



Full length article

A mesh-free framework for high-order simulations of viscoelastic flows in complex geometries

J.R.C. King^{a,*}, S.J. Lind^b^a School of Engineering, University of Manchester, UK^b School of Engineering, Cardiff University, UK

ARTICLE INFO

Keywords:

Viscoelastic
Mesh-free
High-order
Numerical method
Log-conformation
LABFM

ABSTRACT

The accurate and stable simulation of viscoelastic flows remains a significant computational challenge, exacerbated for flows in non-trivial and practical geometries. Here we present a new high-order meshless approach with variable resolution for the solution of viscoelastic flows across a range of Weissenberg numbers. Based on the Local Anisotropic Basis Function Method (LABFM) of King et al. (2020), highly accurate viscoelastic flow solutions are found using Oldroyd B and PPT models for a range of two dimensional problems — including Kolmogorov flow, planar Poiseuille flow, and flow in a representative porous media geometry. Convergence rates up to 9th order are shown. Three treatments for the conformation tensor evolution are investigated for use in this new high-order meshless context (direct integration, Cholesky decomposition, and log-conformation), with log-conformation providing consistently stable solutions across test cases, and direct integration yielding better accuracy for simpler unidirectional flows. The final test considers symmetry breaking in the porous media flow at moderate Weissenberg number, as a precursor to a future study of fully 3D high-fidelity simulations of elastic flow instabilities in complex geometries. The results herein demonstrate the potential of a viscoelastic flow solver that is both high-order (for accuracy) and meshless (for straightforward discretisation of non-trivial geometries including variable resolution). In the near-term, extension of this approach to three dimensional solutions promises to yield important insights into a range of viscoelastic flow problems, and especially the fundamental challenge of understanding elastic instabilities in practical settings.

1. Introduction

Despite decades of research effort, the determination of accurate viscoelastic flow solutions remains a key challenge in computational rheology. There are a wide variety of techniques to improve numerical stability at higher levels of elasticity. Many early approaches sought to adjust the balance of elliptic and parabolic terms, with examples being elastic-viscous split-stress (EVSS) schemes [1,2], adaptive viscoelastic stress splitting (AVSS) [3,4] and “both sides diffusion” (BSD) [5,6]. Whilst these can provide increased stability, especially near the limit of zero solvent viscosity, they do not provide stability in complex transient flows at high levels of elasticity. Another approach, commonly used in pseudo-spectral methods (e.g. [7,8]) is to add some form of artificial diffusivity to the system. Whilst this can provide stability by imposing a limit on the smallest lengthscales of the flow, this is at the expense of accuracy, though justifications can be made by analogy with molecular diffusion. Perhaps the most significant development in the state-of-the-art has arisen from approaches which seek to transform the equations governing the evolution of the conformation tensor (or

polymeric stress), such that (some of the) physical constraints are respected. Two significant examples of this are the log conformation formulation [9,10], which guarantees the conformation tensor remains symmetric positive definite, and the Cholesky decomposition approach [11].

The above developments have improved capability greatly by increasing stability of simulations, especially as one enters higher Weissenberg, Wi , number regimes, where historically simulations would quickly fail approaching $Wi = \mathcal{O}(1)$. Whilst the stability of viscoelastic flow simulations has been greatly improved, *accuracy* of these simulations is now a primary concern. Stable simulations, particularly those at higher Wi number face considerable difficulty in attaining accurate, converged solutions, due to extremely large elastic stress gradients (for example near solid boundaries) and/or the development of very thin transient elastic stress filaments often a precursor to 2D or 3D (visco-)elastic instability and potentially the onset of elastic or elasto-inertial turbulence [12,13]. The resolution required to resolve the elastic stresses is considerable — and as the Wi number increases,

* Corresponding author.

E-mail address: jack.king@manchester.ac.uk (J.R.C. King).<https://doi.org/10.1016/j.jnnfm.2024.105278>

Received 20 December 2023; Received in revised form 2 May 2024; Accepted 10 June 2024

Available online 13 June 2024

0377-0257/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

obtaining fully converged solutions can become prohibitively expensive, with computational grid requirements dwarfing those required for the equivalent Newtonian turbulence simulation at the same Reynolds number. Given that understanding *Newtonian* turbulence remains one of the great open challenges in fluid mechanics, the challenge facing computational rheology in this regard considerable.

