 A.5 FlowDiagnostics | Dean cross-section diagnostics (FlowDiagnostics)

Honest caveat. TWO real caveats, not concealed. (1) The ‘`theta_resid`’ development check (mid-arc vs. 75%-to-outlet) shows 10-53% disagreement at every De, worst at the lowest De (0.53 at De=2) – the entrance transient may not be fully decayed by mid-arc over this study’s ‘`angle=pi/4`’ arc, meaning the very points that match Dean’s exponent most cleanly are also the least confirmed ‘fully-developed’ by this study’s own diagnostic; a longer arc or a direct entrance-length study would retire this, not attempted this session. (2) ‘`dean_vortex_count`’’s raw core count (4 to 30 across the sweep) is NOT a small integer at any De – diagnosed as a resolution/threshold limitation of the local-omega-s-extrema criterion at this cross-section grid (`n_r=n_z=25`, `frac=0.1`), not evidence against a genuine base 2-vortex structure (the field-level `sym_residual` evidence is more trustworthy here); a coarser/adaptive threshold or a smoothed field is the natural fix, not attempted this session. The last 3 of 13 De targets (190/237/260) did not converge at this staging and are excluded from the exponent/ratio numbers above.

6 FlowSolver2D

6.1 Entrance-length scaling and non-Newtonian consistency / viscous-form gap

Purpose

Two auxiliary solver checks from the same session. The entrance length is compared with the semi-empirical scaling $L_e/H \approx 0.05 Re$, using the measured outlet flux as the fully-developed reference (the inlet plug value is overridden at the wall-shared degrees of freedom). The non-Newtonian path is checked for self-consistency (a constant viscosity function must reduce exactly to constant viscosity) and the gap between the Laplacian and stress viscous forms is quantified rather than assumed negligible.

Literature

No primary source is cited: $L_e/H \approx 0.05 \text{Re}$ is a widely used textbook scaling that was not re-verified against a primary reference here, and the viscous-form comparison is a self-consistency check with no external reference. The $\text{Re} = 10$ and 100 deviations are reported as measured.

Analytical solution

Mixed: entrance-length scaling uses a semi-empirical correlation ($L_e/H \sim 0.05 \text{Re}$, not a strict closed form); the non-Newtonian viscous-form gap is a self-consistency cross-check with no external analytical reference.

Result

Category `cross_check`

Tolerance / bar $\text{Re}=40$: exact match to 0.05Re ; $\text{Re}=100$: within $\pm 30\%$; $\text{Re}=10$: loose, correlation breakdown expected at low Re , documented not hidden; non-Newtonian: machine precision; `:laplacian` vs `:stress` gap: quantified, not compared to an external bar

Measured result L_e/H : 2.00 ($\text{Re}=40$, vs $0.05 \cdot 40 = 2.00$ exact); 4.57 ($\text{Re}=100$, vs 5.00); 0.86 ($\text{Re}=10$, vs 0.5 nominal); non-Newtonian consistency 0.0 diff (machine precision); `:laplacian` vs `:stress` relative L2 gap 1.9%

Verdict **PASS**

Traceability MF.1 FlowSolver2D | entrance length $L_e/H \sim 0.05 \text{Re}$ (FlowSolver2D)

6.2 Curved duct: fully-developed azimuthal Stokes flow vs hand-derived closed form

Purpose

Curvature enters the study through the curved-duct particle cases, so the flow solver must be checked on a curved boundary with a known answer. A hand-derived azimuthal Stokes profile is imposed as a Dirichlet condition throughout (a manufactured-solution instance with a ForwardDiff-derived body force). Agreement checks the solver on curved geometry with curved-boundary tagging, and a mismatch would point to the geometry mapping or the boundary treatment rather than to the physics.

Literature

No citation: the closed form is derived in this study from the azimuthal momentum equation and verified with ForwardDiff (see the `fb_forwarddiff_selfcheck` case).

Analytical solution

Yes: a hand-derived fully-developed azimuthal Stokes-flow closed form for a curved annular duct, independently confirmed via ForwardDiff before use.

Result

Category closed_form

Tolerance / bar exact-solution comparison, closed-form velocity profile (not the true curved-channel pressure field)

Measured result L2 velocity error 1.3e-6

Verdict PASS

Traceability MF.1 FlowSolver2D | curved duct: fully-developed azimuthal Stokes flow (FlowSolver2D)

6.3 Confined-cylinder Stokes drag linearity and DFG 2D-1 (Schafer-Turek, Re=20) certified intervals

Purpose

Every prior check in this ladder (Poiseuille/Couette, Kovasznay, the curved-duct closed form) validates the solver on a domain with either a trivial or an analytically-tractable geometry. DFG 2D-1 is the first check that combines a genuinely curved, resolution-sensitive boundary (a cylinder) with the reaction-force identity (**reaction_force**, which extracts drag/lift by integrating the weak-form momentum residual against an indicator test function rather than a direct boundary-traction integral) under a real graded mesh. It would catch: a wrong sign or scale factor in the reaction-force identity itself, an under-resolved boundary layer producing a plausible-looking but wrong drag/lift pair, or a graded-mesh generation bug that silently under-tags the cylinder surface – none of which a closed-form or Kovasznay check, run on a simple domain, is positioned to expose.

Literature

Schäfer & Turek (1996) [35], case 2D-1 (steady, Re=20): the DFG Priority Research Programme’s own community benchmark, Table 3. Chosen over a single group’s published number because it reports a certified interval built from 17 independent research groups across finite-element, finite-volume, finite-difference, and lattice-Boltzmann methods (Table 1/2 of the paper) – landing inside that interval is evidence of agreement with the converged community answer, not with one competing implementation’s own possible bias.

Analytical solution

No closed form; Schafer & Turek (1996)’s certified, community-converged reference intervals (Table 3).

Result

Category certified_literature

Tolerance / bar $F(2U)/F(U)=2.0$ exact; c_D certified [5.57,5.59]; c_L certified [0.0104,0.0110]; δp certified [0.1172,0.1176]; $GCI \leq 0.01$

Measured result $F(2U)/F(U)=2.0000000$, lift/drag=0.0096; $c_D=5.577$, $c_L=0.01063$, $\delta p=0.11748$ (all in-interval); Richardson-extrapolated $c_D=5.582$, $GCI=0.0023$

Verdict **PASS**

Traceability MF.1 FlowSolver2D | DFG 2D-1 (Schäfer-Turek, Re=20): cD/cL/dp vs certified intervals (FlowSolver2D)

6.4 Plane Poiseuille and Couette: exact to near machine precision

Purpose

Taylor–Hood P_2 velocity represents a quadratic profile exactly, so plane Poiseuille and Couette flow should be recovered to near machine precision, not merely to a loose engineering tolerance. This is the tightest possible test of the assembly, the boundary-condition imposition and the linear solve together: any discrepancy above round-off indicates a real defect (a sign, a missing boundary term, a wrong quadrature degree) rather than discretisation error.

Literature

No citation is needed: the plane Poiseuille and Couette profiles are textbook closed forms, and the point of the check is agreement to round-off with them.

Analytical solution

Yes: plane Poiseuille and Couette flow, the textbook closed-form parabolic and linear velocity profiles.

Result

Category closed_form

Tolerance / **bar** machine precision (target $\sim 1e-8$ to $1e-11$)

Measured result Poiseuille $L_2(u)=1.5e-11$, $L_2(p)=1.6e-8$; finite-Re NS identical to Stokes to $1e-8$ ($u.\text{grad}(u)=0$ exactly for fully-developed flow); Couette $L_2(u)=1.0e-11$

Verdict **PASS**

Traceability MF.1 FlowSolver2D | Poiseuille and Couette: exact to near machine precision (FlowSolver2D)

6.5 Lid-driven cavity Re=100 vs Ghia et al. (1982)

Purpose

The lid-driven cavity is the canonical test of the coupled velocity–pressure solve with a corner singularity and no analytical solution. At Re = 100 it checks that the bulk velocity field matches the community reference profiles along the cavity centrelines, catching errors in the convective term, the lid boundary condition or the pressure treatment that a smooth-profile test cannot expose.

Literature

Ghia, Ghia and Shin [16] tabulate centreline velocities from a high-resolution multigrid solution; they remain the standard reference for this benchmark.

Analytical solution

No closed form; Ghia et al. (1982)'s own numerically-converged benchmark tables.

Result

Category certified_literature

Tolerance / bar max bulk error ; 0.10 (this project's own established bar)

Measured result max bulk error 0.087

Verdict **PASS**

Traceability MF.1 FlowSolver2D | lid-driven cavity vs Ghia et al. 1982 (FlowSolver2D)

6.6 Lid-driven cavity $Re=400/1000$ vs Ghia et al. (1982) – documented miss

Purpose

This case documents a miss rather than a pass. At $Re=400$ and 1000 the same bulk comparison against the reference tables does not meet the bar that holds at $Re=100$, and no acceptance bar had been established at these Reynolds numbers, so the result is reported as measured instead of being judged against a bar chosen afterwards. Its purpose in the dossier is to keep a known, mesh-independent disagreement visible.

Literature

The reference is again Ghia et al. [16]. The miss is a statement about this solver's agreement with those tables, not about the reference.

Analytical solution

No closed form; same Ghia et al. (1982) numerical reference tables as fs_ghia_re100.

Result

Category certified_literature

Tolerance / bar no bar established at $Re=400/1000$; reported honestly

Measured result max bulk error 0.299 ($Re=400$); 0.366 ($Re=1000$)

Verdict **MISS (documented)**

Traceability MF.1 FlowSolver2D | lid-driven cavity vs Ghia et al. 1982 (FlowSolver2D)

Honest caveat. Ghia cavity $Re=400/1000$ do not tightly match the bulk-region reference, and the discrepancy is mesh-independent. A refinement sweep ($Re=1000$, ndofs 21k-¿37k-¿58k) left the error essentially flat (0.366-¿0.366-¿0.366) – ruling out plain under-resolution. Two BC families were tried (regularized leaky-lid and literal discontinuous uniform lid, the latter solved with this module’s damped-Newton + rho-continuation driver, which DOES converge cleanly on the literal lid at every Re tested – a genuine fix for the divergence this project’s memory records for Gridap’s undamped NLSolver). Neither BC choice closes the gap. Direct inspection of the $Re=1000$ centerline shows the mismatch concentrated in the mid-cavity secondary-vortex region (the FEM solution’s return-flow is markedly weaker than Ghia’s, even reversing sign where Ghia shows strong negative velocity around $y=0.6-0.85$) – this points to boundary-condition-shape sensitivity of the high- Re secondary vortex structure (a documented sensitivity in the lid-cavity literature), not a discretization defect, but this session did not have budget to pursue a boundary-layer-graded near-wall mesh to test that hypothesis further. Left open for a later session, flagged rather than hidden.

6.7 Kovasznay $Re=40$ and ForwardDiff-manufactured MMS: measured order of accuracy

Purpose

A solver can agree with one exact solution by luck and still be formally inconsistent. Measuring the convergence rate on a refinement sequence against an independently verified exact solution (Kovasznay) and against a ForwardDiff-manufactured source tests the discretisation’s actual order: about third order in velocity and second order in pressure are expected for P_2/P_1 . A lower measured order would indicate an error in the convective term, the weak form or the quadrature.

Literature

Kovasznay’s exact solution [22] is the standard Navier–Stokes verification case; the order-of-accuracy procedure with a manufactured source follows Roache [32].

Analytical solution

Yes: Kovasznay’s exact solution and a ForwardDiff-manufactured MMS solution (also exact by construction); this entry’s own check is the resulting order of accuracy, not a direct value comparison.

Result

Category cross_check

Tolerance / bar order ¿ 2.5 (velocity, ~3rd order expected for P_2); order ¿ 1.5 (pressure, ~2nd order expected for P_1)

Measured result Kovasznay: $L_2(u)$ order ¿ 2.5; MMS: $L_2(u)$ order ¿ 2.5, $L_2(p)$ order ¿ 1.5

Verdict **PASS**

Traceability MF.1 FlowSolver2D | Kovasznay and MMS: measured order of accuracy (FlowSolver2D)

6.8 3D direct-LU feasibility gate: measured DOF vs (RAM, time), full-eel-channel extrapolation

Purpose

Before any three-dimensional production code was written, the direct-LU Stokes solve was measured at two real DOF counts on the production aspect ratio, a power law was fitted (time about $\text{DOF}^{1.79}$, memory about $\text{DOF}^{1.31}$), and the result extrapolated to the full channel's roughly 34 million degrees of freedom. The purpose is to replace a guess about feasibility with a measurement; the finding, that a monolithic direct solve is decisively infeasible, is what motivated the leg-relay decomposition and, later, the preconditioned solver of MF.13.

Literature

No external literature applies; the exponents are this machine's own measurements. They are consistent in kind with the superlinear fill-in expected of three-dimensional sparse direct factorisation.

Analytical solution

Not applicable – a compute-cost measurement (DOF vs RAM/time); no physics solution is involved.

Result

Category diagnostic

Tolerance / bar measured, not guessed – two real data points (96,826 and 240,050 tets), both solved to machine-precision mass conservation

Measured result extrapolated isotropic direct-LU solve: ~19.2 hours, ~1.73 TB RAM, against this machine's 251GB total – decisively infeasible, not a close call

Verdict **FINDING**

Traceability MF.10 MicroMesh | 3D tagged gmsh meshes (FlowSolver2D)

Honest caveat. Isotropic direct-LU on the full-length eel.500um channel is infeasible on this machine. Production A.7 sweep MUST use the anisotropic-streamwise-mesh lever (or a block-preconditioned Krylov fallback), justified by a real streamwise-refinement study, not asserted from this extrapolation alone. Full account in notes.md's "MF.10 A.0" entry.

6.9 3D absolute-ftol/volume-integral-residual false-convergence bug: found, diagnosed, documented

Purpose

A calibration solve at micrometre scale reported `newton_converged = true` with a machine-precision mass residual while every interior degree of freedom was exactly zero. The absolute tolerance `ftol`, tuned in two dimensions, is satisfied by the natural magnitude of a three-dimensional volume-integral residual at the untouched initial state. This case documents the diagnosis so that any 3D solve requiring at least one Newton iteration and a tiny tolerance is understood as a requirement,

not a preference. (MF.13 met a second instance of the same species of trap inside the Krylov solver; see `mf13_krylov_absolute_tolerance_trap`.)

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Not applicable – a solver-convergence-criterion bug diagnosis; no physics solution comparison is involved.

Result

Category diagnostic

Tolerance / bar n/a (a methodology finding, not a pass/fail bar)

Measured result root-caused and reproduced on demand; documented on `solve_flow_3d/solve_flow_3d_transient`'s own docstrings, not silently patched (the shared, already-validated `_damped_newton_solve` stays byte-for-byte unchanged, preserving the 2D regression gate)

Verdict **FINDING**

Traceability MF.10 FlowSolver2D | 3D solver: `dimension guards, exact-solution match, mass conservation (FlowSolver2D)`

Honest caveat. Both `mass_conservation_residual` AND `newton_iterationsi=1` are now the required pair of sanity checks for any 3D FlowSolution at this length scale – mass conservation ALONE is not sufficient (a trivial zero field trivially conserves mass). Every future 3D production case (A.5/A.6/A.7) must check both, not assume the default `ftol` is safe.

6.10 MF.12 Step 5 (Tier 4): 3D exact-Stokes square duct – velocity L2 error converges monotonically; classical Richardson/GCI not applicable

Purpose

First 3D mesh-convergence evidence in this program. It checks that the 3D Taylor-Hood solver's velocity error against the exact rectangular-duct solution decreases monotonically under refinement (a solver that failed to converge, or returned the interpolated boundary data, would fail it) and records that the classical Richardson/GCI procedure is not usable for point or integral quantities at these resolutions because they converge non-monotonically.

Literature

Roache (1994) grid convergence index procedure, applied via the existing two-dimensional helper; the exact solution is the Fourier-series solution for pressure-driven flow in a rectangular duct (White, *Viscous Fluid Flow*; Berker 1963). The observed L2 order of 1.6 to 1.8 is below the nominal order of 3 for this element pair and has not been explained.

Analytical solution

Exact fully-developed rectangular-duct Stokes solution (Fourier series, White/Berker), compared directly at every level.

Result

Category closed_form

Tolerance / bar monotone decrease of the exact-solution L2 error under refinement; GCI validity requires monotone convergence and an observed order near the nominal one

Measured result L2 rel error 0.94% -; 0.50% -; 0.24%, observed order 1.6, 1.8; GCI not applicable to u_probe / u_center / the integral of u squared

Verdict FINDING

Traceability MF.12 Tier 4 | square-duct exact-Stokes 3D convergence (velocity L2 error vs the exact solution) (FlowSolver2D)

Honest caveat. Observed L2 order 1.6-1.8 is below P2’s nominal 3. (MF.13 update: localised to the zero-traction OUTLET region – see :mf13_t4.outlet.localisation; the mechanism is still a hypothesis.) The finer triple H/8, H/12, H/18 was refused here on a memory gate while the sweep ran, and was run in MF.13 on an idle machine: L2 0.497%, 0.238%, 0.122%; peak RSS at H/18 (208k DOF) MEASURED 16.72 GB against 14.89 GB predicted (12% low, still fitted). It did not change the picture, because all three levels share the same natural outlet. The eel.500um production geometry still has no 3D convergence study (its 3-level GCI was already found infeasible on this machine, see MF.12). The velocity field also carries a spurious TRANSVERSE velocity (exactly 0 in the exact solution) of $\sim 1e-3 \cdot U_m$ at H/8 that matters for lift gates (see :mf12_di_carlo_gate_clean_field).

6.11 MF.13 Phase 5: Krylov.jl’s GMRES declares convergence on an ABSOLUTE machine-precision test; unit-norm scaling plus a true-residual check make the Krylov Newton step safe at micrometre scale

Purpose

The first version of the block-preconditioned Newton step reproduced the direct solution at step one and then failed to converge: GMRES reported convergence after four iterations and then one per step, while the true linear residual was 0.27, 0.81, 0.88 and 0.91. The library’s GMRES applies an absolute machine-precision stopping test that fires once the residual norm falls below about 1.1×10^{-16} , and Newton residuals at micrometre scale run from 10^{-8} down to 10^{-25} . The fix (solve the unit-norm problem and always recompute the true residual) makes a false “solved” impossible to pass silently; this case pins that behaviour with a synthetic saddle system at three residual scales.

Literature

GMRES is Saad and Schultz’s algorithm [33]; the block preconditioner is of the type analysed by Silvester and Wathen [38]. The absolute test was found by reading the installed library’s source (Krylov.jl 0.9.10) and confirmed by measurement; it is not described in that library’s documentation.

Analytical solution

No closed form; the assertion is on the computed TRUE residual.