In order to attain high degrees of solution accuracy in a practical time frame, high-order methods become essential. Spectral, spectral element, and hp element methods have become established in computational rheology over the years [14], and there are many examples of their usage in solving a range of challenging viscoelastic flow problems and with a high degree of success (see for example, [15–20]). In simple domain geometries (i.e. rectangular) spectral methods have few competitors: relatively fast and extremely accurate they have been used with great success to model higher Wi number problems and fundamental flow studies in elastic turbulence (see [7,21], for example). For more practical contexts however – namely complicated geometries perhaps resembling industrial processing/mixing devices where accurate flow solutions have broad utility and benefit (e.g. [22]) – spectral methods are inapplicable. Spectral and p-finite element methods offer greater geometric flexibility, but, like mesh-based methods generally, constructing a mesh in a very complicated geometry that results in stable and converged solutions is particularly challenging and a significantly time-consuming task at pre-processing. Generally, element sizes and shapes have to sufficiently regular and well-distributed for accuracy and stability, which can be particularly difficult to achieve in very complicated geometries and makes an effective high-order adaptive or dynamic meshing scheme, e.g. for resolving thin transient elastic flow structures, particularly difficult to implement.

Meshless methods circumvent many of these challenges and make the process of domain discretisation much simpler in comparison. Meshless computational nodes have limited connectivity or requirements on topology and can often be scattered across a domain, then diffused or advected (by the flow or some transport velocity or otherwise) to improve node distribution. This may be done at pre-processing for Eulerian (fixed node) approaches or, for Lagrangian or Arbitrary-Lagrangian-Eulerian (ALE) simulations, during the simulation itself. Smoothed Particle Hydrodynamics is perhaps one of the most well know meshless methods and has been used to solve viscoelastic flow problems in the Lagrangian context for many years (see for example [23–26]), with impressive results in terms of stability, and showing considerable promise for free-surface flows. The computational nodes are simultaneously Lagrangian fluid elements, which can offer stability benefits in the context of viscoelastic flow simulation by effectively removing the advective term in the governing equations. There are also related methods, such as Dissipative Particle Dynamics (DPD) and Smoothed Dissipative Particle Dynamics (SDPD), which are also subject to a concerted research effort in computational rheology [27–31], but these tend to apply on the physical length scales where the search for converged continuum solutions becomes less relevant.

While offering enviable geometric flexibility and stability, one of the primary limitations of SPH and related approaches is accuracy. In its traditional form the SPH method is formally low order [32], and as discussed above, without high-order resolving power, the resolutions required for practical simulation become prohibitively expensive (especially as SPH is more computationally expensive than most grid-based methods at an equivalent resolution). One of the key benefits of SPH is that it is (usually) Lagrangian - particles move (approximately) with the flow, advection terms are included implicitly through the particle motion, and the simulation of flows with dynamically evolving free surfaces or topological changes is straightforward. We note that the relatively low accuracy of SPH is a consequence of the smoothing kernel, and even for simple shear flows, is largely unaffected by the frame of reference (i.e. Eulerian or Lagrangian) [26]. An additional strength of SPH is that in the context of viscoelastic flows, it is possible that the Lagrangian nature of SPH provides increased stability at high

Wi [25,26]. With the above limitations and strengths in mind, the computational rheology community would benefit from a method that is both meshless (to provide the improved geometric flexibility) and high-order (to provide the accuracy and resolving power). In recent years the authors and co-workers have been developing such a method; originally motivated by the need to create a high-order version of the SPH method [33,34], the Local Anisotropic Basis Function Method (LABFM) has emerged as a generalised high-order meshless scheme [35, 36] with arbitrary orders of convergence possible (but 6th or 8th order spatial convergence is typical). Following the analysis of prototypical Newtonian flow cases in [35,36], LABFM has recently been extended to the study of combustion physics and flame-turbulence interactions in complex geometries in [37].

The potential of the LABFM approach as a geometrically flexible, multi-physics, high-order solver is such that the aim and focus of this manuscript is the extension of LABFM to the solution of viscoelastic fluid flow. Herein we demonstrate high-order solutions of viscoelastic flow in both simple and non-trivial geometries with convergence rates of up to 9th order possible. We also consider three numerical approaches for the viscoelastic stresses (based on direct integration, Cholesky decomposition, and log conformation) to assess suitability in this new high-order mesh-free context for different test cases. The paper concludes with a preliminary study of a two-dimensional symmetry breaking elastic instability in a representative porous media geometry at moderate Wi , demonstrating the potential of high-order meshless schemes for the fundamental study of elastic instabilities and elastic turbulence in non-trivial geometries in the near future. It is hoped that in the longer term the method may serve as a practical tool for the computational analysis, optimisation and design of challenging industrial viscoelastic fluid processing, an activity which underpins healthcare product and foodstuff manufacturing, energy supply, and many other important industries worldwide.

The remainder of this paper is set out as follows. In Section 2 we introduce the governing equations, and their Cholesky and log-conformation formulations, and in Section 3 we describe the numerical implementation. Section 4 contains a set of numerical results providing validation for the model against two-dimensional Kolmogorov flow, Poiseuille flow, and flows past cylinders in a channel and representative porous media geometry. Section 5 is a summary of conclusions.

Before continuing further, we briefly comment on our notation. To avoid ambiguity, we use Einstein notation where possible, and, of the Latin characters, subscripts i, j, k, l, n are reserved for this purpose, with repetition implying summation. Subscripts a and b are used for particle/node indexes. Bold fonts are used to refer to tensors in their entirety (e.g. c) rather than individual components. The order of the spatial discretisation scheme is denoted by m .

2. Governing equations

In the present work we limit our focus to the two-dimensional problem. The governing equations for the density, momentum and conformation tensor (in Einstein notation) are

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho u_k}{\partial x_k} = 0, \quad (1)$$

$$\frac{\partial \rho u_i}{\partial t} + \frac{\partial \rho u_i u_k}{\partial x_k} = -\frac{\partial p}{\partial x_i} + \beta \eta \frac{\partial^2 u_i}{\partial x_k \partial x_k} + \frac{(1-\beta)\eta}{\lambda} \frac{\partial c_{ki}}{\partial x_k} + \rho f_i, \quad (2)$$

$$\frac{\partial c_{ij}}{\partial t} + u_k \frac{\partial c_{ij}}{\partial x_k} - \frac{\partial u_i}{\partial x_k} c_{kj} - \frac{\partial u_j}{\partial x_k} c_{ik} = -\frac{(c_{ij} - \delta_{ij})}{\lambda} [1 - 2\epsilon + \epsilon c_{kk}], \quad (3)$$

in which x_i is the i th coordinate ($x_i = x, y$ for $i = 1, 2$), u_i is the i th component of velocity (u, v for $i = 1, 2$), ρ is the density, p the pressure, c_{ij} the ij -th element of the conformation tensor, f_i is a body force, λ the polymer relaxation time, η the total viscosity, β the ratio of solvent to total viscosity, and ϵ is a non-linearity parameter. The system is closed with an isothermal equation of state $p = c_s^2 (\rho - \rho_0)$,

where ρ_0 is a reference density and c_s is the sound speed. Taking U and L to be characteristic velocity and length scales, and $T = L/U$ the characteristic time scale, the governing dimensionless parameters are:

$$Re = \frac{\rho_0 U L}{\eta}; \quad Wi = \frac{\lambda U}{L}; \quad Ma = \frac{U}{c_s}; \quad (4)$$

the viscosity ratio β , and the PTT nonlinearity parameter ϵ . Additional terms which might arise in (2) and (3) due to compressibility (see e.g. [38,39]) are neglected as we operate near the incompressible limit, with Mach number $Ma < 0.05$ for all considered cases. In this paper, three different formulations are investigated for the conformation tensor evolution equation (3) to inform optimal use in the high-order meshless context across test cases. In particular, we employ direct numerical integration of (3), Cholesky Decomposition [11], and log-conformation [9,10], summaries of which are provided in the sections below.