Result

Category regression

Tolerance / bar true relative residual of every Krylov Newton step at most 10 times `krylov_rtol`, at right-hand-side norms $1e-8$, $1e-16$ and $1e-25$

Measured result true relative residual at most $1e-7$ at all three scales; pre-fix 0.27 to 0.91

Verdict **FINDING**

Traceability MF.13 StokesPreconditioner | Krylov.jl's hard-coded absolute convergence test cannot fire on micrometre-scale residuals (FlowSolver2D)

Honest caveat. Reproduced on a synthetic saddle system (FE-independent) and on the real 3D duct, but only against the installed Krylov.jl 0.9.10; a different Krylov.jl version may or may not carry the same test, and the unit-norm scaling is harmless either way. This is a defect in the FIRST draft of this project's own code that measurement caught, and a hazard for any future Krylov use at this scale, not a bug report against Krylov.jl (its documented tolerances are respected; the surprise is the undocumented absolute test).

6.12 MF.13: Navier (Maxwell) slip and electroosmotic wall conditions on solve_flow_2d/3d, validated against the exact slip and EOF-plus-pressure Poiseuille profiles

Purpose

Slip flow and electroosmotic flow are boundary features in COMSOL's microfluidics module, but here they existed only as analytic one-dimensional closed forms and could not drive the finite-element solver. This case validates the new wall conditions against exact profiles: a Maxwell slip length on the walls must reproduce the slip Poiseuille profile, a free-slip wall must preserve a plug, and the Helmholtz–Smoluchowski slip velocity must reproduce both the pure electroosmotic plug and its superposition with a pressure-driven flow. It also records a real defect found on the way: a node shared by a Dirichlet inlet and a slip wall was left free, and that single loose degree of freedom per corner shifted the whole exact profile by 4.3%, which evaluating the discrete residual of the exact solution located immediately.

Literature

No external literature is cited: the slip Poiseuille profile and the thin-double-layer electroosmotic velocity are textbook closed forms already implemented and tested in `SlipFlow.jl` and `ElectroosmoticFlow.jl`, and the comparison is against those closed forms, not against a published numerical result. A COMSOL result for these conditions was not located.

Analytical solution

Plane Poiseuille flow with Maxwell slip length l at both walls, and the thin-EDL superposition u_{HS} plus Poiseuille (Karniadakis and Beskok's slip-flow closed forms; see `SlipFlow.jl` and `ElectroosmoticFlow.jl`).

Result

Category `closed_form`

Tolerance / bar exact slip profile to 1e-5 of scale; EOF profiles to 1e-6; free-slip plug to 1e-3

Measured result slip profile error about 2e-6; EOF profiles 1e-6; free-slip plug 1e-3 (2D) and 2e-3 (3D); 4.3 percent before the corner fix

Verdict **PASS**

Traceability MF.13 WallBCs | Navier slip reproduces the exact slip Poiseuille profile (corner nodes pinned), free slip preserves a plug (FlowSolver2D)

Honest caveat. STEADY solves only (the transient solver does not take `navier_slip`). The no-penetration condition is a penalty (default 1e8 μ over the smallest bounding-box extent), consistent for solutions with $u.n = 0$ and leaving a leakage of order p/β , and the Newton loop runs to its cap at $ftol = 1e-18$ because that lies below the penalty’s roundoff floor (use $ftol$ of about 1e-13 to 1e-15, as the tests do). The EOF condition is the thin-double-layer limit with a PRESCRIBED tangential field: no electrostatics is solved, the double layer is not resolved, and a non-uniform zeta or field is not supported. Slip is validated only on straight walls (the exact reference is a straight channel); curved-wall behaviour and the slip-length-versus-Knudsen mapping (second-order slip, temperature jump, thermal creep) are not addressed. Not validated against a COMSOL result (none located).

6.13 MF.13 Phase 5: block-preconditioned GMRES Newton step with the Stokes Schur approximation – correct, mesh-independent and 3.6x lighter in memory at Stokes, but it fails at Re of about 15 or more (DFG 3D-1Z); see `:mf13_saddle_point_pcd` for the convection-aware fix

Purpose

The direct sparse solve inside every Newton step was measured (MF.10) to cost about $\text{DOF}^{1.8}$ in time and to need 20 GB at 283k degrees of freedom, which is what made production-scale three-dimensional cases and the unsteady DFG 3D benchmarks infeasible. This case measures a replacement: GMRES preconditioned by a block-triangular operator that approximates the Schur complement by the pressure mass matrix and the velocity block by AMG or an exact LU. It is correct (identical solutions), mesh-independent at Stokes with an exact velocity block, and 3.6 times lighter in memory at 283k DOF with AMG. It fails at Reynolds number of roughly 15 or more (DFG 3D-1Z), so the benchmarks that needed it remain open, and the reason is now quantified.

Literature

GMRES is Saad and Schultz’s algorithm [33]; the pressure-mass-matrix Schur approximation with a block preconditioner is analysed by Silvester and Wathen [38]. The DFG 3D-1Z test case is defined by Schäfer and Turek [35].

Analytical solution

Exact fully-developed rectangular-duct Stokes solution for the Stokes levels; DFG 3D-1Z has no closed form (certified interval, see `:mfv_dfg_3d1z`).

Result

Category `cross_check`

Tolerance / **bar** Krylov solution equals the direct solution; GMRES iterations flat under refinement at Stokes; true linear residual at most 10x the tolerance on every step

Measured result identical to 10 digits; iterations flat (exact LU) or mildly growing (AMG) at Stokes; 3.6x less memory with AMG at 283k DOF; FAILS at Re of about 15-20 in 3D

Verdict **FINDING**

Traceability MF.13 StokesPreconditioner | `:block_lu` and `:block_amg` reproduce the direct solve on a 3D exact-Stokes duct (FlowSolver2D)

Honest caveat. This entry describes the STOKES (pressure-mass-matrix) Schur approximation only. It is the classical Stokes one and degrades with Reynolds number, which is why convection-aware approximations exist; the pressure convection-diffusion variant WAS built afterwards and is recorded in `:mf13_saddle_point_pcd`, which supersedes the finite-Re statements here for the exact-LU velocity block. Newton iteration counts are NOT comparable across solvers (the inexact solves need two steps where the direct solve needs one). The first Krylov call at each size includes compilation; the WARM (second-step) linear solves take 0.1, 2.8, 31.8 s with AMG against 0.6, 3.4, 169.3 s direct at 3.7k, 33k, 281k DOF (time exponents about 1.1 and 1.8), so AMG is NOT faster than direct at 33k DOF (it needs two steps to direct's one) and wins clearly only at 281k. One geometry family (a straight square duct for Stokes, one cylinder-in-channel case for finite Re), one tolerance ladder. A first draft of this solver failed silently because of Krylov.jl's absolute convergence test (see `:mf13_krylov_absolute_tolerance_trap`).

6.14 MF.13: pressure convection-diffusion (PCD) Schur approximation – removes the finite-Re failure of the Stokes approximation (DFG 3D-1Z solved, duct iterations flat under refinement) but only with an exact-LU velocity block; DFG 3D-2Z at Re 100 stays out of reach

Purpose

The block-preconditioned Krylov solve built earlier in this unit is correct at Stokes flow but fails at a Reynolds number of about 15 to 20, because its Schur-complement approximation, the pressure mass matrix, is a Stokes result. This case records the convection-aware replacement, the pressure convection-diffusion approximation, and measures exactly what it fixes: the duct at Reynolds number 22.5 (iteration counts under a nine-fold refinement) and the DFG 3D-1Z benchmark, which the mass approximation could not solve at all. It is a *finding*, with a hard limit stated in the caveat: the gain requires an exact factorisation of the velocity block, the scalable multigrid block still fails at finite Reynolds number, and a bounded feasibility run did not reach the Reynolds number 100 of the DFG 3D-2Z case, which therefore remains unrun.

Literature

The pressure convection-diffusion preconditioner is that of Kay, Loghin and Wathen [21], who also give the mesh-independence result and report GMRES iteration growth roughly like the square root of the Reynolds number; the Stokes mass-matrix approximation it replaces is the Silvester-Wathen

result [38]. The choice of a Dirichlet pressure condition on both inflow and outflow for the discrete Laplacian was made by measurement in this study, after a natural inflow condition proved worse than the Stokes approximation.

Analytical solution

Exact fully-developed rectangular-duct solution (a Navier-Stokes solution at any Re because $\mathbf{u} \cdot \nabla \mathbf{u} = 0$); DFG 3D-1Z has no closed form (certified interval, see :mfv_dfg_3d1z).

Result

Category cross_check

Tolerance / bar Krylov solution equals the direct solution; GMRES iterations bounded and not growing with refinement at finite Re

Measured result same solution to 12 digits; duct iterations 35 to 38 flat over 9x more DOF; DFG 3D-1Z solved; AMG velocity block fails at every finite-Re level

Verdict **FINDING**

Traceability MF.13 StokesPreconditioner | pressure convection-diffusion (PCD) Schur approximation: same solution as the direct solve, fewer GMRES iterations than the Stokes mass approximation at Re 22 (FlowSolver2D)

Honest caveat. A FINDING WITH A HARD LIMIT. (1) The gain needs an EXACT-LU velocity block, which is the expensive part: the peak memory is a saving of only about 20 percent at 59k DOF and wall time is 30 percent WORSE, so nothing here makes a large 3D finite-Re case cheaper than the direct solve; the AMG variant that would be the one that fails. (2) It does not reach Re = 100: a bounded feasibility run on the coarse DFG 3D-1Z mesh (mf13_dfg3d2z_feasibility.jl, NOT the benchmark) stepped the steady solution to about a third of the target (Re about 33) with GMRES iterations growing from 44 to 256-348, and stopped with an error at the next continuation stage (Re about 42) on the 400-iteration cap; no transient step, drag or lift was obtained, so DFG 3D-2Z and 3Z remain unrun. The iteration growth with Re is the known behaviour (Kay et al. report growth like the square root of Re); the measured growth here is faster. (3) The failing PCD runs were logged before the driver recorded the Schur type in failure records, so their 'pcd' label is inferred from the order the runs were executed in. (4) Two geometries, one mesh level for DFG. (5) The transient path detects the inflow tag at $t = t_{\text{end}}/2$ and is exercised only by a 2D unit test (three steps at Re 10, same answer as the direct solve to 1e-6); it was never run in 3D.

6.15 MF.13 Phase 1 (Tier 4): the sub-nominal global L2 order of the 3D square-duct study is inherited from the zero-traction OUTLET region

Purpose

MF.12's first three-dimensional convergence study measured a global velocity- L_2 order of 1.6–1.8 against the nominal 3 of P_2 elements and left it unexplained; a finer triple on an idle machine (H/8, H/12, H/18) reproduced it (1.8, 1.6), so more resolution alone could not explain it. On an exactly nested mesh sequence the error was split by streamwise region and compared with the P_2 interpolation error of the exact field on the same mesh, which is the best any solution in that space

can do. The interior and inlet regions track the interpolation error at order 3, while the outlet quarter, where the natural zero-traction condition applies, does not; imposing exact Dirichlet data there removes the excess. The purpose is to tell a defect in the discretisation (which would need fixing) from an artefact of how the error was measured (which needs only a different measurement).

Literature

The grid-convergence index of Roache [31] presupposes monotone convergence at a near-nominal order, which is why it is reported as not applicable to these quantities; the use of a known exact solution follows the verification literature [32]. No published source characterises this specific outflow effect for this element pair, and none is claimed.

Analytical solution

Exact fully-developed rectangular-duct Stokes solution (Fourier series), plus its P2 interpolant on the same mesh as the best-approximation reference.

Result

Category closed_form

Tolerance / **bar** FEM error on each streamwise region within 1.5x of the P2 interpolation error of the exact field

Measured result interior and inlet FEM/interp 0.99-1.21 (order 2.98 and 3.00 at the finest pair); outlet quarter FEM/interp 3.38, 7.09, 15.8 (order 1.77); Dirichlet outlet 1.01, 1.15

Verdict FINDING

Traceability MF.13 Tier 4 | the sub-nominal global L2 order is inherited from the zero-traction outlet region (nested H/4 -> H/8) (FlowSolver2D)

Honest caveat. This LOCALISES the sub-nominal order; it does not establish the mechanism. The exact solution satisfies the homogeneous natural outlet condition ($p = 0$, $du_x/dx = 0$ there), so this is not an inconsistency of the exact solution with the BC; a consistency error of the weakly-imposed discrete traction is the candidate and was NOT isolated. The leading alternative – an edge-ring effect where the essential wall BC meets the natural outlet BC – was tested and is NOT supported: at H/4 and H/8 the outlet excess is the same in the outlet core (FEM/interp 3.56, 7.67) as in a wall shell of thickness 0.15 H (3.31, 6.90). The all-Dirichlet intervention was run at H/4 and H/8 only (the finest, 283k-DOF level of it was not run beside other jobs). One geometry (a straight square duct, $L_x = 2 H$), one BC pair. Practical consequence: the global L2 error of a duct with a natural outlet is not a valid convergence-order measurement; use an interior slab or an all-Dirichlet outlet. Production A.7 legs are unaffected in their interior by construction (they overlap and hand off), but this was not re-tested there.

6.16 MF.13: density-based topology optimisation of a Stokes flow (Borrvall-Petersson) – envelope gradient equals finite differences to 1e-7, dissipation falls monotonically by a factor 7.8 on a pipe bend, volume constraint exact

Purpose

Density-based topology optimisation lets a flow solver design the geometry that carries the flow, and COMSOL couples its flow interfaces to an optimisation module for it. Earlier work in this study had prepared the pieces (an exact Brinkman channel solution, the dissipated-power identity and a finite-difference check of the gradient) but no design variable or optimiser. This case checks the complete loop on a pipe-bend problem: the analytic gradient against finite differences, the monotone fall of the dissipated power, the volume constraint and the connectivity of the resulting design. It is a check of the machinery, not a claim about the optimum.

Literature

The formulation, a Brinkman friction that penalises solid regions in a Stokes flow and a dissipated-power objective with a fluid-volume constraint, is that of Borrvall and Petersson [3], from whom the self-adjoint gradient follows. No optimum from their paper is reproduced.

Analytical solution

Envelope theorem for the self-adjoint problem: $d\Phi/d\gamma_e = \alpha'(\gamma_e)$ times the cell integral of $-u^2$, with $\alpha(\gamma) = \alpha_{\max} q (1 - \gamma)/(q + \gamma)$.

Result

Category cross_check

Tolerance / **bar** gradient to 1e-4 of finite differences; monotone descent; volume constraint exact; connected channel

Measured result gradient agrees to 1e-7; dissipation 124.7 to 15.9; volume exact; 8.5 percent grey cells

Verdict **PASS**

Traceability MF.13 TopologyOpt | Stokes dissipation minimisation on a pipe bend: gradient vs finite differences, monotone descent, volume constraint, connected channel (FlowSolver2D)

Honest caveat. One geometry on a COARSE mesh, no design filter or projection (so the design is mesh-dependent, may carry checkerboard cells and 8.5 percent are grey), one RAMP parameter, Stokes only, 2D, dissipation objective only. The optimum is NOT compared with a published Borrvall-Petersson result and it is not claimed to be the global optimum: OC is a local method and a different start or volume fraction could give a different topology. The objective uses the solver's Laplacian viscous form, which equals the stress form only for divergence-free Dirichlet data (using the stress form first left a 0.3 percent gradient mismatch, which is how that was found).

6.17 MF.14: DFG 3D-2Z ($Re = 100$, $\sin(\pi t/8)$ inflow) run end to end on a coarse mesh (16,960 tets) – cD_max 4.146 is 24-26% above the published code-vs-code band and cL_max 0.244 is about 100 times the published values: a disagreement

Purpose

This is the first end-to-end run of the unsteady three-dimensional cylinder benchmark at Reynolds number 100 with a time-ramped inflow, using the solver configuration selected by the cost gate (exact velocity block with a reusable factorisation, pressure convection–diffusion Schur approximation with the pressure condition on the outflow only, one linearised Newton step per time step, backward Euler, 800 steps). It answers two questions: does the solver complete the benchmark, and do its drag and lift maxima agree with the values computed by other codes? The first answer is yes: every step converged, in 2.97 hours and 3.5 GB. The second is no: the drag maximum is 24 to 26 per cent above the published range and the lift maximum is about a hundred times the published values. The case is therefore registered as a disagreement. It would show a wrong drag or lift extraction, a mesh that is too coarse, a time scheme that is too dissipative or a linearisation that is too crude; it does not say which, because those three were not separated.

Literature

The benchmark is that of Schäfer and Turek [35]. The values compared with are those FeatFlow publishes for three codes (FeatFlow, OpenFOAM and CFX) at two mesh levels; they are a comparison between codes, not a certified interval, and the same source states that the case has not been rigorously studied. The case definition was read from FeatFlow’s configuration page rather than from the original chapter. The pressure preconditioner follows Kay, Loghin and Wathen [21].

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar cD_max within 5% of the 3.2877-3.3396 band (pre-committed); cL reported without a pass/fail

Measured result cD_max 4.146 (+24 to +26%), cL_max 0.244 (about 100x), cL_min -0.0106 (inside the published range); 800/800 steps converged, 2.97 h

Verdict **MISS (documented)**

Traceability MF.14 solve_flow_2d_transient | newton_mode=:linearized: exact for a linear problem, default untouched, fewer Newton iterations, error shrinking with dt (FlowSolver2D)

Honest caveat. A DISAGREEMENT with a weak reference. The reference is code-vs-code (three codes, two mesh levels) from FeatFlow’s results page, not a certified interval, and that site says the case has not been rigorously studied. The mesh is coarser than theirs (16,960 tets against 49,152 cells) and no finer mesh was run (an lc 0.035 run is projected at about 13 h), so mesh

convergence is unknown; backward Euler is first order and the step is linearised (accuracy established on DFG 2D-3 only), and none of the three contributions (mesh, time scheme, linearisation) was separated. The force region excludes 2 mm at each end of the cylinder while the coefficient is normalised by the full height, which lowers rather than raises c_D . c_L is at the level of discretisation noise for these meshes. The DFG 3D-2Z definition was read from FeatFlow’s configuration page (mean velocity 1.0 at the peak, $c_D = 2F/(\rho U^2 D H)$), not from the primary paper. No experiment or COMSOL comparison.

6.18 MF.14: DFG 3D-2Z measured go/no-go – transient PCD-preconditioned Krylov is correct and its iteration count is flat, the scalable AMG velocity block fails at $Re = 100$, the full-Newton exact-LU path projects to about 32-37 h against a 24 h bar (NO-GO as built), and one linearised step per time step brings it to about 12-14 h (GO on cost)

Purpose

The second unsteady three-dimensional cylinder benchmark of Schäfer and Turek starts from rest and ramps the inflow to a Reynolds number of 100. It had never been run: the direct sparse factorisation is too costly in three dimensions, and the block-preconditioned Krylov solve failed at Reynolds numbers of about fifteen or more. This case is a measured go/no-go with a verdict rule fixed before any run, not the benchmark: it checks that the transient block-preconditioned solve reproduces the direct solution, measures how the iteration count depends on the pressure boundary convention, on the time step and on the mesh, and measures the memory and time per step at the mesh and time step declared benchmark-grade. It would catch a preconditioner that is fast but wrong, a pressure condition that makes the iteration count grow as the time step shrinks (the steady-state convention does), and a cost that puts the benchmark out of reach.