2.1. Cholesky decomposition

First considered in the context of viscoelastic numerical simulation by [11], the Cholesky decomposition of the conformation tensor offers a convenient way to maintain symmetric positive definiteness. Consider the Cholesky decomposition of the 2-D conformation tensor, viz.,

$$c = \begin{bmatrix} l_{11}^2 & l_{11}l_{12} \\ l_{11}l_{12} & l_{12}^2 + l_{22}^2 \end{bmatrix}, \quad (5)$$

with l_{ij} denoting the components of the lower triangular matrix L , such that $c = LL^T$. Defining S_{ij} as the right hand side of (3)

$$S_{ij} = -\frac{(c_{ij} - \delta_{ij})}{\lambda} [1 - 2\epsilon + \epsilon c_{kk}], \quad (6)$$

then the equations for the evolution of the Cholesky decomposition components are:

$$\frac{\partial l_{11}}{\partial t} + u_k \frac{\partial l_{11}}{\partial x_k} = \frac{\partial u}{\partial x} l_{11} + \frac{\partial u}{\partial y} l_{12} + \frac{S_{11}}{2l_{11}} \quad (7a)$$

$$\frac{\partial l_{12}}{\partial t} + u_k \frac{\partial l_{12}}{\partial x_k} = \frac{\partial v}{\partial y} l_{12} + \frac{\partial u}{\partial y} l_{22} + \frac{\partial v}{\partial x} l_{11} + \frac{S_{12}}{l_{11}} - \frac{l_{12}}{2l_{11}^2} S_{11} \quad (7b)$$

$$\frac{\partial l_{22}}{\partial t} + u_k \frac{\partial l_{22}}{\partial x_k} = \frac{\partial v}{\partial y} l_{22} - \frac{\partial u}{\partial y} \frac{l_{12}l_{22}}{l_{11}} + \frac{S_{22}}{2l_{22}} - \frac{S_{12}l_{12}}{l_{22}l_{11}} + \frac{l_{12}^2 S_{11}}{2l_{11}^2 l_{22}} \quad (7c)$$

Eqs. (7a)–(7c) can then be discretised in space and integrated in time as detailed in Section 3. We refer to this formulation as CH. It is common practice to evolve the natural logarithms of l_{11} and l_{22} . The corresponding evolution equations can be obtained by simply dividing (7a) by l_{11} and (7c) by l_{22} (with (7b) unchanged). We refer to this formulation as Cholesky-log, or CH-L.

2.2. Log-conformation

The log-conformation formulation employed directly follows that of [9,10], to which readers may refer for further details. We define the diagonalisation of the conformation tensor as

$$c_{ij} = R_{ik} \Lambda_{kl} R_{jl}, \quad (8)$$

in which R_{ij} contains the eigenvectors of the conformation tensor, and the diagonal matrix Λ_{ij} contains the eigenvalues. We then denote the log-conformation tensor as

$$\Psi_{ij} = R_{ik} \ln \Lambda_{kl} R_{jl}, \quad (9)$$

where the logarithm is applied independently to each diagonal element Λ_{ii} . The log-conformation tensor is then evolved according to

$$\frac{\partial \Psi_{ij}}{\partial t} + u_k \frac{\partial \Psi_{ij}}{\partial x_k} = \Omega_{ik} \Psi_{kj} - \Psi_{ik} \Omega_{kj} + 2B_{ij} - \frac{1}{\lambda} (1 - 2\epsilon + \epsilon c_{kk}) \{c^{-1}\}_{ij} (c_{ij} - \delta_{ij}), \quad (10)$$

where

$$\Omega_{ij} = R_{ik} (\omega_{kn} (1 - \delta_{nl})) R_{jl}, \quad (11)$$

and

$$B_{ij} = R_{ik} (m_{kn} \delta_{nl}) R_{jl}, \quad (12)$$

in which

$$\omega_{ij} = (\Lambda_j m_{ij} + \Lambda_i m_{ji}) / (\Lambda_j - \Lambda_i) \quad (13)$$

and

$$m_{ij} = R_{ki} \frac{\partial u_k}{\partial x_j} R_{lj}. \quad (14)$$

The subsequent spatial and temporal discretisation of the log-conformation scheme is described in Section 3. As mentioned, both the Cholesky decomposition and log-conformation approaches will be compared with a scheme that directly integrates the conformation tensor evolution Eq. (3), with no explicit constraints given on tensor positive definiteness. The aim will be to compare the three schemes across different flow test cases to determine optimal usage in the high-order meshless framework.