Literature

The benchmark is that of Schäfer and Turek [35]; the geometry was taken from that source in the three-dimensional steady case, but no reference table for the unsteady case is available here, so no drag or lift coefficient is compared. The preconditioner is the block-triangular form with a pressure convection–diffusion Schur approximation [21, 38], whose pressure boundary conditions and the time-step term are the subject of the measurement.

Analytical solution

Not yet characterized – pending review.

Result

Category diagnostic

Tolerance / bar iteration count per solve $j = 60$ in steady state; 800 x wall per step $j = 24$ h; peak RSS $j = 24$ GB (pre-committed)

Measured result correctness $3.4e-15$; 36 iterations (outflow-only) vs 212 (default) at dt 0.01; exact-LU 48-55 iterations at lc 0.035; AMG fails; full Newton 142 s/step - $j \sim 32$ -37 h projected; linearised step (:mf14_linearized_transient_step) 52-62 s/step - $j \sim 11.5$ -13.7 h; 10.3-10.5 GB

Verdict **FINDING**

Traceability MF.14 StokesPreconditioner | VelocityLUCache: a reused (stale) velocity-block LU changes cost, not the answer (FlowSolver2D)

Honest caveat. A NO-GO on cost for the full-Newton solver, with a margin of only 1.3-1.5x; about 60% of a step is Newton assembly, so the linearised step (one Newton step per time step, built afterwards and registered in `:mf14_linearized_transient_step`, accuracy checked on DFG 2D-3 only) was measured at `lc 0.035`: 52-62 s per step, about 11.5-13.7 h for 800 steps, 10.5 GB, inside the bar (timed on a machine shared with another job). The flow state is an IMPULSIVE START at full $Re = 100$ inflow (peak-magnitude velocities), not the developed wake of the real ramp, so iteration counts in the recirculating wake are not measured; the coarse-dt ramp itself failed for every block (`dt 0.1` is a weaker mass shift than the benchmark's). Timings were taken on a machine loaded with three to five other Julia processes (+30%). The 24 h / 24 GB bar and the benchmark-grade mesh and `dt` are MY declaration, not the paper's. No `cD` or `cL` was computed and no 3D-2Z reference table is available in this repository, so nothing here validates a result. Only the exact-LU block was measured at `lc 0.035` (one mesh pair); 3D-3Z was not attempted.

6.19 MF.14: linearised transient step (`newton_mode=:linearized`, one Newton step per time step) – exact for a linear problem, default untouched, and on DFG 2D-3 its `cD` error is 0.19 of the backward-Euler scheme's own temporal error

Purpose

The unsteady three-dimensional cylinder benchmark is too expensive with a fully converged Newton solve at every time step, and most of a step is spent assembling the Jacobian. This case checks the cheaper alternative, a linearly implicit step that takes one Newton step per time step from the previous solution, and asks the one question that decides whether it may be used: is its error smaller than the error the time discretisation already makes? For a linear problem the single step must be exact. On the unsteady two-dimensional cylinder flow with a time-ramped inflow, the difference from the fully converged solution, compared with the difference between the same converged scheme at the time step and at half of it, is the criterion, fixed before the runs. It would catch a linearisation that adds an error larger than the scheme's own, an accepted step that increases the residual, and any change to the default fully converged behaviour.

Literature

Backward Euler with a converged Newton iteration is the reference the linearised step is compared with; no published result exists for the linearised variant on this problem, so the comparison is against the fully converged solution of the same scheme rather than against a literature value. The flow problem is the time-ramped cylinder flow of Schäfer and Turek [35], used here for its unsteadiness and its Reynolds number of up to 100, not for its reference coefficients.

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar linearisation error $\bar{\epsilon}$ = the scheme's own temporal error and $\bar{\epsilon}$ = 1% of peak cD; exact for a linear problem; default bitwise unchanged

Measured result cD error 2.4e-4 vs temporal error 1.3e-3 (ratio 0.19), 100 vs 200 Newton iterations on DFG 2D-3; 7 of 7 unit checks

Verdict **PASS**

Traceability MF.14 solve_flow_2d_transient | newton_mode=:linearized: exact for a linear problem, default untouched, fewer Newton iterations, error shrinking with dt (FlowSolver2D)

Honest caveat. First-order only: with $\theta = 0.5$ (Crank-Nicolson) one Newton step does not preserve second-order accuracy, and the accuracy was established with backward Euler ($\theta = 1$). Verified on ONE 2D problem over t in $[0, 2]$ (Re up to about 71); the 3D use rests on that plus a single-solve equivalence, and a first unit-test criterion (error below the temporal error) failed on a start-up ramp because backward Euler is nearly exact for a ramp linear in time, so that criterion was replaced there by a convergence-with-dt check and kept for the DFG 2D-3 driver where it is meaningful. Compared against fully converged Newton of the same scheme, not against a reference table. No experiment or COMSOL comparison.

6.20 MF.15 (recorded MF.16): DFG 3D-2Z re-run at a finer mesh ($lc=0.035$, 47,088 tets) than the earlier $lc=0.05$ gate – cD_max disagreement with the published band shrinks from +24.2% to +7.1%, a real convergence signal, though two mesh points cannot fix an order

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar cD_max within 5% of the 3.2877-3.3396 band (pre-committed, same bar as :mf14_dfg3d2z_run_lc0p05)

Measured result cD_max 3.522 (+7.1%, improved from +24.2% at $lc=0.05$); cL_max 0.109 (improved from 0.244); cL_min -0.0024 (right sign, ~4-5x too weak)

Verdict **MISS (documented)**

Traceability MF.14 solve_flow_2d_transient | newton_mode=:linearized: exact for a linear problem, default untouched, fewer Newton iterations, error shrinking with dt (FlowSolver2D)

Honest caveat. Still a registered MISS against the pre-committed 5% bar, though closer. Two mesh points (lc 0.05, 0.035) cannot determine a convergence order; the naive $p=2$ extrapolation undershooting the band suggests the true order is lower than 2 or another error source (time scheme, linearisation) still contributes. No new run this session – writing up a result the finisher already produced on 2026-09-25.

6.21 MF.16: the Borrvall-Petersson (2003) double-pipe topology switch (two separate pipes at low aspect ratio, merging into one channel at $\delta=1.5$) was NOT reproduced under three different q-continuation schedules on a self-consistent double-pipe geometry – primary-source port coordinates/objective values could not be confirmed

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar the $\delta=1.5$ design should show a single merged fluid channel (gap_all_solid=false at the vertical midline)

Measured result $\delta=1.0$: two separate pipes (grey 0.089, $q=0.1$); $\delta=1.5$: two separate pipes under all 3 q-schedules (grey 0.097/0.005/0.14) – switch NOT reproduced

Verdict **MISS (documented)**

Traceability MF.13 TopologyOpt | Stokes dissipation minimisation on a pipe bend: gradient vs finite differences, monotone descent, volume constraint, connected channel (FlowSolver2D)

Honest caveat. The most likely explanation is that the self-consistent port geometry (width/position) differs enough from the paper's own unreproduced exact values that the double-pipe's well-known bistability doesn't land on the merged solution here; a systematic port-geometry sweep was not attempted (out of this session's remaining budget). Not a claim that CFDCore's topology optimiser cannot reproduce the switch under any geometry.

6.22 MF.16: one 3D leg of eel_500um Q0.5 (MF.9's own low-shear case) solved under Carreau-Yasuda blood viscosity through the production sweep's own mesh/solve path – mean speed deviation from water is 1.6% (rms 2.0%, max 7.3%), smaller than MF.9's 2D finding (18.9%) for the same nominal case

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar both solves newton_converged with newton_iterations $i=1$ (stop rule); comparison reported as a measurement, no pass/fail threshold

Measured result water: 4 Newton iterations, converged; Carreau-Yasuda: 3 Newton iterations (5-stage continuation), converged; 200/200 probe points resolved; mean 1.6%, rms 2.0%, max 7.3% speed deviation

Verdict FINDING

Traceability MF.13 Rheology | shear-thinning matters where shear is LOW (MF.9 9d mechanism, 1D Carreau-Yasuda vs Newtonian at the plateau viscosity) (FlowSolver2D)

Honest caveat. One leg, not a re-run of the sweep under blood rheology (out of this session's scope). No particle was traced through this leg – the $D=7/15$ tracer drag reading is a stated future step, not done here. Discrepancy with MF.9's own 2D finding is reported, not explained.

6.23 MF.16: Pries et al. (1990/1992) hematocrit/diameter blood-viscosity correlations (in-vitro relative viscosity, the Fahraeus tube-hematocrit effect), new to this codebase – reproduces the qualitative Fahraeus-Lindqvist viscosity minimum and a literature anchor ($H_T/H_D \sim 0.65$ at $D \sim 10\text{-}20\mu\text{m}$, H_D 0.3-0.4)

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar H.T/H.D in [0.55,0.80] at D=10/15/20um, H.D=0.3/0.35/0.4 (band around the literature-reported average 0.65, not a point match); qualitative F-L minimum must be an interior minimum, not monotonic

Measured result 9 testsets, 60+ assertions, all pass; fahraeus_tube_hematocrit(15,0.35)=0.671; F-L minimum at D;20um confirmed on a 594-point diameter sweep

Verdict **PASS**

Traceability fahraeus_tube_hematocrit: reproduces the literature anchor -- H.T/H.D ~ 0.65 (averaged) for H.D in [0.3,0.4] at D~10-20um (Pries et al. 1986) (FlowSolver2D)

Honest caveat. No published digit-for-digit reproduction was possible this session (primary-source access failed); the anchor check is a band, not an exact value. Circular-tube correlations applied to rectangular microchannels via Dh – unvalidated for this geometry. mu_plasma defaults to 1.2 mPa.s, not independently sourced. The in-vivo (endothelial surface layer) correction from Pries & Secomb 2005 was NOT implemented (its fetched fragment was too garbled to trust); only the in-vitro relation and the Fahraeus effect are here.

6.24 MF.16: the S7.15 topology-optimisation density filter is NOT promoted into TopologyOpt.jl – a 4-mesh-level study (lc 0.1/0.07/0.05/0.035) at two filter radii shows the filtered design’s mesh-to-mesh difference decreasing at the same rate as the UNFILTERED design’s, not distinguishably better

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar promote only if the filtered design’s adjacent-mesh difference decreases with refinement while the unfiltered one does not

Measured result R=0.1: filtered [0.0433,0.0295,0.0204] vs unfiltered [0.0371,0.0339,0.0182]; R=0.15: filtered [0.0443,0.0293,0.0217] – both radii FAIL the promotion bar

Verdict **MISS (documented)**

Traceability MF.13 TopologyOpt | Stokes dissipation minimisation on a pipe bend: gradient vs finite differences, monotone descent, volume constraint, connected channel (FlowSolver2D)

Honest caveat. A genuine negative result on ONE problem (the pipe-bend of test_mf.topopt.jl) at two filter radii and four mesh levels up to lc=0.035 (2044 cells) – not a proof the filter never helps, and not compared with a published mesh-independence study. TopologyOpt.jl is not edited; the filter remains a sandbox artefact (mf15_topopt_filter.jl).

6.25 DFG 2D-2 (periodic vortex shedding): `cD_max`/`cL_max`/`St` vs FeatFlow’s tabulated exemplary band

Purpose

DFG 2D-2 is the standard test of an unsteady solver: periodic vortex shedding behind a cylinder at $Re = 100$. It checks the time integrator, the resolution of the wake and, through the shedding frequency, the coupling of the solution’s phase to the numerics. The run was chunked and checkpointed with online periodicity detection, and a real statistics-contamination bug (a perturbation spike misread as a periodic-state force) was caught and fixed.

Literature

Schäfer and Turek [35] define the case; the comparison is against the tabulated exemplary band of the FeatFlow reference solution associated with that benchmark.

Analytical solution

No closed form; FeatFlow’s own tabulated, numerically-converged exemplary results for unsteady periodic vortex shedding.

Result

Category `certified.literature`

Tolerance / bar FeatFlow’s own tabulated exemplary values (fetched live this session, explicitly the wider ‘exemplary’ band, not the tighter secondhand intervals this session could not independently confirm against a primary source): `cD_max` in $[3.1626, 3.2200]$, `cL_max` in $[0.9805, 0.9859]$, `St` in $[0.29411, 0.30303]$

Measured result `cD_max`=3.229 (0.3% above the upper bound); `cL_max`=0.9985 (1.3% above the upper bound); `St`=0.301 (mean of last 5 periods: $0.30303/0.29412/0.30303/0.30303/0.30303$, each landing almost exactly on one of the two band edges)

Verdict **FINDING**

Traceability `MF.9 reaction_force_transient` | reduces exactly to `reaction_force`; isolated uniform-field closed-form check (`FlowSolver2D`)

Honest caveat. `cD_max` and `St` are essentially on the reference band’s edge – strong independent confirmation `solve_flow_2d_transient`’s Crank-Nicolson integration reproduces genuine vortex-shedding dynamics, not just DFG 2D-3’s smoother non-oscillatory response. `cL_max` sits 1.3% over the upper bound – a small, disclosed gap consistent with DFG 2D-3’s own finding that shedding-coupled quantities carry a real resolution-dependent sensitivity relative to FeatFlow’s much finer Q2/P1-discontinuous reference discretization, not evidence of a defect (mass conservation held at/near machine precision throughout).

6.26 DFG 2D-3 (transient, ramped inflow): $cD_{\max}$ to 0.10%; $cL_{\max}/dp(t=8)$ disclosed resolution-sensitivity misses

Purpose

DFG 2D-3 tests the transient solver with a time-ramped inflow at Re up to 100. The maximum drag coefficient is matched to 0.10% of the reference. The lift maximum and the pressure drop at $t = 8$ are disclosed misses, traced to resolution sensitivity near a zero crossing rather than to a sign convention, and reported as measured.

Literature

Schäfer and Turek [35] define the case and tabulate the reference values.

Analytical solution

No closed form; FeatFlow’s own tabulated finest-level point values for the time-ramped-inflow transient case.

Result

Category certified literature

Tolerance / bar FeatFlow’s own finest-level point values (not certified intervals): $cD_{\max}=2.9437637214$ at $t=3.9365625$; $cL_{\max}=0.47748781595$ at $t=5.6928125$; $dp(t=8)=0.11154138872$

Measured result $cD_{\max}=2.9468$ at $t=3.930$ (0.10% relative error, timing matches to 0.007 time units); $cL_{\max}=0.4147$ at $t=5.810$ (13.2% relative error); $dp(t=8)=-0.0986$ (sign-flipped vs $+0.1115$, see honest caveat – both values are small relative to the trace’s own peak amplitude of ~ 2.3)

Verdict **FINDING**

Traceability MF.9 `solve_flow_2d_transient` | DFG 2D-3 ($Re=100$, time-ramped inflow): cheap correctness-only regression (FlowSolver2D)

Honest caveat. The transient solver’s dominant, well-resolved response ($cD_{\max}$ and its peak timing) is quantitatively validated to sub-percent accuracy against an independent reference implementation. Its response near a zero-crossing (dp at $t=8$) or during vortex-shedding- influenced peaks ($cL_{\max}$) is directionally/qualitatively correct but not tightly matched at this run’s resolution – a genuine, disclosed limitation, not evidence the machinery is broken (the exact-machine-precision mass conservation throughout all 800 steps is strong independent evidence the underlying discretization itself has no defect).

6.27 MF.9 (registered MF.13): shear-thinning changes the velocity field most where shear is LOW – the MF.9 rheology finding’s mechanism, reproduced on the 1D Carreau-Yasuda solver

Purpose

MF.9 found that a blood-like shear-thinning fluid changes the velocity field by more in the wide, low-shear channel than in the narrow, high-shear one, the opposite of the naive expectation. This

case checks the mechanism cheaply and independently on the one-dimensional Carreau–Yasuda solver: the profile approaches the Newtonian one at both shear limits and departs most in between, so the deviation from the plateau-viscosity profile falls monotonically with shear (0.55%, 4.3%, 19%). It does not re-run the finite-element sweep and its percentages differ from that sweep’s by construction, because the reference here is the plateau viscosity, not water.

Literature

The blood viscosity parameters are the Cho and Kensey fit [8] used as this study’s defaults; the Carreau–Yasuda form itself is a standard generalised-Newtonian model whose original source was not re-verified in this session and is not cited.

Analytical solution

No closed form for the Carreau-Yasuda channel profile; the reference is the Newtonian plane-Poiseuille profile at the plateau viscosity.

Result

Category `cross_check`

Tolerance / bar deviation decreases monotonically with shear, by more than 5x and 3x between the three cases

Measured result 0.55 percent, 4.3 percent, 19 percent; steps 7.9x and 4.4x

Verdict **FINDING**

Traceability MF.13 Rheology | shear-thinning matters where shear is LOW (MF.9 9d mechanism, 1D Carreau-Yasuda vs Newtonian at the plateau viscosity) (FlowSolver2D)

Honest caveat. This is the MECHANISM on a 1D plane channel, not a re-run of the finite-element sweep: the reference here is the Newtonian plateau viscosity (the sweep’s reference was water at $8.9\text{e-}4$ Pa s, so its percentages differ from these by construction), the geometries are straight channels standing in for the eel channels, and the sweep itself (two 2D cells, three models) is neither re-run nor asserted in the suite. Casson is not re-checked here. It supports the direction and the physical reason (the profile approaches Newtonian at both shear limits and departs most in between), not the sweep’s numbers.

6.28 DFG 3D-1Z (Schäfer-Turek, steady $\text{Re}=20$, 3D circular cylinder): $c_D/c_L/dp$ vs certified interval, mesh-convergence campaign

Purpose

`solve_flow_3d` and `reaction_force` were validated through MF.10 only against synthetic exact solutions (Hagen–Poiseuille, the rectangular-duct Fourier series) and mesh/tag-completeness checks – none of those exercise a genuinely 3D, non-axisymmetric flow past a bluff body, the case that actually stresses the P2/P1 Taylor–Hood discretization’s ability to resolve boundary-layer separation and a delicate drag/lift asymmetry in three dimensions. This benchmark is the first time this solver’s 3D path is asked to reproduce a community-converged answer on a case nobody wrote the

equations for by hand – it would catch a 3D-specific discretization defect (an incorrectly assembled 3D stress tensor, a wrong reaction-force sign convention when a boundary indicator is 3D, an under-resolved separation bubble) that no synthetic-solution check could, since a manufactured or closed-form solution is built to be exactly representable by the discretization and therefore cannot expose a resolution defect the way a genuine bluff-body wake can.