3. Numerical implementation

For all three formulations, the numerical implementation closely follows that described in [36] for isothermal Newtonian flows and [37] for turbulent reacting flows, with spatial discretisation based on the Local Anisotropic Basis Function Method (LABFM), time integration via an explicit Runge–Kutta scheme, and acceleration via OpenMP and MPI, with non-uniform structured block domain decomposition.

3.1. Spatial discretisation

The spatial discretisation is based on LABFM, which has been detailed and extensively analysed in [35,36], and we refer the reader to these works for a complete description, but provide sufficient details here to reproduce the method. Briefly summarising, the domain is discretised with a point cloud of N nodes, unstructured internally, and with local structure near boundaries. The node-set is fixed in space; i.e. the nodes do not move during the simulation. Along wall, inflow and outflow boundaries, nodes are distributed uniformly. Near boundaries, an additional 4 rows of uniformly distributed nodes are arranged along boundary normals originating at nodes on boundaries. These are used to construct one-sided difference operators on the boundaries. Internally, the node distribution is generated using the propagating front algorithm of [40]. The generation of the node distribution is akin to a mesh-generation procedure for unstructured mesh-based methods, and we do not repeat details of the algorithm here, which has been described in [36,37]. To ensure repeatability of our work, for each case herein, node distributions are available on request.

Each node a has an associated resolution s_a , which corresponds to the local average distance between nodes. The resolution need not be uniform, and s_a may vary with x and y . Each node also has an associated computational stencil length-scale h_a , which again may vary with x and y . The ratio s_a/h_a is approximately uniform, with some variation due to the stencil optimisation procedure described in [36], which ensures the method uses stencils marginally larger than the smallest stable stencil. Each node holds the evolved variables ρ_a , $\rho u_{i,a}$, and either the components of c , the components of Ψ , or the components of the L . The governing equations are solved on the set of N nodes.

The difference between properties at two nodes is denoted $(\cdot)_{ba} = (\cdot)_b - (\cdot)_a$. The computational stencil for each node a is denoted \mathcal{N}_a , and is constructed to contain all nodes b such that $r_{ba}^2 = x_{ba}^2 + y_{ba}^2 \leq 4h_a^2$. A schematic of the computational stencil is shown in Fig. 1.

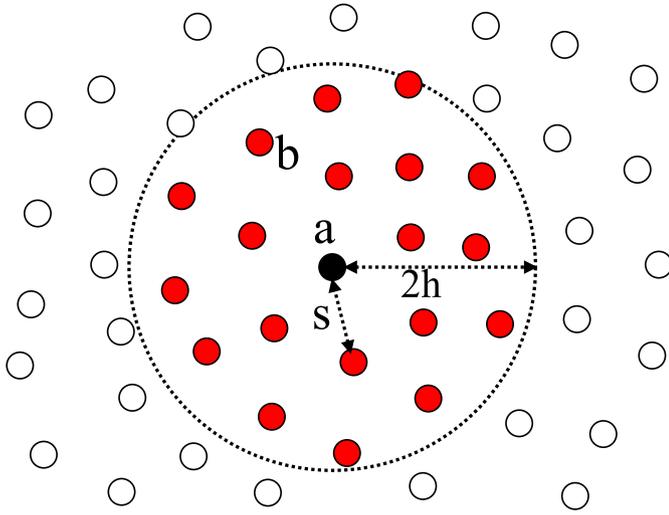


Fig. 1. A schematic of the computational stencil.