Literature

Schäfer & Turek (1996) [35] is the DFG Priority Research Programme’s own laminar-flow-around-a-cylinder benchmark suite – the same primary source already underlying `fs_cylinder_dfg2d1`’s 2D case, extended here to its own dedicated 3D test case (3D-1Z, Table 7 of that paper). It is the right source specifically because it reports the community’s converged answer across 10 independent research groups and solution methods (finite element, finite volume, finite difference, lattice-Boltzmann; Table 1 of the paper), not one group’s own run – exactly the “certified reference interval is stronger evidence than a single vendor’s single run” standard `MF_VALIDATION_PROGRAM.md` §0 states this whole program holds itself to.

Analytical solution

No closed form; Schafer & Turek (1996)’s certified 3D reference interval (Table 7), the same primary source as `fs_cylinder_dfg2d1`’s 2D case.

Result



c_D converges INTO the certified interval [6.05,6.25] with refinement (6.40 coarse → 6.12 fine). Δp sits just below [0.165,0.175] at both levels. c_L is highly resolution-sensitive and changes SIGN between the two meshes tested (+0.093 → -0.012 vs reference [0.008,0.01]) -- not yet converged. A region-indicator hypothesis (xy-only vs z-restricted) was tested and ruled out as the cause (both agree to ~2%).

Category certified_literature

Tolerance / bar c_D , c_L , Δp against DFG_3D1Z_REFERENCE’s certified interval, ASME V&V 20-style

Measured result coarse ($l_c=0.05$): $c_D=6.396$ (2.3% above [6.05,6.25]), $c_L=+0.093$ (order of magnitude high vs [0.008,0.01]), $\Delta p=0.166$ (inside [0.165,0.175]). Finer ($l_c=0.035$): $c_D=6.124$ (INSIDE the certified interval), $c_L=-0.012$ (sign-flipped from the coarse mesh, now the same order of magnitude as the reference but wrong sign — not yet converged), $\Delta p=0.162$ (just below the interval). `mass_conservation_residual` at machine precision ($\approx 1.4e-17$) at both levels; both solves converged in 4 Newton iterations.

Verdict FINDING

Honest caveat. cD converges toward (and lands inside) the certified interval with mesh refinement, matching the established 2D DFG-1 pattern. cL does NOT — it is far more mesh-sensitive than cD or dp at both tested resolutions, even changing sign between them, and has not converged to a stable value at either. This is consistent with the primary source’s own documented characteristic (Table 7 itself shows independent research groups’ cL values spanning a wide range across their own mesh levels for this exact case) rather than a bug in this solver — a near-cancelling small coefficient (drag/lift asymmetry around a 3D cylinder is a delicate residual of much larger normal forces) is intrinsically far more resolution-sensitive than cD or dp. Neither mesh level’s cL should be read as a validated result; a substantially finer mesh (beyond this session’s compute budget — the finer level already took ~1230s to solve, and A.0’s own $\text{DOF}^{1.79}$ cost-scaling finding means the NEXT refinement level would cost several times that) would be needed to see whether cL actually converges toward the certified interval or reveals a genuine, different discretization issue. DFG 3D-2Z/3Z (both unsteady) remain entirely unattempted — explicitly out of this session’s scope, and per A.0’s own feasibility finding, likely infeasible at production mesh resolution without the saddle-point-preconditioned solver this study’s Open Items already name as a prerequisite. UNEXPLAINED DRIFT (S7.13/MF.13, 2026-09-21): re-running the ORIGINAL mf_dfg3d1z.jl at lc = 0.05 (same 16,960 cells) on the current toolchain gives cD = 6.3267, cL = +0.0958, dp = 0.16628, i.e. cD 1.1 percent below the 6.396 recorded above; the driver, viscous form, continuation and region indicator are unchanged and the cheap regression testset still passes. Julia 1.11.1 was used, while the package Manifest was resolved under 1.10.11; a sparse-LU difference alone would move cD by about 1e-10, so no cause is claimed. The verdict is unchanged (both values sit above the certified interval at the coarse level, inside the indicator spread disclosed above); the finer-mesh values were not re-run.

7 FocusingMetrics

7.1 Curved-duct equilibrium-position-vs-diameter: non-monotonic, not a clean separation trend

Purpose

Five particle sizes of twenty particles each are traced through the realistic-scale curved duct under the full composite force model with a confirmed nonzero Du/Dt . The recorded outcome is that equilibrium position versus diameter is non-monotonic and not a clean separation trend. This is reported as measured: it bounds what the two-dimensional composite model can claim about curved-channel separation and motivates the three-dimensional work.

Literature

The size-dependent lateral equilibria that such a curved-channel device is expected to produce are those described by Di Carlo et al. [11]; this result is a measurement of the model, not a comparison with that paper’s data.

Analytical solution

No closed form; a measured, non-monotonic equilibrium-position result reported as a finding, not compared against an external analytical reference.

Result

Category diagnostic

Tolerance / bar no monotonicity or tight-spread bar was written into the test, because the data does not support one – this is the point of the finding

Measured result $d=[5,10,15,20,25]\mu\text{m}$ - $i_{s_eq}=[0.393, 0.554, 0.425, 0.425, 0.553]$; $\sigma=[0.24,0.30,0.29,0.25,0.24]$

Verdict **FINDING**

Traceability MF.4 | curved-duct case: Du/Dt, streamline, and equilibrium-position ensemble (items 6,15,16) (FocusingMetrics)

Honest caveat. Five size classes (5-25 μm) x 20 particles each, released near the inlet of a realistic-velocity-scale (~ 0.5 m/s at mid-radius, matching MF.0’s own operating point) curved duct, traced under the full composite force model with a confirmed-nonzero Du/Dt (0.095% relative error vs analytic centripetal acceleration). Measured final radial fractions: $d=[5,10,15,20,25]\mu\text{m}$ - $i_{s_eq}=[0.393, 0.554, 0.425, 0.425, 0.553]$, $\sigma=[0.24,0.30,0.29,0.25,0.24]$. This is NOT monotonic in diameter, and the population spread at every size class is a large fraction of the full [0,1] range. A live Du/Dt alone does not, at this operating point, produce a clean equilibrium-position trend for this composite lift model – consistent with the study’s own standing caveat: 2D structurally cannot produce genuine size-based inertial separation (no Dean vortices without a resolved third dimension); the lift model through MF.9 remains a 2D empirical stand-in, and MF.10 is where this becomes physically real. No monotonicity or tight-spread assertion was written into the test, because the data does not support one.

7.2 Residence time vs $L/\text{peak_U}$, and force_budget_summary’s gravity/Brownian ranking metric traps

Purpose

This is a diagnostic finding about how not to read the output. The force-budget summary’s ranking of gravity and Brownian terms against raw time-averaged trajectory drag is physically correct and easy to misread as “gravity or Brownian motion is dominant”. Recording the trap, together with a residence-time check against L/U_{max} , protects every later consumer of the budget from that misreading.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Partial: residence time is compared to the exact kinematic identity $L/\text{peak_U}$; the force-budget-ranking finding itself is a methodological diagnostic, not a solution comparison.

Result

Category diagnostic

Tolerance / bar residence time $\text{rtol}=0.5\%$; force-budget ranking: no pass/fail bar – the finding is about which comparison metric is physically meaningful, not a numeric target

Measured result residence time agreement to 0.5% ; gravity vs characteristic flow-drag scale ($\gamma(\text{peak_U}) \cdot \text{peak_U}$): 0.37% (consistent with MF.0’s own screening); gravity vs raw time-averaged trajectory drag: ~ 1 (misleadingly large, see caveat); Brownian vs drag’s time-averaged magnitude: $\sim 2.9\text{x}$ (misleadingly large, see caveat)

Verdict **FINDING**

Traceability MF.4 | residence time and force budget (items 13-14, straight channel)
(FocusingMetrics)

Honest caveat. Comparing gravity’s magnitude against a particle’s own time-averaged trajectory drag is a misleading metric once the particle reaches terminal settling (drag trivially equals gravity at steady state, by force balance) – the physically meaningful comparison is against a characteristic flow-drag scale ($\gamma(\text{peak_U}) \cdot \text{peak_U}$), against which gravity measures 0.37% , consistent with MF.0’s own screening. Brownian’s force-budget entry (a stochastic kick-over-dt quantity) cannot be ranked on equal footing with deterministic forces – it scales $\sim 1/\sqrt{\text{dt}}$ like any white-noise term, not because Brownian motion is physically dominant (measured at $\sim 2.9\text{x}$ drag’s own time-averaged magnitude for a realistic step size, this is NOT evidence Brownian motion is physically dominant; MF.0’s own screening found the opposite via a DISPLACEMENT-over-transit comparison, $\sim 0.05\mu\text{m}$ against a $100\text{--}500\mu\text{m}$ channel). `force_budget_summary`’s docstring now states this caveat explicitly; no assertion is made on `fbs.ranked[1]` for this reason. Any future force-budget consumer (this registry’s own report generator, MF.6’s GUI) should use the characteristic-drag-scale and MF.0-style displacement-budget comparisons, not raw time-averaged-force ratios, when characterizing stochastic or settling-dominated terms.

7.3 FocusingMetrics synthetic validation: focusing index, separation resolution, outlet purity

Purpose

The focusing and separation metrics are pure arithmetic on particle ensembles, so they can be validated exactly against hand-constructed populations whose answers are known by construction: a delta-function final population gives a focusing index of exactly one, an unchanged population gives exactly zero, and a spread population gives a negative index. This would catch an inverted definition, a wrong normalisation or a purity calculation that miscounts the outlet.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Yes: synthetic, hand-constructed particle-position data with an exactly known focusing-index/purity/resolution answer.

Result

Category closed_form

Tolerance / bar atol=1e-12 throughout (synthetic inputs with exact known answers)

Measured result focusing_index: 1.0/0.0/negative exactly; separation_resolution: 0.0 exactly, closed form to atol=1e-12; outlet_purity: 1.0/1.0 and 0.5/0.5 exactly

Verdict **PASS**

Traceability MF.4 FocusingMetrics | synthetic validation (items 10-12) (FocusingMetrics)

7.4 MF.16: FocusingMetrics.jl's focusing/separation/purity functions, never applied to the MF.12 A.7 sweep before this session, run over all 45 v4 cells – separation resolution Rs is generally low (0.001-0.13 for adjacent size classes), and the D=15um depth-equilibrium headline is extended to all 9 channels

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category cross_check

Tolerance / bar Julia-;Python metric agreement to 1e-12 relative; no pass/fail on Rs itself (a measurement, not a gate)

Measured result focusing_index/separation_resolution/outlet_purity match Julia to 1e-15/1e-16 relative; 45/45 cells processed; Rs 0.001-0.13 (di_carlo_3d, adjacent size-class pairs)

Verdict **FINDING**

Traceability MF.4 FocusingMetrics | synthetic validation (items 10-12) (FocusingMetrics)

Honest caveat. n=20 particles per class, rigid spheres, water viscosity, one-way coupling – see :mf16_blood_validity_map for how far that is from real hematocrit. Every cell is tagged with its own leg design (baseline vs long-leg/stop-plane); a 'sharper turns' conclusion is only drawn within the same design (serpentine_sharp_100um vs the eel/serpentine_500um baselines), never across the design boundary. Best-split purity/recovery uses a grid search over candidate thresholds (200 points), not an analytic optimum.

7.5 MF.9 (registered MF.13): two-way-coupling thresholds – Elghobashi volume-fraction bound and the direct force-balance criterion – recomputed in Julia from the stored assessment data

Purpose

MF.9 asked at what real particle concentration one-way coupling stops being defensible. Its answer was an analysis script, which the registry could not trace to a test. This case ports the two closed-form criteria (Elghobashi’s volume-fraction bound and a direct momentum-feedback ratio built from measured particle drag) and recomputes every stored threshold from the assessment’s own inputs. It protects the arithmetic behind the reported finding: the two criteria disagree by 4.8 to 6.2 orders of magnitude, so isolated rare cells are one-way defensible by both, whole-blood concentrations by neither, and white-cell-scale suspensions are the ambiguous middle.

Literature

The volume-fraction classification is Elghobashi’s [13]; its map is defined for turbulent flow, and this device is laminar, so it is used as a reference point and not as a bound. The force-balance criterion is a flow-regime-agnostic ratio defined in MF.9 itself and has no published counterpart.

Analytical solution

Closed forms: $n = \phi / ((\pi/6) d^3)$ and $n = f(\rho U^2/H) / \text{drag_avg}$; the drag input is measured.

Result

Category `cross_check`

Tolerance / **bar** thresholds recomputed from the stored inputs to 1e-9 relative

Measured result 34/34 checks: all Elghobashi and force-balance thresholds reproduce; the two criteria differ by 4.8 to 6.2 orders of magnitude across sizes

Verdict **FINDING**

Traceability MF.13 Coupling | MF.9’s two-way-coupling thresholds recomputed from the stored assessment JSON (FocusingMetrics)

Honest caveat. This anchors the ARITHMETIC of the two criteria, not the physics conclusion. Elghobashi’s map is defined for TURBULENT flow (second axis τ_p over the Kolmogorov time) and this device is laminar, so it is a widely cited reference point, not a rigorous bound. The force-balance threshold takes its drag from MF.8’s force budgets, which are read from the stored JSON and NOT recomputed here; the physiological concentrations and the 1 percent fraction are the assessment’s own choices. No coupling is implemented (one-way by the methodology document’s scope); the particle-to-fluid momentum source remains absent.

8 MicroMesh

8.1 3D tagged gmsh meshes: box/pipe/curved-duct/extruded-polygon/DFG-cylinder, all tag-complete

Purpose

The 3D meshers (box, pipe, curved duct, extruded polygon and the DFG cylinder) are the foundation of every 3D result. Each must produce a watertight, tag-complete tetrahedral mesh, verified by set equality on the tag list, exactly mirroring the two-dimensional check. A missing tag or a leaking surface would make a 3D benchmark fail for a reason unrelated to the physics.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

No closed form; validated via mesh topological/quality invariants (tag completeness, watertightness), the 3D extension of `mm_tagged_shapes`'s own check.

Result

Category `cross_check`

Tolerance / `bar` tag-set equality (exact); `min_sicn` ≥ 0 (no inverted/degenerate tets)

Measured result 5/5 mesh functions tag-complete and watertight

Verdict **PASS**

Traceability MF.10 MicroMesh | 3D tagged gmsh meshes (MicroMesh)

8.2 Tagged-mesh completeness, watertightness, and quality across 5 literature-benchmark shapes

Purpose

Boundary conditions are attached to mesh entities by name, so a mesh that is missing a tag (or labels the wrong faces) fails silently downstream: the solver simply applies no condition where one was intended. This check confirms by exact set equality that each of the five literature-benchmark shape wrappers, and the generic polygon entry point, produces the complete expected tag set and elements of positive quality, so a benchmark that fails later can be attributed to physics or discretisation rather than to a mislabelled boundary.

Literature

The shapes are the standard geometries of the backward-facing-step benchmark [15] and of the cylinder-in-channel benchmark [35]; the checks themselves are topological and reference no numerical value.

Analytical solution

No closed form; validated via mesh topological/quality invariants (tag completeness, watertightness), not a solution comparison.

Result

Category `cross_check`

Tolerance / bar tag-set equality (exact); min_sicn $\leq$ 0.05 (0.02 for the cylinder mesh, tighter geometry near the circular boundary)

Measured result 6/6 shape wrappers tag-complete and watertight; 18 MicroMesh tests green (MF.0 test count)

Verdict **PASS**

Traceability MicroMesh | tagged gmsh meshes for literature-benchmark shapes (MF.0) (MicroMesh)

9 MicroVTK

9.1 VTK/PVD file round-trip: written files parsed back, not merely checked to exist

Purpose

Exported VTK/PVD files feed the visual inspection that catches what numbers miss, so a silent export defect would corrupt every figure. The check writes a file, parses it back, and asserts an invariant (the point count strictly exceeds the mesh node count because P_2 fields are subdivided on P_1 geometry), rather than merely checking that the file exists.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Not applicable – a file-format round-trip check (write then re-parse); no analytical solution is involved.

Result

Category `cross_check`

Tolerance / bar non-trivial file size; correct internal structure (point counts strictly greater than mesh nodes, expected field names present, .pvd DataSet-entry count == n_frames)

Measured result .vtu/.vtp/.pvd all written correctly-structured on a real solved flow field; a real bug (MeshCell[] abstractly-typed vector defeating WriteVTK's polydata dispatch) was caught by this smoke test before it ever reached runtests.jl

Verdict **PASS**

Traceability MF.4 MicroVTK | file round-trip (item 17) (MicroVTK)

10 ParticleBenchmarks2D

10.1 Di Carlo (2007) a^4 wall-lift scaling and Saffman-Mei d^2 low-shear-Re scaling recovered by regression

Purpose

Two different lift mechanisms have different size exponents, and conflating them is a classic error. This check regresses each against its own correct exponent: the Saffman–Mei low-shear-Reynolds limit is regressed by fixing position and sweeping diameter, isolating the scaling from trajectory effects, while the wall-induced lift is checked against its a^4 dependence. It would catch an exponent applied to the wrong length or a coefficient that hides a size dependence.

Literature

The a^4 scaling of the inertial lift is the result of Di Carlo et al. [11]; the low-Reynolds shear-lift scaling is Saffman’s [34] with Mei’s correction [24].

Analytical solution

No closed form; Di Carlo (2007)’s empirically-fitted a^4 scaling law, itself an empirical correlation, recovered here by regression.

Result

Category regression

Tolerance / **bar** Saffman-Mei log-log slope: atol=0.3 around the hand-derived $F \sim d^2$ exponent ($V_p \sim d^3$, $1/\sqrt{\text{Re} \cdot \omega} \sim 1/d$, net d^2); Di Carlo slope: atol=0.05 around DI_CARLO_2007_SCALING.lift

Measured result Saffman-Mei log-log slope 2.0 +/- 0.3 (matches the hand-derived exponent); Di Carlo a^4 scaling recovered to ± 0.05

Verdict **PASS**

Traceability MF.3 | focusing-position measurement (tests 10–11, production-scale channel) (ParticleBenchmarks2D)

11 ParticleForces

11.1 ParticleForces/ParticleTracer 3D: Saffman-Mei/Cherukat-McLaughlin exact 2D reductions, di_carlo_lift_3d, tracer smoke test

Purpose

The 3D force kernels must reduce exactly to their 2D counterparts when the problem is effectively planar (relative tolerance 10^{-12}), since the underlying physics was always dimension-general. The check also records that the new wall-distance lift correlation is a literature substitution (documented as such) and exercises the 3D tracer end to end, so any drift between the 2D and 3D code paths is caught.

Literature

The reduced kernels are those of Saffman [34], Mei [24] and Cherukat and McLaughlin [7]; the substituted lift profile is the a^4 wall-distance form of Di Carlo et al. [11].

Analytical solution

No closed form; a 3D self-consistency/plumbing smoke test against the already-validated 2D force-kernel formulas.