In the following, we use standard multi-index notation with index (e.g.) α representing the ordered pair (α_1, α_2) , with $|\alpha| = \alpha_1 + \alpha_2$. For clarity, as is standard in multi-index notation $\alpha! = \alpha_1! \alpha_2!$, $x^\alpha = x_1^{\alpha_1} x_2^{\alpha_2}$, and the partial derivative

$$\partial^\alpha = \frac{\partial^{\alpha_1}}{\partial x_1} \frac{\partial^{\alpha_2}}{\partial x_2}. \quad (15)$$

In LABFM, all spatial derivative operators take the form

$$L_a^\gamma(\cdot) = \sum_{b \in \mathcal{N}_a} (\cdot)_{ba} w_{ba}^\gamma, \quad (16)$$

where γ is a multi-index which identifies the derivative being approximated by (16), and w_{ba}^γ are a set of inter-node weights for the operator. To evaluate w_{ba}^γ , we first define the vector of monomials X_{ba} element-wise, with the element corresponding to multi-index α

$$X_{ba} = \frac{x_{ba}^\alpha}{\alpha!}, \quad (17)$$

and a vector of anisotropic basis functions $W_{ba} = W(x_{ba})$, with the element corresponding to multi-index α given by

$$W_{ba} = \frac{\psi(r_{ba}/h_a)}{\sqrt{2^{|\alpha|}}} H_\alpha\left(\frac{x_{ba}}{h_a \sqrt{2}}\right), \quad (18)$$

where

$$H_\alpha(x) = H_{\alpha_1}(x_1) H_{\alpha_2}(x_2) \quad (19)$$

are bi-variate Hermite polynomials of the physicists' kind, and the radial basis function (RBF) ψ is a Wendland C2 kernel [41]. The weights w_{ba}^γ in (16) are constructed as

$$w_{ba}^\gamma = W_{ba} \cdot \Psi^\gamma, \quad (20)$$

with Ψ^γ a vector to be determined. To determine Ψ^γ we construct and solve the linear system

$$\left[\sum_{b \in \mathcal{N}_a} X_{ba} \otimes W_{ba} \right] \cdot \Psi^\gamma = C_\gamma, \quad (21)$$

in which C_γ is a unit vector defined element-wise as $C_\gamma^\alpha = \delta_{\alpha\gamma}$, with $\delta_{\alpha\gamma}$ the Dirac-delta function. The consistency of the operator (16) is then determined by the size of the linear system (21). If we include the first $M = \frac{m^2+3m}{2}$ terms (i.e. up to order $|\alpha| \leq m$) then the operator has polynomial consistency of order m . Consequently, first derivative operators converge with s^m , and second derivatives with s^{m-1} . With

our nodes fixed in space, as a preprocessing step, for each node we construct and solve the linear system (21) to obtain Ψ^γ , for γ corresponding to both first spatial derivatives, and the Laplacian, which we then use to calculate and store w_{ba}^γ in (16). The derivatives appearing in (1), (2), and either (3), (7a) to (7c), or (10) are approximated using (16).

We highlight here that the consistency correction procedure described above removes the *discretisation error limit* which causes a saturation of convergence in low-order mesh-free methods, such as SPH (see e.g. [32] for a discussion). In such methods, this limit is caused by the fact that the kernel does not account for the node distribution, and the derivation of the method assumes the equivalence of the integral of the kernel over its support with the sum of the kernel over the particles in a stencil. No such assumptions are made in the present method, and the operators converge until an error limit dictated by the accuracy with which (21) can be solved, typically $\mathcal{O}(10^{-12})$ for order $m = 8$. For in-depth analysis of the LABFM discretisation, we refer the reader to [35,36].

The order of the spatial discretisation can be specified between $m = 4$ and $m = 10$, and although there is capability for this to be spatially (or temporally) varying, in this work we set $m = 8$ uniformly away from boundaries (except where explicitly stated in our investigations of the effects of changing m). Whilst a larger value of m gives greater accuracy for a given resolution, it also requires a larger stencil (larger h_a/s_a), and hence incurs greater computational expense for a given resolution. As in central finite difference schemes, the cost scales with the product of the stencil size and the number of points $\mathcal{N}N$, where \mathcal{N} is the average number of nodes in a computational stencil. In the present implementation, parallelised with MPI, scaling results have shown parallel efficiency above 96% between 4 and 1024 cores. The method exhibits the same scaling performance as our related work in [37]. A value of $m = 8$ provides a good compromise between accuracy and computational cost. At non-periodic boundaries, the consistency of the LABFM reconstruction is smoothly reduced to $m = 4$. The stencil scale h_a is initialised to $h_a = 2.7s_a$ in the bulk of the domain, and $h_a = 2.4s_a$ near boundaries (choices informed through experience as being large enough to ensure stability with $m = 8$). In the bulk of the domain, the stencil scale is then reduced following the optimisation procedure described in [36]. This has the effect of both reducing computational costs, and increasing the resolving power of LABFM, and accordingly h_a takes a value slightly larger than the smallest value for which the discretisation remains stable. Note that as with high-order central finite differences, or pseudo-spectral methods, no upwinding is used in the present scheme.

3.2. Temporal discretisation

For the direct integration of c (hereafter referred to as DI), we evolve only 3 of the 4 (in two-dimensions) components of c , and impose symmetry. There is nothing in this formulation to ensure c remains positive definite. Both the log-conformation (referred to as LC) and Cholesky decomposition (referred to as CH) formulations ensure c is symmetric positive definite by construction.

Time integration is by explicit third order Runge–Kutta scheme. We use a low-storage four-stage Runge–Kutta scheme with an embedded second order error estimator, with the designation RK3(2)4[2R+J]C in the classification system of [42]. The value of the time-step is controlled using a Proportional-Integral-Derivative (PID) controller as described in [37], which ensures errors due to time-integration remain below 10^{-4} . In addition to the PID controller, we impose an upper limit on the time step with

$$\delta t \leq \delta t_{\max} = \max(\delta t_{cfl}, \delta t_{visc}) \quad (22)$$

where $\delta t_{cfl} = \min(s/(u + c_s))$ and $\delta t_{visc} = \min(C_{visc} \eta s^2)$ denote the time steps due to the CFL and viscous diffusion constraints, respectively (with the min taken over the entire domain). We set the coefficient C_{visc}

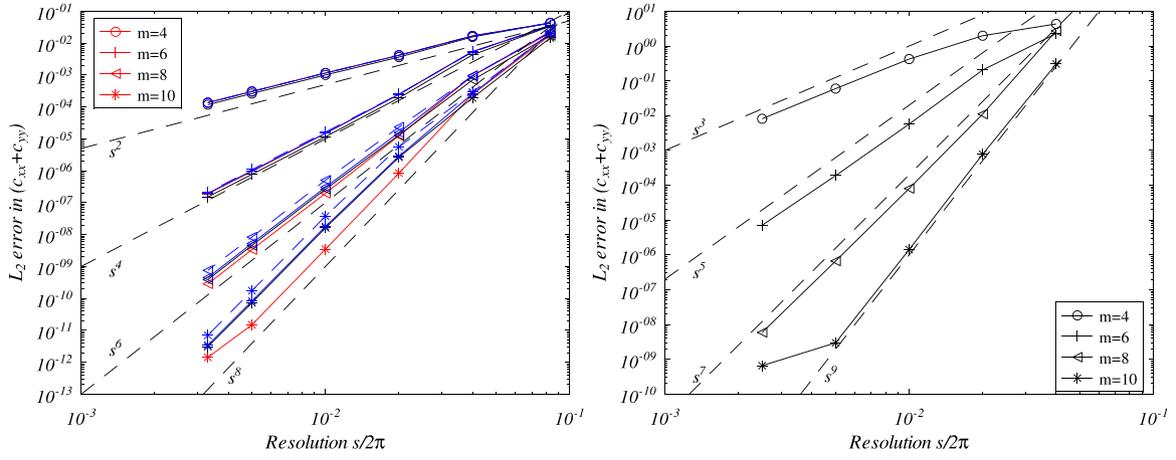


Fig. 2. Kolmogorov flow. Variation of L_2 error in steady state conformation tensor trace $c_{xx} + c_{yy}$, for different orders m of the numerical discretisation scheme. The left panel shows the errors for all formulations DI (red lines), CH (solid blue lines), CH-L (dashed blue lines) and LC (black lines), for $Re = 1$ and $Wi = 0.1$. The right panel shows the errors for DI only, for $Re = 10$ and $Wi = 1$.

to the limiting value $C_{visc} = 1$, with the PID controller reducing the time-step size as necessary to keep time-integration errors bounded.