Result

Category `cross_check`

Tolerance / **bar** 2D-reduction `rtol=1e-12` (exact); `di_carlo_lift_3d` correctly zero at `s=0` (wall) and `s=0.5` (centerline), nonzero at `s=0.25` (stable equilibrium)

Measured result 2D reductions pass exactly; tracer smoke test advects all 4 particles downstream with consistent trajectory/force-budget bookkeeping

Verdict **FINDING**

Traceability MF.10 `ParticleForces` | 3D force-kernel self-checks (`ParticleForces`)

Honest caveat. `di_carlo_lift_3d` is a documented literature substitution (InertialFocusing.jl's existing scaling law + shape, not a fitted DNS/experimental $f_L(\text{Re}, \text{position})$ profile) – see this module's own MF.10 section header and `di_carlo_lift_3d`'s docstring for the full provenance. A.6's Segre-Silberberg pipe trace is the real test of whether this produces genuine literature-quantitative inertial focusing.

11.2 MF.12 Step 4: `:ho_leal_3d_multiwall` wired through `lift_force_3d` and `trace_particles_3d!` – exact reductions and the two-wall factor

Purpose

Guards the wiring of the multi-wall inertial lift into the particle tracer. The wiring touches the tracer's force-evaluation loop, so the failure modes are: a silent single-wall fallback, a changed result for the existing lift models, and a wrong sum over walls. The checks are exact algebraic reductions (dispatch equals the direct multi-wall call; one wall reduces exactly to the single-wall model; the new keyword is a byte-for-byte no-op for the other models; a missing wall list is an error) plus a geometric consequence: two equidistant walls must displace a particle twice as far as one wall.

Literature

Internal consistency check with no external reference; the lift profile it sums is the Ho and Leal (1974) / Segre and Silberberg (1962) collapse already registered as `mfv_ho_leal_normalization_collapse`.

Analytical solution

No closed form for the trajectory; the checks are exact algebraic reductions plus the geometric two-wall factor of 2.

Result

Category regression

Tolerance / bar exact equality for the dispatch/reduction/no-op checks; corner displacement ratio 2.0 within 2%

Measured result force-dispatch testset 8/8; tracer-wiring testset 5/5; corner displacement 1.8216e-10 m vs 9.109e-11 m (ratio 1.9998)

Verdict **PASS**

Traceability MF.12 ParticleForces | :ho_leal_3d_multiwall dispatch (Step 4) (ParticleForces)

11.3 MF.13 Phase 4: acoustic radiation force on a small sphere in a 1D standing wave against closed forms, and node focusing through the real tracer

Purpose

Acoustic radiation force was screened out for this device (no transducer) and is built here on request. The check confirms the kernel against its closed forms: the contrast factor, the peak force, the zeros at nodes and antinodes, node stability for a positive contrast factor and antinode stability for a negative one, half-wavelength spacing, and, through the real tracer in quiescent fluid, migration of a bead to within $0.1\text{ }\mu\text{m}$ of a node from a $5\text{ }\mu\text{m}$ start. It is a kernel-level verification of an analytically prescribed standing wave, not a validation of acoustofluidic separation.

Literature

The contrast factor and force follow Gor'kov [17] in the standing-wave form given by Bruus [5].

Analytical solution

Gor'kov (1962) contrast factor and the 1D standing-wave radiation force (Bruus 2012).

Result

Category closed_form

Tolerance / bar closed-form zeros/peak to 1e-12; node focusing to within 0.1 um from a 5 um start

Measured result 11/11 kernel checks and the tracer focusing test pass

Verdict **PASS**

Traceability MF.13 Phase 4 | acoustophoresis: Gor'kov contrast factor, node/antinode structure, half-wavelength spacing (ParticleForces)

Honest caveat. BUILT ON REQUEST, NOT NEEDED BY THIS DEVICE (MF.0 screened acoustic radiation force out: no transducer). Closed-form verification only, for an analytically PRESCRIBED planar standing wave; no acoustic field solver (the existing HelmholtzInterface is real-valued pressure acoustics and is not coupled), no acoustic streaming, no thermoviscous effects, no viscous-boundary-layer correction, point-particle limit (a much smaller than the wavelength). It carries NO Tier-3 physical-validation evidence.

11.4 MF.13 Phase 4: dielectrophoresis force and Clausius-Mossotti factor against closed forms (limits, crossover frequency, a^3 scaling)

Purpose

Dielectrophoresis was screened out for this device (no electrodes or applied field), and is built here on request as a general capability. The purpose of the check is therefore narrow and stated as such: to confirm that the force and the Clausius–Mossotti factor implement their closed forms exactly (the low- and high-frequency limits, the closed-form crossover frequency at which the real part vanishes and the force changes sign, exact a^3 scaling), so that a later user of the module inherits a verified kernel. It gives no evidence that dielectrophoretic separation works in any real device.

Literature

The force and the complex-permittivity form of the Clausius–Mossotti factor are treated in Jones [20]. The crossover-frequency expression is derived from that form in this study and verified numerically against it.

Analytical solution

Clausius-Mossotti factor and its crossover frequency in closed form (Jones 1995; Pohl 1978).

Result

Category closed_form

Tolerance / **bar** closed-form limits to 1e-6, crossover to 1e-12, exact scaling

Measured result 13/13 checks pass

Verdict **PASS**

Traceability MF.13 Phase 4 | dielectrophoresis: Clausius-Mossotti limits, closed-form crossover frequency, force scaling (ParticleForces)

Honest caveat. BUILT ON REQUEST, NOT NEEDED BY THIS DEVICE: MF.0 screened DEP out (no electrodes or applied field in the thesis device, reports/mf0_physics_relevance_screening.md). Closed-form verification only: no device, no experiment and no electrostatics field solver exist here, the field gradient is an ANALYTICALLY PRESCRIBED input, and the model is the point-dipole limit (a much smaller than the field length scale, no multipole or particle-particle terms, no electrothermal flow). It carries NO Tier-3 physical-validation evidence.

11.5 MF.13: DLVO colloidal interaction – sphere-plate wall adhesion and sphere-sphere forces (van der Waals plus weak-potential double layer): force equals minus the energy gradient, exponent laws, barrier, Newton’s third law

Purpose

A finding of the particle sweep was that every stalled chain is a wall capture, decided in the tracer by a geometric rule with no adhesion physics behind it. This case checks the colloidal interaction forces that a physical adhesion model would use: sphere–plate and sphere–sphere forces from the

sum of a van der Waals attraction and a screened electrostatic repulsion. They are checked against their own energy (the force must be the negative gradient), their exponent laws, the position of the energy barrier and Newton’s third law for a cloud of spheres. It is a kernel-level check: the forces are not yet wired into the tracer’s wall handling, and no trajectory with adhesion has been run.

Literature

The sum of an attractive dispersion term and a repulsive double-layer term as the interaction between charged colloidal surfaces is the theory of Derjaguin and Landau [10]. The sphere–plate and sphere–sphere forms used here are the Derjaguin approximation of that pair interaction, in its weak-potential limit; the closed forms are checked against their own energy gradient rather than against a measured force curve.

Analytical solution

$U(\text{gap}) = -A_H a/(6 \text{ gap}) + 4 \pi \epsilon a \psi_p \psi_w \exp(-\kappa \text{ gap})$ (Derjaguin, weak-potential linear superposition); $F = -dU/d\text{gap}$.

Result

Category closed_form

Tolerance / bar force equals minus $dU/d\text{gap}$ to $1e-6$; exact exponent laws; barrier at the force sign change; pair sum zero

Measured result 22 of 22 checks

Verdict **PASS**

Traceability MF.13 ParticleForces | DLVO wall adhesion and particle-particle interaction:
force = $-dU/dh$, exponent laws, barrier, Newton’s third law (ParticleForces)

Honest caveat. KERNEL LEVEL (the wall CONTACT is wired into the 3D tracer as `wall_model=:dlvo` in MF.14, see `:mf14.dlvo.wall_contact`; the force-field factory is still not usable to resolve a contact in the explicit integrator, whose steps are far larger than the Debye length). The pair sum is not called by the tracer, and this entry itself changes no recorded result. Validity limits stated in the source: gap much smaller than the radius, weak potentials (below about 25 mV), κ times gap of order one or more, no retardation, hydration or steric force, no near-wall drag correction, smooth rigid spheres. The Hamaker constant and potentials are inputs. Verified only against its own closed-form energy: no measured force curve, no experiment, no COMSOL comparison.

11.6 MF.13 Phase 4: magnetophoretic force against its closed form, and exponential quadrupole growth through the real tracer

Purpose

Magnetophoresis was screened out for this device (no magnet or magnetic particles) and is built here on request. The check confirms the force $V \Delta \chi \nabla(B^2)/(2\mu_0)$, its equivalence to $(\mathbf{B} \cdot \nabla)\mathbf{B}$ for a curl-free field (verified explicitly on a quadrupole), and, through the real tracer in quiescent fluid, the closed-form exponential growth of the particle position under Stokes drag within 3%. It is a kernel-level verification against an analytically prescribed field.

Literature

The force on weakly magnetic particles in microfluidic devices is reviewed by Pamme [28].

Analytical solution

Magnetophoretic force on a weakly magnetic sphere (Pamme 2006; Zborowski et al.); quadrupole $(\mathbf{B} \cdot \text{grad})\mathbf{B}$ in closed form.

Result

Category closed_form

Tolerance / **bar** closed-form force to 1e-12; tracer growth rate within 3 percent of $V \, d\chi / (2 \mu_0 \nabla \cdot \mathbf{G})$

Measured result 5/5 kernel checks; tracer growth rate within 3 percent

Verdict **PASS**

Traceability MF.13 Phase 4 | magnetophoresis: $\mathbf{F} = V \, d\chi / (2 \mu_0) \, \text{grad}(\mathbf{B}^2)$, and the $(\mathbf{B} \cdot \text{grad})\mathbf{B}$ equivalence for a curl-free field (ParticleForces)

Honest caveat. BUILT ON REQUEST, NOT NEEDED BY THIS DEVICE (MF.0 screened magnetophoresis out: no magnet, no magnetic particles). Closed-form verification only, on an ANALYTICALLY PRESCRIBED quadrupole field; no magnetostatics solver, no magnetisation saturation or particle-particle dipole interaction, and the linear (small χ) limit only. It carries NO Tier-3 physical-validation evidence.

11.7 MF.13 Phase 3: a field-directed tangential lift term (grad of shear-rate squared) does NOT robustly select Di Carlo’s face-centres; the converged model leaves a bimodal ensemble

Purpose

Every three-dimensional lift kernel in the study returns a scalar times the wall normal, so lift direction is fixed by wall geometry alone. That can only produce a ring in a square duct (single wall) or corner attractors (summed walls), whereas Di Carlo et al. report four face-centre positions. This case tests whether letting the velocity field set a tangential push, along the gradient of the squared shear rate, can select the face centres. The verdict criterion was fixed before any ensemble ran. The exact (analytic) gradient leaves a bimodal ensemble (about 57% at face centres, 43% on a diagonal), a sign map shows why (the tangential push reverses direction near the diagonal and at larger wall distance), and a clean pass obtained with a coarse sampler disappears when the sampler is refined. It is reported as a negative finding, not a fix.

Literature

The target behaviour is the four-face-centre equilibrium of Di Carlo et al. [11]; the wall-normal part of the model is Ho and Leal’s [18] lift profile. The tangential term is an ansatz of this study and has no published counterpart.

Analytical solution

No closed form for the 3D inertial lift map; the shear-rate-squared gradient is checked against a finite-difference of the closed-form duct solution (relative error 8e-8).

Result

Category diagnostic

Tolerance / bar at least 2/3 of unfrozen particles within 5 degrees of a face-centre axis AND median angle below 10 degrees (fixed before any ensemble was run)

Measured result analytic gradient: 24/42 within 5 degrees, 16-18/42 on a diagonal, NOT selected at $f_t = 0.5, 1, 2$; sampled H/8: 41/42 (spurious); sampled H/12: 24/42 and 18/42, NOT selected

Verdict FINDING

Traceability MF.13 Phase 3 | `inertial_lift_3d_shear_gradient`: algebra, dispatch, and refusal to substitute (ParticleForces)

Honest caveat. This is an ANSATZ, not a derived lift law: the tangential coefficient f_t is free (pre-set to 1, with 0.5 and 2 declared in advance) and the sign was chosen to point toward stronger shear, so even a pass would show only that such a term is sufficient, not that inertial lift in a square duct has this form. It is NOT a fix for the Di Carlo specificity miss. The clean H/8 sampled result is a coarse-sampler artefact (the piecewise-constant Hessian's discretisation error corrupts the sign structure near the diagonal, where the force is small): refining the sampler to H/12 reproduces the analytic bimodal outcome exactly, and the median angle alone would have badly overstated every one of these results (0 to 1.3 degrees for a 57/43 split). One particle size (10 μm), one Re (22.5), one quadrant ensemble on the exact duct field only (a FEM field's noise was not tested); the H/16 sampler run was stopped unfinished. The B.2 multiwall result (corner attractors) stands.

11.8 MF.13: particle forces on SOLVED fields – DEP centring in a solved quadrupole at the closed-form rate, electrophoretic drift μ_{ep} $\mathbf{E}$, thermophoretic drift from a solved Joule temperature

Purpose

This case joins the field solvers to the particle tracer. Electrophoresis and thermophoresis are added as Stokes-friction forces (a drift $\mu_{ep}\mathbf{E}$ and a drift $-D_T\nabla T$), and dielectrophoresis is run on a field that has been *solved* rather than prescribed. Three checks go through the real tracer in quiescent fluid: a negative-DEP bead released off-centre in a solved quadrupole decays to the field null at the closed-form rate, the electrophoretic drift in a solved uniform field equals $\mu_{ep}E$ with no transverse motion, and the thermophoretic force computed from a solved Joule-heating temperature has the closed-form gradient and points toward the cold wall. These modules were screened out for the device of this study, so the case carries closed-form verification only and no physical-validation evidence.

Literature

The point-dipole dielectrophoretic force and the Clausius–Mossotti factor are taken from Jones [20]; the drift-velocity forms of electrophoresis and thermophoresis are first-order models with the mobility and the thermodiffusion coefficient supplied as inputs. The closed-form decay rate against which the tracer is checked follows from balancing that force against Stokes drag and is derived in this study.

Analytical solution

Closed forms: nDEP restoring rate $4 \pi \epsilon a^3 \text{abs}(\text{Re } K) G^2 / (6 \pi \mu a)$; electrophoretic velocity $\mu_{ep} E$; thermophoretic drift $-D_T \text{grad } T$.

Result

Category closed_form

Tolerance / bar DEP decay rate within 5 percent; electrophoretic terminal velocity within 3 percent; thermophoretic force exact for a given gradient

Measured result 8/8 checks: DEP rate and drift within tolerance; thermophoretic force exact

Verdict **PASS**

Traceability MF.13 ElectroThermal | particle forces on SOLVED fields: DEP centring in a solved quadrupole, electrophoretic drift, thermophoretic drift (ParticleForces)

Honest caveat. BUILT ON REQUEST, NOT NEEDED BY THIS DEVICE (MF.0 screened electrokinetic and thermal particle forces out). First-order drift-velocity models with the mobility as an INPUT (the particle’s charge, zeta potential and double layer are folded into μ_{ep} ; no relaxation or retardation correction; D_T is a given number); DEP is the point-dipole limit with no electrothermal flow. The potential is a conduction solution (see :mf13_electric_potential_solver). Closed-form verification only: no experiment, and the drift velocities are not compared with any published measurement. It carries NO Tier-3 physical-validation evidence.

11.9 MF.13: particle-to-fluid momentum coupling fields (static two-way coupling) – the smoothed Stokes-drag reaction reproduces the exact Brinkman channel for frozen and for moving particles

Purpose

The particle tracer treats the fluid as unaffected by the particles, and the earlier coupling assessment found that the two validity criteria for that assumption disagree by five to six orders of magnitude. This case checks the missing direction, built as fields: a cloud of point particles is turned into a smoothed friction coefficient and a momentum source for the flow solver, so that the reaction on the fluid is the Stokes drag of the particles. Because that reaction is exactly a Darcy–Brinkman friction, the flow solution with frozen particles, and with particles moving at a prescribed speed, can be compared with the exact Brinkman channel profile. It is a static coupling: the loop that moves the particles and re-solves the flow is not built.

Literature

The frictional force of a dilute suspension on a fluid in the Stokes limit is the standard point-particle drag; the exact profile used for the comparison is the Brinkman channel solution already established in this study for the topology-optimisation readiness work, and no external result is reproduced.

Analytical solution

Brinkman channel: $\mu \mathbf{u}'' - \alpha \mathbf{u} + \mathbf{G} = 0$ with no slip (`brinkman_channel_velocity`), $\alpha = 6 \pi \mu a n$.

Result

Category `closed_form`

Tolerance / **bar** particle friction integral to 1e-6; uniform lattice friction to 1e-3; Brinkman channel profile within 2 percent

Measured result 11 of 11 checks

Verdict **PASS**

Traceability MF.13 ParticleForces | particle->fluid coupling fields: kernel normalisation, uniform lattice, and the exact Brinkman channel (static two-way coupling) (ParticleForces)

Honest caveat. STATIC coupling only: positions and velocities are given, and the moving-cloud fixed-point loop (advance the particles, re-solve, repeat) is NOT built, so no coupled focusing run exists. Stokes drag only (no Schiller-Naumann correction, lift, added mass or history), dilute limit (particle volume displacement ignored), a smoothing bandwidth that must exceed the particle spacing and be resolved by the mesh, and the 2D depth factor is a modelling choice. Verified against the exact Brinkman solution only.

11.10 Segre-Silberberg gate RE-RUN under the new `:ho_leal_3d` lift model: bidirectional mechanism confirmed, equilibrium prediction improved from 0.50R to 0.61R (literature 0.6R)

Purpose

The pipe gate is repeated unchanged except for the lift model: the local-half-width `:ho_leal_3d` correlation replaces the hydraulic-diameter one. The bidirectional mechanism is confirmed again and the model's self-consistent target moves from $0.50R$ to $0.61R$, within 1.7% of the literature's $0.6R$. It is cheap-geometry evidence for adopting the alternative model, not a production change.

Literature

The reference equilibrium is the Segré–Silberberg $0.6R$ [37]; the lift profile's stable zero is set from Ho and Leal [18] and the same reference.

Analytical solution

No closed form; compared against the classical Segre & Silberberg (1962) equilibrium-position result, same as `:mf10_a6_segre_silberberg`.

Result

Category `certified_literature`

Tolerance / bar `certified reference: r_eq/R=0.6` (Segre & Silberberg, single canonical value, no stated band); this model's own self-consistent target: `r_eq/R=1-0.39=0.61`

Measured result `particle 1: r0/R=0.1 -> r_final/R=0.116` (outward, correct direction); `particle 2: r0/R=0.3 -> r_final/R=0.3326` (outward); `particle 3: r0/R=0.7 -> r_final/R=0.6673` (inward, correct direction); `particle 4: r0/R=0.75 -> r_final/R=0.6994` (inward) – genuine bidirectional convergence toward an intermediate radius reproduced again under the new model. The model's own self-consistent target (0.61R) is within 1.7% of the certified 0.6R value, versus `:di_carlo_3d`'s own target (0.50R, 17% off). Mechanism CONFIRMED (as before); as with `:mf10_a6_segre_silberberg`, none of the 4 particles fully equilibrated within the one-transit residence time available at this mesh length – the SAME residence-time limitation already disclosed there, not a new caveat.

Verdict **FINDING**

Traceability `ForceModel3D lift dispatch: :ho_leal_3d` gives an identical result via `lift_force_3d` and `inertial_lift_3d` directly, `:di_carlo_3d` untouched (`ParticleForces`)

Honest caveat. This is NOT yet the production A.7 sweep's lift model – the sweep is running on `:di_carlo_3d` and stays there (see `:mf10_particle_forces_tracer_3d/:mf10_a6_segre_silberberg` above); changing the running sweep's own physics mid-flight would produce an aggregate mixing two different models across its 30 cells. This entry is separate validation evidence for a NOT-YET-PRODUCTION candidate, run on the cheap pipe benchmark only – it has not been run against the 6.5mm eel/serpentine production geometries, and `fL=0.5` remains the same single representative constant `:di_carlo_3d` already used (Di Carlo 2009/Asmolov 1999's exact fitted coefficients remain paywalled, unchanged from before).

11.11 MF.V: Ho-Leal (1974)/Segre-Silberberg (1962) equilibria collapse to one normalization when expressed as a fraction of the LOCAL half-width from the wall

Purpose

The production lift correlation normalises wall distance by the hydraulic diameter, which grows with channel width for the sweep's fixed 50 μm depth, so its predicted equilibrium drifts from 67% to 91% of the half-depth across the sweep even though the physical depth is constant. This case shows that the Ho-Leal and Segré-Silberberg equilibria collapse onto a single value once expressed in a local wall-normal half-width, identifying the normalisation, not the physics, as the source of the drift.

Literature

Ho and Leal [18] and Segré and Silberberg [37] agree to about 5% when expressed in the same normalisation, which is what licenses the shared stable-zero fraction.

Analytical solution

Two independent certified references (HO_LEAL_1974, SEGRE_SILBERBERG_1962), not a single closed-form solution – the check is that they AGREE under one normalization, not that either alone is a derivation.

Result

Category `cross_check`

Tolerance / bar the two references’ fractions must agree to within 6% (measured: 5.1%, inside the literature’s own documented replication spread on the 0.6R figure)

Measured result `s_ho_leal=0.380`, `s_segre=0.400`, relative spread=5.1%, averaged to `s_star=0.39`
– the stable-zero location `inertial_lift_profile/inertial_lift_3d` use by default

Verdict **FINDING**

Traceability MF.V (2026-09-18) `inertial_lift_profile` | Ho-Leal/Segre-Silberberg collapse gate (ParticleForces)

Honest caveat. This licenses the NORMALIZATION only (where the stable zero sits, as a fraction of the LOCAL half-width) – it is not, by itself, evidence that `inertial_lift_3d`’s overall shape or force MAGNITUDE match either reference’s own fitted profile; `:mfv_a6_ho_leal_gate` below is the actual numerical re-run this normalization licenses.

11.12 Brownian MSD vs Stokes-Einstein diffusivity, and Basset order-of-magnitude budget

Purpose

Brownian motion is implemented as an exact Ornstein–Uhlenbeck velocity kick, and its correctness is statistical: the mean-square displacement of a 60-member ensemble must equal $2Dt$ with the Stokes–Einstein diffusivity, with a tolerance derived from the ensemble’s own sampling statistics rather than chosen by hand. The Basset term is checked only for order-of-magnitude consistency with the size-independent criterion of the screening report, which is all that its diagnostic-only role requires.

Literature

The diffusivity relation is Einstein’s [12]; the Basset history term belongs to the Maxey–Riley equation of motion [23].

Analytical solution

Partial: the expected mean-squared displacement follows from the closed-form Stokes-Einstein diffusivity; the Basset-force budget is an order-of-magnitude comparison with no closed form of its own.

Result

Category cross_check

Tolerance / bar Brownian MSD: $\text{rtol} = 4 \cdot \sqrt{2/N_{\text{ens}}}$ (ensemble-statistics-derived); Basset: within one order of magnitude of $\tau_p/t_{\text{nu}} \sim 0.059$

Measured result Brownian MSD matches $2 \cdot D \cdot t$ within its ensemble-derived tolerance ($N_{\text{ens}}=60$); Basset ratio order-of-magnitude consistent with 0.059

Verdict **PASS**

Traceability MF.3 | Brownian MSD and Basset budget (statistical / diagnostic) (ParticleForces)

11.13 Composite-model (Saffman-Mei + Cherukat-McLaughlin) focusing position: honest measurement, not a tuned target

Purpose

This case records a measurement, not a pass. A $15 \mu\text{m}$ particle released inside the basin where the literature places a near-wall equilibrium is traced through a production-scale channel under the full composite lift model. The point is to report honestly where the model actually takes the particle, so that any claim of literature-quantitative focusing rests on a measured position and not on an expectation.

Literature

The composite model is the sum of the Saffman–Mei shear lift [34, 24] and the Cherukat–McLaughlin wall lift [7]; the inertial-focusing equilibrium picture it is compared against is that of Ho and Leal [18].

Analytical solution

No closed form; a measured, non-monotonic result compared qualitatively against the Ho & Leal / Segre-Silberberg literature mechanism, not a quantitative closed-form benchmark.

Result

Category cross_check

Tolerance / bar measurement, not a pre-registered pass bar (the plan’s own risk section anticipated this outcome as a possible finding, not a required pass)

Measured result s: 0.15 -j 0.1506 (0.4% relative change) – negligible net lateral migration

Verdict **FINDING**

Traceability MF.3 | focusing-position measurement (tests 10–11, production-scale channel) (ParticleForces)

Honest caveat. At this study’s real operating point, a near-zero-slip tracer (v_0 =local fluid velocity) released in a straight channel ($Du/Dt=0$ everywhere, confirmed by test 7) transits the entire 2300um production-scale channel and shows negligible net lateral migration under the composite lift model (s: 0.15 to 0.1506, a 0.4% relative change). Not a bug: Saffman-Mei is a slip-shear force, identically zero at zero slip, and the particle re-equilibrates to near-zero slip almost instantly ($\tau_p \sim 1.5e-5s$, four orders of magnitude shorter than the transit time); Cherukat & McLaughlin (1994)’s wall-lift term is gated to its own documented inner-region validity boundary ($1/a \lesssim \min(Re_\gamma^{-1/2}, Re_{slip}^{-1})$), which at this operating point works out to only $\sim 1-2$ particle radii from the wall – the particle released at 30um (4 radii) from the wall never enters that zone. The classical Ho & Leal/Segre-Silberberg inertial-focusing mechanism does not arise from plain Saffman lift on a near-zero-slip point particle – it requires finite-size stresslet/curved-shear interaction effects (Asmolov, McLaughlin, Rubinow-Keller-type corrections) that neither correlation captures alone. Not swept under the rug: the second unresolved-but-documented finding in this study, alongside MF.1’s high-Re lid-cavity mismatch.

11.14 Force-kernel self-checks: added-mass prefactor, drag limits, Saffman-Mei low-Re coefficient, Cherukat-McLaughlin asymptote, Brownian variance

Purpose

Each particle force is a small formula that is easy to get wrong by a constant factor. The kernel-level checks compare each against a hand-derivable identity or a verified limit: the added-mass term equals half the displaced fluid mass, the drag reduces to Stokes drag as the particle Reynolds number vanishes, drag is friction times slip velocity by definition, and the shear-lift coefficient tends to its classical low-Reynolds value. Together they would expose a wrong prefactor or sign before any trajectory-level test blames the integrator.

Literature

The drag correction is the Schiller–Naumann correlation [36]; the added-mass and pressure-gradient terms follow the Maxey–Riley equation [23]; the shear lift follows Saffman [34] with Mei’s finite-Reynolds correction [24]; the wall lift follows Cherukat and McLaughlin [7].

Analytical solution

Yes: Stokes-drag, added-mass, and Saffman-lift closed-form algebraic identities.

Result

Category closed_form

Tolerance / bar added-mass term $rtol=1e-12$; Stokes limit $rtol=1e-6$; drag_force definition $rtol=1e-12$; Saffman-Mei low-Re coefficient $rtol=1e-4$; Cherukat-McLaughlin contact asymptote $rtol=1e-3$, far-field exactly (0,0); Brownian variance $rtol=1e-6/1e-9$; $6\pi\mu a$ identity $rtol=1e-14$

Measured result all 8 sub-checks pass to their stated tolerances

Verdict **PASS**

Traceability MF.2 ParticleForces | force-kernel self-checks (ParticleForces)

11.15 MF.0 physics-relevance-screening force budget: Stokes relaxation, added-mass prefactor, Basset criterion, gravity/Brownian displacement budgets

Purpose

Which forces to switch on by default was decided by an order-of-magnitude budget at the study's real operating point. Left as report prose, that budget could silently drift from the code. This case reproduces it as a runnable arithmetic check (Stokes relaxation, the added-mass prefactor, the Basset criterion, gravity and Brownian displacement), so a change to the operating point or to a kernel cannot invalidate the default term set without a test failing.

Literature

The force set follows the Maxey–Riley equation [23] and the drag correlation of Schiller and Naumann [36]. The operating point is the study's own (reports/mf0_physics_relevance_screening.md).

Analytical solution

Partial: the Stokes relaxation time is closed form; the measured force-ratio budget itself is a numerical comparison against MF.0's own order-of-magnitude estimate, not a solution match.

Result

Category closed_form

Tolerance / bar Stokes relaxation vs methodology doc: 3 significant figures; particle Re_p spans documented as 1-17 (finite- Re correction required, confirmed not assumed); added-mass prefactor and Basset criterion: order-of-magnitude, not tight targets; gravity/Brownian displacement $\ddot{y}$ channel width (100-500 μ m)

Measured result $\tau_p(2\mu\text{m})=2.647\text{e-}7\text{s}$ (methodology doc states $\sim 2.6\text{e-}7\text{s}$); Re_p spans 1.12-16.85 for $d_p=2\text{-}30\mu\text{m}$; added-mass prefactor $0.5*\rho_f/\rho_p=0.47$; Basset criterion $\rho_p/(18*\rho_f)=0.059$ (exactly size-independent); gravity displacement 0.15 μm , Brownian rms displacement 0.048 μm (2 μm class) down to 0.012 μm (30 μm class), both $\ddot{y}$ H

Verdict **PASS**

Traceability PhysicsScreening | arithmetic cross-check for reports/mf0_physics_relevance_screening.md (ParticleForces)

11.16 Saffman lift in isolation matches the closed-form classical coefficient (1.615) across the size range

Purpose

The Saffman–Mei lift is evaluated directly, not through a trajectory, and compared with the closed-form classical coefficient across the 2–30 μm size range. Isolating the kernel from the integrator means a disagreement can only come from the formula or its coefficients, so this would catch a mis-transcribed constant or an exponent applied to the wrong quantity.

Literature

Saffman [34] gives the low-Reynolds shear-lift coefficient; Mei [24] gives the finite-Reynolds approximation used here.

Analytical solution

Yes: Saffman (1965)’s closed-form lift coefficient (1.615), independently hand-derived and confirmed this study.

Result

Category closed_form

Tolerance / **bar** rtol=1e-3

Measured result matches to $\pm 0.1\%$ across d_p in $\{2,15,30\}\mu\text{m}$

Verdict **PASS**

Traceability MF.3 | particle validation suite (ParticleForces)

12 ParticleTracer

12.1 Segre-Silberberg tubular-pinch gate: genuine bidirectional migration confirmed; literature 0.6R value missed (disclosed)

Purpose

The tubular-pinch gate asks whether the tracer reproduces genuine inertial migration: neutrally buoyant particles released both inside and outside the model’s equilibrium must migrate toward it. The mechanism passed (bidirectional migration). The magnitude did not: the model’s own equilibrium is near $0.5R$ where the literature reports about $0.6R$, the disclosed consequence of the width-dependent normalisation of the lift correlation.

Literature

The equilibrium annulus at roughly $0.6R$ is the Segré-Silberberg result [37]; the theoretical counterpart in a channel is that of Ho and Leal [18]. The pipe number is used here only as a reference for the depth direction.

Analytical solution

No closed form; compared against the classical Segre & Silberberg (1962) equilibrium-position result – a single canonical experimental/numerical value ($0.6R$), not a derived closed-form solution.

Result

Category certified_literature

Tolerance / **bar** mechanism: bidirectional migration observed (qualitative pass); magnitude: literature canonical equilibrium is $0.6R$ (independently web-confirmed this session)

Measured result mechanism CONFIRMED (all 4 particles moved the correct direction); magnitude MISSED – model’s own equilibrium is $\sim 0.5R$, not $0.6R$, none of the 4 particles fully equilibrated within the one-transit residence time available at this mesh length

Verdict FINDING

Traceability MF.10 ParticleTracer | 3D smoke test: particles move through the solved pipe field (ParticleTracer)

Honest caveat. The $0.5R$ -vs- $0.6R$ gap is NOT a new finding – it is the predicted, disclosed consequence of `di_carlo_lift_3d` being a documented substitution (see `:mf10_particle_forces_tracer_3d` above), now confirmed by an actual run rather than left as a hypothesis. A literature-fitted `f_L(Re,position)` profile (Liu et al. 2016 or equivalent, if its coefficients become independently verifiable) would be the direct next step to close this gap – not a re-tuning of this stand-in’s `f_L` constant, which would fit the number without fixing the underlying substitution.

12.2 MF.12 Step 4: Di Carlo square-duct gate on a NOISE-FREE field – single-wall rings confirmed, multiwall selects corners, not face-centres

Purpose

Tests the lift models against the Di Carlo et al. (2007) square-duct result while removing a confound. The flow field is the exact duct solution interpolated into the solver’s own velocity space, so the transverse velocity is identically zero. Through the finite-element field the same runs showed particles drifting to the wall; that drift was discretization noise (spurious transverse velocity comparable to the lift migration speed near equilibrium), and this entry exists so the model is judged without it. It would catch a lift model that fails to converge onto its own predicted equilibrium ring, and shows that none of the tested models selects the literature’s four face-centre positions.

Literature

Di Carlo et al. (2007), *Continuous inertial focusing, ordering, and separation of particles in microchannels*, PNAS 104, 18892: four focusing positions at the centres of the faces of a square channel. The literature target was confirmed from published summaries of that paper, not by reading it in full; earlier text in this study called it “corners”, which was a mislabel.

Analytical solution

Exact rectangular-duct Fourier-series Stokes solution (transverse velocity identically 0). Predicted rings from the lift models’ own zero-crossings: $wall_d/H = 0.25$ (`di_carlo`, `s=wall_d/Dh`) and 0.195 (`ho_leal`, `s*=0.39` of the half-width). Literature target: four face-centre positions (Di Carlo et al. 2007).

Result

Category `cross_check`

Tolerance / bar each single-wall model converges to its own predicted ring within $0.005H$ for a large majority; the literature target is selection of the four FACE-CENTRE positions (angle from a face-centre axis $\neq 0$)

Measured result clean field, K=120: di_carlo 0.2492 (pred 0.25), ho_leal median 0.1954 (pred 0.195); face-centre selection NOT achieved by any lift; multiwall median angle 38.3 -j 44.9 deg (diagonal)

Verdict **FINDING**

Traceability MF.12 ParticleTracer | :ho_leal_3d_multiwall wiring through trace_particles_3d!
(Step 4) (ParticleTracer)

Honest caveat. One particle size (10um), one Re (22.5), one quadrant ensemble (N=44, 2 diagonal probes released inside the wall zone freeze at once), f_L=0.5 and s*=0.39 unchanged; 120 legs (36 mm) is converged for the rings but the slowest near-centre :ho_leal_3d particles were still arriving. The clean field is an exact-interpolated idealization: it tests the lift MODEL, not the FEM solver, and the earlier FEM-field runs are retained only as the record of what misled (di-carlo_v2_lift_comparison.json, dicarlo_v2_k120_sw.json). The mean-wall-distance falling BELOW the lift zero-crossing in those FEM runs is attributed to the transverse-velocity noise by the clean-field rerun, but the individual candidate mechanisms (nearest-wall tie-breaking, the per-leg local-fluid restart) were not separately isolated. The literature target was confirmed by web search (PNAS 2007 and a review), not by reading the paper; older text in this study says 'corners'. Sweep-level audits (out-of-channel containment, relay stranding) are documented in notes.md and the validation program, not registered here, because they are not module benchmarks.

12.3 MF.13 Phase 2: three particle-tracer defects found by reading the source, each with a failing-first regression test

Purpose

Reading the tracer and force source while planning MF.13 turned up four suspected defects. Three are addressed with a regression test that was run against the unfixed source and fails there: the diagnostic force-budget call omitted the wall list, so any tracing with diagnostics under the multiwall lift threw an error; the single-wall Ho-Leal lift silently substituted the hydraulic-diameter normalisation for a missing local half-width; and the minimum time step could override the never-step-halfway-to-a-wall guard. The fourth suspected defect (Brownian noise built from plain Stokes friction) was examined and withdrawn, and the reason is recorded in the code.

Literature

The Brownian argument rests on Einstein's relation between diffusivity and the zero-slip friction [12]; the drag correction whose slip-based Reynolds number is small in this study is Schiller and Naumann's [36].

Analytical solution

No closed form; behavioural regression tests.

Result

Category regression

Tolerance / bar each regression test fails on the pre-fix source and passes after; existing suites unchanged

Measured result B-1 3/3, B-2 2/2, B-4 2/2 after the fix; existing ParticleForces (17, 33, 8) and tracer suites unchanged

Verdict **PASS**

Traceability MF.13 ParticleTracer | B-1: `:ho_leal_3d_multiwall` with non-empty `diag_particles` no longer throws (ParticleTracer)

Honest caveat. B-4 is LATENT, not live: the default `dt_min` of $1e-14$ s can only beat the wall limiter at speeds above about $5e7$ m/s, so no prior result is affected; it was fixed for consistency and tested with an artificially large `dt_min`. B-2 was reachable only by direct kernel callers (the tracer always supplies `h_local`). A fourth candidate, B-3 (Brownian kick built from plain Stokes friction while drag uses the Schiller-Naumann correction), was EXAMINED AND WITHDRAWN as a defect, not fixed: the Einstein relation ties the noise to the zero-slip friction, and the tracer's slip-based `Re_p` is about $1e-3$ so the correction is about 1.001; the reasoning is now in the `brownian_kick_3d` docstring.

12.4 MF.14: DLVO contact versus the geometric freeze on a 44-particle clean-field square-duct ensemble – capture becomes sliding contact only when the barrier's peak FORCE exceeds the wall-directed lift; the particles stay on the wall

Purpose

This case puts the contact model to the finding it was built for: are the walls the particle sweep stalls at real, irreversible captures? On one small clean-field ensemble in a square duct of the sweep's depth, it counts captured and held particles under the geometric freeze and under the DLVO contact at three Debye lengths, for two particle sizes and both lift models. It is a comparison, not a validation, and it would show a wrong contact model by producing a difference that the force balance cannot explain. What it shows is that nothing is captured at 15 micrometres under any rule, and that at 30 micrometres the outcome is decided by whether the barrier's peak force exceeds the wall-directed lift, not by the barrier's height in units of thermal energy, with the particles remaining pressed on the wall when the barrier holds.

Literature

The interaction is that of Derjaguin and Landau [10], with surface potentials, Hamaker constant and ionic strength assumed because none is measured for these channels. There is no published wall-capture experiment for this geometry to compare with, so no external reference is claimed; the finding is a statement about this model and its assumed inputs.

Analytical solution

Not yet characterized – pending review.

Result

Category diagnostic

Tolerance / bar counts of captured and held particles per scenario; compared, not against a reference

Measured result D=15 um: 0/44 captured in every scenario. D=30 um: 44/44 captured under :freeze and at 100 nm; 0/44 captured, 44/44 held on the wall at 10 nm and 1 nm (both lifts)

Verdict FINDING

Traceability MF.14 DLVO wall contact | tracer: held vs captured vs :freeze on a wall-directed push in quiescent fluid (ParticleTracer)

Honest caveat. A COMPARISON on one small clean-field ensemble, not a replacement of any recorded number: the sweep's results stay as recorded under :freeze. The surface potentials, Hamaker constant and ionic strength are ASSUMED (nothing here measures them), so the D = 30 um outcome is a function of chemistry that is not known for these channels. A 'held' particle is not a freed one: the wall-directed push that drove it to the wall (a point-particle lift model at blockage 0.6, outside its validity, and not interpretable per FLUDICTS.txt) is unchanged, so DLVO turns irreversible capture into sliding contact without moving the particle off the wall. It does NOT explain or repair the ~44% stall fraction of the sweep, which was not re-run under :dlvo. The wall push is lift-only here (no Brownian escape). No experiment or COMSOL comparison.

12.5 MF.14: DLVO wall contact in the 3D tracer (wall_model=:dlvo) – a sub-barrier push is HELD at the force-balance gap, a push above the peak repulsion (or no barrier) is CAPTURED exactly as :freeze

Purpose

A finding of the particle sweep was that every stalled chain is a wall capture, decided by a purely geometric freeze rule. The colloidal force kernels of the previous unit could not simply be added as a force in the tracer: its explicit steps are set by the cell size, of order ten micrometres, while the interaction varies over the Debye length, of order nanometres, so a step would tunnel through the barrier. This case checks the contact model that replaces the freeze at the tracer's existing contact event: the wall-directed push on a particle at rest normal to the wall is balanced against the repulsion curve, the particle is held at the stable balance gap if the push is below the peak repulsion, and captured exactly as before otherwise. It would catch a held gap that is not the root of the force balance, a held particle that creeps into the wall or saw-tooths away from its gap (the first version did, measured 7.7 nm against 18.7 nm), a capture that differs from the geometric rule, and any change to the default behaviour.

Literature

The interaction curve is the sum of a van der Waals attraction and a screened double-layer repulsion in the theory of Derjaguin and Landau [10], in the Derjaguin sphere–plate and weak-potential form used by the kernels this case builds on. The contact model itself, a quasi-static force balance at the contact event, is this study's own construction and has no literature counterpart to be compared with; it is therefore checked against the closed-form force curve and against a quiescent-fluid tracer run with a known constant push, not against a published trajectory.

Analytical solution

Held gap g solves $F_{DLVO}(g) = P$ on the branch between the repulsion peak and the secondary zero, with $F_{DLVO}(g) = -A_H a / (6 g^2) + 4 \pi \epsilon a \kappa \psi_p \psi_w \exp(-\kappa g)$; captured when $P \geq \max F_{DLVO}$.

Result

Category closed_form

Tolerance / bar held gap within one step's terminal travel of the closed-form balance gap; capture position identical to :freeze; :freeze bitwise unchanged

Measured result 43 of 43 checks (27 kernel, 16 tracer)

Verdict PASS

Traceability MF.14 DLVO wall contact | tracer: held vs captured vs :freeze on a wall-directed push in quiescent fluid (ParticleTracer)

Honest caveat. 3D tracer only (the 2D tracer refuses :dlvo). Quasi-static: no inertia, no Brownian escape over the barrier, no lubrication correction to the near-wall drag, no contact mechanics or adhesion hysteresis; a held particle is restored to its balance gap at contact events, not carried through a continuous DLVO force. The Derjaguin/weak-potential limits of the kernels apply (gap much smaller than the radius, potentials below about 25 mV). The Hamaker constant, surface potentials and Debye length are the caller's ASSUMPTIONS: this study measures none of them. Verified against the closed-form DLVO force curve only; no measured adhesion or force curve, no experiment, no COMSOL comparison.

12.6 MF.15: in the leg-relay sweep, a leg's flow field is 7-27% off a long reference within ~125 um of its own zero-traction outlet on the sharp serpentine channels, and every particle path runs through that zone; a 250 um overlap makes the hand-off plane clean (0.3% leg-to-leg) and a trace-stop plane keeps particles out of the zone

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category diagnostic

Tolerance / bar hand-off plane error at most 8% (the sweep's MAX_ACCEPTABLE_OVERLAP_REL_ERR, unchanged)

Measured result hand-off plane 0.62% / 0.42% / 0.50% (sharp 100 / 200 / 300 um, pre-flight at the tightest bend, Q = 0.7); leg 1 versus a long reference 24-28% off 10 um before its outlet at 300 um legs on the 200 and 300 um channels

Verdict FINDING

Traceability MF.15 ParticleTracer -- stop_fn ends a 3D trace at a caller-defined surface (leg-relay trusted-zone stop) (ParticleTracer)

Honest caveat. A diagnostic finding, not a validation against literature or experiment. The reference is itself a numerical solution (a longer leg whose inlet is a uniform plug 200-450 um upstream of the probe points and whose own outlet is 300 um farther on), so ‘error against the reference’ is a consistency measure, not a bound on the true error; leg 2 sitting within about 1% of it across the overlap in the 450 um design is the evidence it is the better field. Probes are the relative velocity-vector difference on the leg-1 centreline at z/depth 0.3, 0.5, 0.7 only, at one crest per channel and (for the reference study) $Q = 0.7$; the lateral and near-wall regions were sampled at a few points in the overlap study, not mapped. Not measured: the effect of the outlet zone on particle equilibria. The v3 sweep and every baseline (non-sharp) channel of v4 still trace to the leg outlet (a stop would change their physics relative to v3), so their particle results carry this zone; for the v3 eel geometry the field error there is small at the narrow widths measured (7.2% worst, at 10 um) but v3’s own gate only ever probed eel_500um at $Q = 0.5$. The 450 um legs cost about 4-5 times as much per leg as 300 um legs. sharp_100um keeps the v3 design (300 um legs, no stop) because its measured errors are already 1-3.4%. Mechanism is the zero-traction outlet artifact already registered as :mf13_t4_outlet_localisation, here amplified by curvature-driven secondary flow; no cause beyond that was isolated.

12.7 MF.16: interior-leg-outlet capture clustering is robust and per-lift-significant on eel_500um (di_carlo 12/12, ho_leal 14/15, both $p \leq 0.0001$) but weak/unconfirmed on eel_100um, whose own seed+jitter control reproduces NO clustering at all (6/24, exactly the 25% null) – a properly tested, per-lift, tangent-plane-distance version of the MF.15 finding, correcting BOTH a prior session’s mis-summary AND this session’s own first-pass pooled/raw-x-distance analysis

Purpose

Literature

Analytical solution

Not yet characterized – pending review.

Result

Category diagnostic

Tolerance / bar binomial $p \leq 0.05$ PER LIFT MODEL (not pooled) for the interior-leg clustering claim; the control must not show the same clustering if the baseline signal is to be treated as seed-independent

Measured result eel_500um baseline: di_carlo 12/12 $p \leq 0.0001$, ho_leal 14/15 $p \leq 0.0001$ (robust); eel_100um baseline: composite 9/18 $p = 0.025$ only (weak); eel_100um control: 6/24 $p = 1.0$ (no clustering) – eel_100um finding NOT confirmed seed-robust

Verdict **FINDING**

Traceability MF.15 ParticleTracer -- stop_fn ends a 3D trace at a caller-defined surface (leg-relay trusted-zone stop) (ParticleTracer)

Honest caveat. Restricted to the two cases `mf15_outlet_zone.jl` ran (not a sweep-wide result); the `eel_100um D=30 di_carlo leg-32` cluster (17- $\dot{i}$ 2 frozen) is a DIFFERENT mechanism (the real channel outlet) and is excluded, reported separately. Relay-reprojection totals differ substantially between designs on `eel_500um` (348 baseline vs 976 long-leg) – a candidate confound, reported, not resolved. No control run exists for `eel_500um`, so its strong result cannot be checked against seed noise the same way `eel_100um`’s was. Per-lift sample sizes are small ($n=3-18$), so several cells are simply inconclusive rather than negative.

12.8 Stokes relaxation, added-mass consistency, neutrally-buoyant limit, terminal settling, IMEX-vs-RK4 cross-check, wall models, and COMSOL-baseline-vs-full-model regression

Purpose

These are the tracer’s own kinematic correctness checks: velocity relaxation follows the exponential closed form, the added-mass-corrected relaxation time is wired through the integrator (not merely algebraically), a neutrally buoyant particle follows the fluid, terminal settling matches Stokes, and the implicit–explicit integrator agrees with RK4. A wrong effective mass, a mis-split stiff term or a step-size limiter that interferes would each show up as a failure of a closed-form comparison.

Literature

The force decomposition is the Maxey–Riley equation [23] with the Schiller–Naumann drag correction [36]; the closed forms compared against are textbook.

Analytical solution

Yes: the exact exponential-drag Stokes-relaxation solution $v(t)=v_{\text{terminal}}*(1-e^{-(t/\tau)})$.

Result

Category `closed_form`

Tolerance / bar Stokes relaxation `rtol=5e-3` (over 6 relaxation times); added-mass `rtol=5e-3`; neutrally-buoyant `rtol=1e-3` (velocity), `atol=1e-9` (position); terminal settling `rtol=1e-2`; IMEX-vs-RK4 `rtol=1e-4`; wall models: exact non-penetration + correct active/inactive state per model; COMSOL-baseline spread $\dot{i}$ 1% of channel width, full-model spread strictly larger

Measured result Stokes relaxation $\dot{i}$ 0.5% over 6 relaxation times; added-mass effective- τ $\dot{i}$ 0.5%; neutrally-buoyant $\dot{i}$ 0.1% velocity match, zero measurable drift; terminal settling $\dot{i}$ 1%; IMEX vs RK4 agree to $\sim 5e-10$ relative on a Stokes-relaxation case; all 4 wall models pass containment + state checks; COMSOL baseline spread $\dot{i}$ 1% of H, full-model spread $\dot{i}$ baseline

Verdict **PASS**

Traceability MF.3 | particle validation suite (ParticleTracer)

12.9 trace_particles!'s GUI-facing progress_cb/cancellation hooks are a pure, non-perturbing query and produce self-consistent partial results when cancelled

Purpose

The graphical interface needed a progress bar and a working Stop button, which meant editing the already-validated particle driver loop. That is only acceptable if the new hook is provably non-perturbing. Three properties are checked: a run with no callback is identical to before, a callback that only observes leaves results unchanged, and cancellation stops cleanly with consistent partial output.

Literature

No external literature applies: this is an internal engineering or self-consistency check, and no published reference defines its pass criterion.

Analytical solution

Not applicable – a bit-for-bit regression/identity check on new callback plumbing; no physics solution is involved.

Result

Category regression

Tolerance / bar identity: exact equality (`==`), not `isapprox`, on every returned array; progress: `issorted(completed)` and `calls[end]==(N,N)`; cancel: `cancelled==true`, `sum(steps_taken)` strictly less than the uncanceled run, all post-cancellation `steps_taken==0`, `traj` arrays parallel-length-consistent

Measured result identity holds exactly; progress sequence monotonic and terminates at `(N,N)`; cancellation after 3 of 6 particles leaves the remaining 3 untouched (`steps_taken=0`) and roughly halves the total step count

Verdict **PASS**

Traceability MF.6 ParticleTracer | progress_cb / cancellation hooks (bridge re-validation)
(ParticleTracer)

13 PhaseField2D

13.1 MF.13: wetting boundary condition of the phase-field solver and the Lippmann-Young electrowetting law – a drop relaxes to the imposed equilibrium contact angle within 0.5 degrees at 120, 90 and 60 degrees

Purpose

Electrowetting on dielectric changes the apparent contact angle of a conducting drop with the applied voltage. This case checks the two pieces built here: a wetting boundary condition for the phase-field solver, and the Lippmann-Young law that maps a voltage to an angle. The test is

that a drop relaxing under the imposed wall angle reaches that angle, at three voltages spanning hydrophobic to hydrophilic. It is a check of an *equilibrium boundary condition*: no electrostatic field is solved, and saturation, hysteresis and contact-line dynamics are outside it. The record also carries a measurement lesson, kept in the caveat: an angle inferred from a diffuse drop’s area and its zero-contour height was biased by about 20° and had to be replaced by a contour-based estimate before the boundary condition could be judged.

Literature

The Lippmann–Young equation and its range of validity, including the saturation this implementation does not model, are reviewed by Mugele and Baret [25]. The wall free-energy form used for the wetting boundary condition, with the angle measured inside the phase carrying $\phi = +1$, follows the diffuse-interface treatment of Jacqmin [19].

Analytical solution

Young’s law (the imposed equilibrium angle) and the Lippmann-Young law $\cos(\theta) = \cos(\theta_0) + \frac{\epsilon_0 \epsilon_r V^2}{2 \sigma d}$; the equilibrium drop is a circular cap.

Result

Category closed_form

Tolerance / **bar** measured equilibrium angle within 2 degrees of the Lippmann-Young value

Measured result 119.98, 89.62, 59.5 degrees against 120, 90, 60 (within 0.5 degrees)

Verdict **PASS**

Traceability MF.13 ElectroThermal | electrowetting: Lippmann-Young contact angle, closed forms and refusal beyond full wetting (PhaseField2D)

Honest caveat. This validates the EQUILIBRIUM angle imposed by a wall boundary condition, not an electrowetting device. No electrostatic field is solved (the voltage enters only through the closed-form law), and contact-angle saturation, hysteresis, contact-line friction and drop actuation are absent. A first draft measured the angle from the drop’s area and height and reported 89, 71 and 50 degrees; that was a MEASUREMENT bias (a half-disc under the natural 90 degree condition read 73 degrees, because the area of $\int (1 + \phi)/2$ and the height of the $\phi = 0$ contour describe different contours of a diffuse drop, $\epsilon_0/R = 0.2$), not a boundary-condition error, and the first-draft numbers are kept in `mf13.ewod_first_draft_area.height_method.json`. The angle measurement assumes a circular cap. One geometry, one interface thickness, one mesh (0.6 μm , 1870 cells), one drop size; the bulk free energy printed in the JSON is not monotone in two of the three cases because it excludes the wall energy, and the 0.25 to 0.29 m/s peak speed there is the physical contact-line relaxation (about 1 percent of σ/η), not a spurious current. No comparison with an experiment or COMSOL.

13.2 MF.13: Cahn-Hilliard-Navier-Stokes phase-field two-phase flow on the unstructured 2D path – Laplace pressure jump within 1 percent, mass and energy behave; the ellipse relaxation is monotone but unfinished

Purpose

COMSOL offers three two-phase methods (level set, phase field, moving mesh); the only interface-capturing solver in this code base was a structured-mesh volume-of-fluid scheme, which cannot represent a body-fitted microchannel. This case records the first unstructured two-phase solver, a monolithic backward-Euler Cahn–Hilliard–Navier–Stokes discretisation, and the checks made on it: the Laplace pressure jump across a static drop against σ/R , the behaviour of the phase integral and of the free energy, the size of the spurious velocity, and the relaxation of an elliptical drop toward a circle. It is a *finding* and not a pass: the pressure jump is within one percent, but the ellipse was still relaxing when the run ended, the chemical potential was not yet uniform, the phase integral is conserved to about 3×10^{-6} of the box area and not to round-off, and no droplet-generation (T-junction) case has been run.

Literature

The free energy with a gradient term and a double-well potential is the model of Cahn and Hilliard [6]; its coupling to the Navier–Stokes equations through a capillary body force $-\phi \nabla \mu$, as used here, is the diffuse-interface formulation of Jacqmin [19]. The Laplace law used as the reference, $\Delta p = \sigma/R$ for a two-dimensional drop, is the standard sharp-interface result and is not cited to a source here. No published Cahn–Hilliard–Navier–Stokes benchmark was reproduced.

Analytical solution

Laplace law: pressure jump σ/R for a 2D drop (sharp-interface limit; diffuse-interface corrections are $O(\epsilon\sigma/R)$); planar tanh profile with interfacial energy σ .

Result

Category diagnostic

Tolerance / bar Laplace jump within a few percent of σ/R_{eff} ; integral of ϕ conserved to a small fraction of the box area; energy monotone; aspect ratio decreasing toward 1

Measured result Laplace ratio 0.990; integral of ϕ drift $2.6\text{e-}6$ of box area (not round-off); energy monotone; aspect 1.87 -> 1.11 (not yet 1)

Verdict **FINDING**

Traceability MF.13 PhaseField2D | drop relaxation: energy decays monotonically, phase integral drift stays small, refusals (PhaseField2D)

Honest caveat. A FINDING, not a pass: the ellipse was still relaxing when the run ended (aspect 1.11, not 1), so a full circularisation and the relaxation TIME were not measured, and the chemical potential was not yet uniform (relative spread 0.16 across inside and outside samples) even though the Laplace jump built from its mean is within 1 percent. The phase integral is conserved only to about $3\text{e-}6$ of the box area, NOT to round-off: Taylor-Hood is divergence-free only weakly, so the discrete convective term does not integrate to zero against a P2 phase field; the JSON field

mass_drift is an ABSOLUTE m^2 number ($2\text{e-}15$), not a relative one. The peak speed during the ellipse relaxation, 0.66 m/s, is the physical capillary relaxation (2 percent of σ/η), not a spurious current. 2D only, laminar, first-order in time, one mesh (about 1 element per ϵ s at the interface) and one mobility; no wetting model beyond the natural 90 degree contact angle, so NO wall-wetting droplet generation (a T-junction) was run; the sharp-interface limit is not asserted; the convective form is not conservative for large density ratios; not compared with a published CHNS benchmark or a COMSOL result. The static-drop run took 255 s for 12 steps (about 21 s per step, direct LU of the four-field Jacobian), so this is a small-domain tool.

13.3 MF.13: T-junction on the phase-field solver – the open-flow two-phase case RUNS, but no droplet was generated in any of four configurations (the dispersed phase forms a thread that reaches the outlet)

Purpose

The COMSOL comparison named droplet generation in a T-junction as the missing two-phase case. With inflow velocities, a phase-carrying inlet, a traction-free outlet and a wall contact angle now available on the phase-field solver, this case runs a two-dimensional T-junction on an unstructured mesh and records what happens. It is a *finding* and not a droplet-generation result: in none of four configurations did a droplet detach, the dispersed phase forming a thread that reaches the outlet. The purpose of recording it is to establish that the open-flow two-phase machinery runs end to end and to state, without inflating it, what was and was not obtained.

Literature

Droplet breakup in a T-junction at low capillary number is governed, in three dimensions, by the pressure that builds behind a plug filling the channel cross-section and squeezes the neck through the thin films leaking past it, as shown by Garstecki and co-workers [14], who also give the size-scaling law. A two-dimensional model has no analogue of the thin films, so the absence of pinch-off here should not be read as a statement about a three-dimensional device, and the scaling law was not reproduced.

Analytical solution

none: no closed form exists for droplet generation; the scaling law of Garstecki et al. is a 3D squeezing-regime correlation and was not reproduced.

Result

Category diagnostic

Tolerance / bar n/a (a case run; the droplet formation was not reached)

Measured result no droplet generated in four configurations; parallel (thread) flow in the last, converged run

Verdict **FINDING**

Traceability MF.13 PhaseField2D | open flow: an inflow carries phase 1 in at the flux the boundary data prescribe, and the outlet pins the pressure (PhaseField2D)

Honest caveat. A FINDING, NOT A DROPLET-GENERATION RESULT. Four points in a large parameter space (viscosity ratio, wall angle, capillary number, channel-width ratio, flow-rate ratio, interface thickness and mobility were mostly not varied), one mesh, backward Euler with a time step whose CFL number is above one. The 2D geometry has no top and bottom walls, so the mechanism that sets breakup in real T-junctions at low capillary number, a pressure build-up behind a plug that fills the cross-section and squeezes the neck through the thin films leaking past it (Garstecki and co-workers), has no analogue here; a 2D failure to pinch off therefore says nothing about a 3D device. Three of the four runs had steps where Newton hit its 15-iteration cap without meeting the residual reduction: run (1) at step 71, run (2) at the LAST SIX steps (95 to 100, so its final thread history is not converged), run (3) at steps 1 to 3 (the impulsive start) and 101 to 102 (as the thread front reached the outlet); run (4) had none, and its repeat after the corner-node fix hit the cap only at step 1 and reproduced the same steady parallel flow (dispersed area 231.75 square micrometres, thread front at the outlet from 37 us). Runs (1) to (3) were stopped at 100 to 120 steps and (2) had not reached a steady state. The droplet detector labels a component detached when it does not touch the top row of the side channel and was never triggered; it was never validated on a case with a known detached droplet, so a true detachment could in principle have been missed; a coarse-grid snapshot could also miss a thin neck (grid 0.5 um). No comparison with the droplet-length scaling law, an experiment or COMSOL was made. The live testset anchor covers the open-flow BOUNDARY-CONDITION wiring only (an exact flux balance, the wall and the outlet); the T-junction case itself is measured by the driver and recorded in the JSON. Runs (1) to (3) used the solver BEFORE the corner-node weak enforcement of the velocity conditions was added (a wall-flux leak of 0.3 percent of the inflow at the outlet corners); run (4) was repeated after it, see the JSON with the suffix cornerfix.

14 SpeciesTransport2D

14.1 MF.13: first-order surface reaction (Robin wall condition) on species transport – exact linear profile and wall flux for a stagnant slab, from insulating to perfect-sink limits

Purpose

COMSOL's transport interfaces carry surface reactions at walls, while this code base had only volume reactions. The purpose of the case is to check the new wall condition, a first-order irreversible reaction expressed as a Robin flux condition, against the one situation where the steady solution is exact and linear: a stagnant slab held at fixed concentration on one wall and reacting at the other. The check spans the insulating wall, intermediate rates and the perfect-sink limit, so any error belongs to the boundary term. It says nothing about saturable adsorption kinetics, which are not modelled.

Literature

The Robin form of a first-order surface reaction and its two limits are elementary transport theory, and the slab solution is derived directly from it in this study; no external source is being reproduced.

Analytical solution

Steady stagnant slab: $C(y) = 1 - k y / (D + k H)$; wall flux $k D / (D + k H)$.

Result

Category closed_form

Tolerance / bar profile relative error below 1e-9; wall flux within 1e-8 relative of $k D / (D + k H)$

Measured result profile below 1e-9, flux within 1e-8, at four rate ratios; refusals

Verdict **PASS**

Traceability MF.13 SpeciesTransport2D | first-order wall reaction (Robin): exact linear profile and flux, sink and insulating limits (SpeciesTransport2D)

Honest caveat. First-order, irreversible, linear only. Saturable (Langmuir) adsorption, desorption, site depletion and wall-bound species are NOT modelled, so this is not adsorption kinetics. Verified in a stagnant slab (no advection) against a closed form; the Graetz-Leveque limit with flow was not tested. No experiment, no COMSOL comparison.

14.2 SpeciesTransport2D closed-form patch tests + cross-check against CFD-Core.species_transport_2d + Taylor-Aris cross-check

Purpose

Four complementary checks on the species solver: pure diffusion recovers the exact linear profile; a uniform inlet concentration is a trivial fixed point for any velocity field; the outlet profile agrees with the older species solver on the same mesh (stated plainly as a consistency check, since the two weak forms are near-identical by construction); and the dispersion coefficient extracted from the solver's own sampled Poiseuille profile matches the closed-form Taylor-Aris coefficient $K = \bar{U}^2 H^2 / (210 D_m)$. That last check is the benchmark validation of the existing dispersion module that the screening report itself requested.

Literature

The dispersion coefficient is the classical result of Taylor [39] and Aris [1]; the stabilised formulation is that of Brooks and Hughes [4].

Analytical solution

Yes for the pure-diffusion and uniform-inlet cases (the exact linear profile and trivial fixed point are both closed form); the Taylor-Aris item is a cross-check against this codebase's own already-validated dispersion_bvp_planar module, not an independent analytical solution.

Result

Category cross_check

Tolerance / bar pure diffusion: L2 error $\leq 1e-10$; uniform fixed point: L2 error $\leq 1e-8$; cross-check vs CFDCore: max pointwise diff $\leq 1e-6$; Taylor-Aris: rtol=0.02 (this study's own established first-order-discretization bar for a sampled-profile comparison)

Measured result all four checks passed

Verdict **PASS**

Traceability MF.8 SpeciesTransport2D | pure diffusion ($u=0$) recovers the exact linear profile (SpeciesTransport2D)

14.3 SpeciesTransport2D MMS convergence: ForwardDiff-manufactured source recovers exact rates for both P1 SUPG and P2 plain Galerkin

Purpose

A structural gap forced this module to exist: the older species solver had no source-term hook, so a manufactured-solution test was impossible there. Here a ForwardDiff-manufactured source on the unit square is refined ($n = 8, 16, 32$) and the measured convergence rate is compared with theory for streamline-stabilised P_1 and plain-Galerkin P_2 . A wrong stabilisation parameter or a mis-weighted source term would show as a degraded rate.

Literature

The stabilisation is the streamline-upwind Petrov–Galerkin method of Brooks and Hughes [4]; the order-of-accuracy procedure with manufactured sources follows Roache [32].

Analytical solution

Yes: a ForwardDiff-manufactured exact solution for the advection-diffusion species-transport equation.

Result

Category closed_form

Tolerance / **bar** measured order ≥ 1.8 (P1 SUPG), ≥ 2.8 (P2 Galerkin)

Measured result both bars met on the $n=8,16,32$ sequence

Verdict **PASS**

Traceability MF.8 SpeciesTransport2D | MMS: ForwardDiff-manufactured source recovers exact rates (SpeciesTransport2D)

15 ThinFilm

15.1 MF.13: Reynolds-equation thin-film flow on a 2D tagged mesh – parallel plates, linear-wedge slider bearing and squeeze strip against their closed forms

Purpose

COMSOL’s Microfluidics Module includes a thin-film flow interface for slender gaps, where a depth-averaged equation replaces a full three-dimensional solve. This case checks the implementation of that equation, the Reynolds lubrication equation for an incompressible, isoviscous film, against three closed-form solutions: the pressure-driven flow between parallel plates, the pressure in a linear-wedge slider bearing, and the pressure in a squeeze film. The reversed sliding direction is also checked, to confirm that a diverging film produces suction and that the pressure is not clipped.

Literature

The equation is Reynolds' theory of lubrication [30]. The parallel-plate, wedge and squeeze-strip solutions are the standard closed forms, derived in this study from the equation so that the comparison is exact; no measured bearing data are used.

Analytical solution

Parallel plates: linear p . Linear wedge: $p(x) = 6 \mu U/s (1/h - 1/h_1) + 6 \mu C/s (1/h^2 - 1/h_1^2)$, $C = -U h_1 h_2/(h_1 + h_2)$. Squeeze strip: $p = 6 \mu \dot{h}/h^3 (x^2 - L^2/4)$.

Result

Category closed_form

Tolerance / **bar** parallel plates to 1e-6; wedge bearing within 2e-3; squeeze strip to 1e-6

Measured result 11 of 11 checks

Verdict **PASS**

Traceability MF.13 ThinFilm | Reynolds equation: parallel plates, linear-wedge slider bearing and squeeze strip against their closed forms (ThinFilm)

Honest caveat. Incompressible, isoviscous, Newtonian; slender gap only; no inertia, no cavitation (the film pressure can go negative and is NOT clipped), no slip, no gap curvature. Pressure Dirichlet conditions are also enforced weakly (leak of about 1e-8 relative). One wedge geometry, one mesh (coarse: the wedge tolerance is set by it). Not compared with an experiment or with a full Stokes solve of the same gap, so the slenderness error of the Reynolds approximation itself was not measured.

16 Viscoelastic1D

16.1 MF.13: one-dimensional Oldroyd-B start-up channel flow – Newtonian-limit exact series, steady parabola, and the pure-elastic shear-wave arrival time

Purpose

COMSOL's Microfluidics Module offers viscoelastic fluid models; the full two- and three-dimensional versions need stabilisation methods for high Weissenberg numbers that are a research effort of their own. What is checked here is the one viscoelastic case that can be verified exactly, unidirectional start-up flow of an Oldroyd-B fluid between plates, against three limits: the Newtonian limit against its exact eigenfunction series, the steady parabola, and the arrival time of the shear wave in the purely elastic regime. It is a rheology check and not a viscoelastic flow solver.

Literature

The constitutive model is Oldroyd's [26]. In unidirectional shear flow its nonlinear terms vanish and it reduces to a linear relaxation equation for the shear stress, which is what is solved and checked here.

Analytical solution

Newtonian start-up: $u = G/(2\eta)(H^2/4 - y^2) + \sum \text{over odd } n \text{ of } b_n \cos(n\pi y/H) \exp(-(n\pi/H)^2 t/\rho)$; elastic shear-wave speed $\sqrt{\eta_p/(\rho\lambda)}$.

Result

Category closed_form

Tolerance / bar Newtonian limit within $1e-3$ of the exact series; steady parabola within $1e-5$; free acceleration exact until the wave arrival

Measured result 25 of 25 checks

Verdict **PASS**

Traceability MF.13 Viscoelastic1D | Oldroyd-B start-up channel flow: Newtonian limit, steady parabola and the elastic shear-wave speed (Viscoelastic1D)

Honest caveat. ONE-DIMENSIONAL and unidirectional ONLY: no normal-stress differences, no die swell, no 2D or 3D flow, no Weissenberg-number stabilisation (log-conformation, DEVSS), first order in time. It is a rheology check, not a viscoelastic flow solver, and says nothing about viscoelastic flow in a real microchannel or about other models (Giesekus, PTT, FENE-P). No experiment, no comparison with a published viscoelastic result.

References

- [1] R. Aris. On the dispersion of a solute in a fluid flowing through a tube. *Proceedings of the Royal Society of London. Series A*, 235(1200):67–77, 1956.
- [2] S. A. Berger, L. Talbot, and L. S. Yao. Flow in curved pipes. *Annual Review of Fluid Mechanics*, 15:461–512, 1983.
- [3] Thomas Borrvall and Joakim Petersson. Topology optimization of fluids in Stokes flow. *International Journal for Numerical Methods in Fluids*, 41(1):77–107, 2003.
- [4] A. N. Brooks and T. J. R. Hughes. Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 32(1–3):199–259, 1982.
- [5] H. Bruus. Acoustofluidics 7: The acoustic radiation force on small particles. *Lab on a Chip*, 12(6):1014–1021, 2012.
- [6] John W. Cahn and John E. Hilliard. Free energy of a nonuniform system. I. interfacial free energy. *The Journal of Chemical Physics*, 28(2):258–267, 1958.
- [7] P. Cherukat and J. B. McLaughlin. The inertial lift on a rigid sphere in a linear shear flow field near a flat wall. *Journal of Fluid Mechanics*, 263:1–18, 1994. Corrigendum: *J. Fluid Mech.* 285 (1995), 407.
- [8] Y. I. Cho and K. R. Kenney. Effects of the non-newtonian viscosity of blood on flows in a diseased arterial vessel. Part 1: Steady flows. *Biorheology*, 28(3–4):241–262, 1991.

- [9] W. R. Dean. The stream-line motion of fluid in a curved pipe (second paper). *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 5(30):673–695, 1928.
- [10] B. Derjaguin and L. Landau. Theory of the stability of strongly charged lyophobic sols and of the adhesion of strongly charged particles in solutions of electrolytes. *Acta Physicochimica URSS*, 14:633, 1941.
- [11] Dino Di Carlo, Daniel Irimia, Ronald G. Tompkins, and Mehmet Toner. Continuous inertial focusing, ordering, and separation of particles in microchannels. *Proceedings of the National Academy of Sciences*, 104(48):18892–18897, 2007.
- [12] A. Einstein. Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der Physik*, 322(8):549–560, 1905.
- [13] S. Elghobashi. On predicting particle-laden turbulent flows. *Applied Scientific Research*, 52(4):309–329, 1994.
- [14] Piotr Garstecki, Michael J. Fuerstman, Howard A. Stone, and George M. Whitesides. Formation of droplets and bubbles in a microfluidic T-junction—scaling and mechanism of break-up. *Lab on a Chip*, 6(3):437–446, 2006.
- [15] D. K. Gartling. A test problem for outflow boundary conditions—flow over a backward-facing step. *International Journal for Numerical Methods in Fluids*, 11(7):953–967, 1990.
- [16] U. Ghia, K. N. Ghia, and C. T. Shin. High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method. *Journal of Computational Physics*, 48(3):387–411, 1982.
- [17] L. P. Gor’kov. On the forces acting on a small particle in an acoustical field in an ideal fluid. *Soviet Physics Doklady*, 6:773–775, 1962.
- [18] B. P. Ho and L. G. Leal. Inertial migration of rigid spheres in two-dimensional unidirectional flows. *Journal of Fluid Mechanics*, 65(2):365–400, 1974.
- [19] David Jacqmin. Calculation of two-phase Navier–Stokes flows using phase-field modeling. *Journal of Computational Physics*, 155(1):96–127, 1999.
- [20] T. B. Jones. *Electromechanics of Particles*. Cambridge University Press, Cambridge, 1995.
- [21] David Kay, Daniel Loghin, and Andrew Wathen. A preconditioner for the steady-state Navier–Stokes equations. *SIAM Journal on Scientific Computing*, 24(1):237–256, 2002.
- [22] L. I. G. Kovásznyai. Laminar flow behind a two-dimensional grid. *Mathematical Proceedings of the Cambridge Philosophical Society*, 44(1):58–62, 1948.
- [23] M. R. Maxey and J. J. Riley. Equation of motion for a small rigid sphere in a nonuniform flow. *The Physics of Fluids*, 26(4):883–889, 1983.
- [24] Renwei Mei. An approximate expression for the shear lift force on a spherical particle at finite reynolds number. *International Journal of Multiphase Flow*, 18(1):145–147, 1992.
- [25] F. Mugele and J.-C. Baret. Electrowetting: from basics to applications. *Journal of Physics: Condensed Matter*, 17(28):R705–R774, 2005.

- [26] J. G. Oldroyd. On the formulation of rheological equations of state. *Proceedings of the Royal Society of London. Series A*, 200(1063):523–541, 1950.
- [27] S. Ookawara, R. Higashi, D. Street, and K. Ogawa. Feasibility study on concentration of slurry and classification of contained particles by microchannel. *Chemical Engineering Journal*, 101:171–178, 2004.
- [28] N. Pamme. Magnetism and microfluidics. *Lab on a Chip*, 6(1):24–38, 2006.
- [29] A. Ramos, H. Morgan, N. G. Green, and A. Castellanos. AC electrokinetics: a review of forces in microelectrode structures. *Journal of Physics D: Applied Physics*, 31(18):2338, 1998.
- [30] Osborne Reynolds. On the theory of lubrication and its application to Mr. Beauchamp Tower’s experiments, including an experimental determination of the viscosity of olive oil. *Philosophical Transactions of the Royal Society of London*, 177:157–234, 1886.
- [31] P. J. Roache. Perspective: A method for uniform reporting of grid refinement studies. *Journal of Fluids Engineering*, 116(3):405–413, 1994.
- [32] P. J. Roache. Code verification by the method of manufactured solutions. *Journal of Fluids Engineering*, 124(1):4–10, 2002.
- [33] Y. Saad and M. H. Schultz. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.
- [34] P. G. Saffman. The lift on a small sphere in a slow shear flow. *Journal of Fluid Mechanics*, 22(2):385–400, 1965. Corrigendum: *J. Fluid Mech.* 31 (1968), 624.
- [35] M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher. Benchmark computations of laminar flow around a cylinder. In *Flow Simulation with High-Performance Computers II*, volume 48 of *Notes on Numerical Fluid Mechanics*, pages 547–566. Vieweg, Wiesbaden, 1996. DFG Priority Research Programme “Flow Simulation with High-Performance Computers” results 1993–1995.
- [36] L. Schiller and A. Naumann. A drag coefficient correlation. *Zeitschrift des Vereins Deutscher Ingenieure*, 77:318–320, 1935.
- [37] G. Segré and A. Silberberg. Behaviour of macroscopic rigid spheres in Poiseuille flow. Part 2. Experimental results and interpretation. *Journal of Fluid Mechanics*, 14(1):136–157, 1962. Part 1 (concentration measurement method): *J. Fluid Mech.* 14(1), 115–135 (1962); original observation: *Nature* 189, 209–210 (1961).
- [38] D. Silvester and A. Wathen. Fast iterative solution of stabilised Stokes systems. Part II: using general block preconditioners. *SIAM Journal on Numerical Analysis*, 31(5):1352–1367, 1994.
- [39] G. I. Taylor. Dispersion of soluble matter in solvent flowing slowly through a tube. *Proceedings of the Royal Society of London. Series A*, 219(1137):186–203, 1953.