We note here that because we use an explicit time integration scheme, at small Re , $\delta t \propto Re$, making simulations in the creeping flow limit prohibitively expensive. An implicit time-integration scheme could be implemented with the present spatial discretisation scheme. This would remove the viscous constraint on the time-step, at the expense of the solution of a large sparse linear system every step. Implicit schemes have been implemented in Smoothed Particle Hydrodynamics [26,43], which although low-order and Lagrangian, has effectively the same stencils as the present method, including on massively parallel architectures [44], and for high-order variants on multiple GPUs [45]. In Lagrangian schemes the linear system must be reconstructed each time-step, incurring considerable cost. In the present method, where the nodes are fixed, the linear system would need to be built only once. The development of an implicit, incompressible formulation of our numerical framework is planned, but beyond the scope of the present work.

In common with other high-order collocated methods, the discretisation admits solutions with energy at the wavenumber of the resolution, and some form of de-aliasing or filtering is required. In the present work, the solution is de-aliased every time-step by a high-order filter, as is commonly used for high-order central finite differences applied to compressible flows. For a field ϕ , the filtering procedure is defined as

$$\hat{\phi} = (1 + \mathcal{F}_a) \phi, \quad (23)$$

in which the operator \mathcal{F}_a is

$$\mathcal{F}_a(\phi) = \kappa_{m,a} \sum_{b \in \mathcal{N}_a} \phi_{ba} w_{ba}^{\gamma_m}, \quad (24)$$

where γ_m is a multi-index such that the operator using weights $w_{ba}^{\gamma_m}$ approximates ∇^m , and the pre-factor $\kappa_{m,a}$ is calculated in a pre-processing step as

$$\kappa_{m,a} = \frac{2}{3} \left\{ \sum_{b \in \mathcal{N}_a} \left[1 - \cos\left(\frac{3\pi x_{ba}}{2s_a}\right) \cos\left(\frac{3\pi y_{ba}}{2s_a}\right) \right] w_{ba}^{\gamma_m} \right\}^{-1}. \quad (25)$$

By construction, $\kappa_{m,a}$ ensures that the amplitude responses of the filter at a wavenumber two thirds of the Nyquist wavenumber (defined by the resolution s) is equal to $1/3$. Further details of the procedure can be found in [36,37]. The filter primarily acts on large wavenumbers, and has little effect on wavenumbers which are small relative to the Nyquist wavenumber of the discretisation (π/s). Provided the flow field

is sufficiently resolved that the physical information in the solution lies at small wavenumbers (relative to π/s), the effect of the filter on the physical solution is negligible. In the present work, the filter (23) is applied every time step to the density and momentum fields, to either the components of c (for DI), the components of Ψ (LC) or the components of L (CH).

3.3. Boundary treatment

Throughout this work, the computational domain is discretised with a strip of uniformly arranged nodes near solid boundaries, and an unstructured node-set internally. The discretisation procedure is the same as that used in [36,37], to which we refer the interested reader for details. As in those works, numerical boundary conditions for no-slip walls are implemented via the Navier–Stokes Characteristic boundary condition formalism following [46], the implementation of which directly follows [36]. For the conformation tensor (or its decompositions), the hyperbolic term is zero on solid boundaries, and requires no additional treatment. The upper-convected terms may be directly evaluated using the values of c (or its decompositions) on the boundary, and the local velocity gradients (evaluated using one-sided derivatives as detailed in [36]).

4. Numerical results

For the following, we reiterate use of the acronyms DI, CH, CH-L and LC to indicate direct integration of the conformation tensor, Cholesky decomposition, log-Cholesky decomposition, and the log-conformation formulation, respectively. Where a figure provides a comparison of these formulations, we use red, black and blue lines for DI, LC and CH respectively.

4.1. Kolmogorov flow

Our first test is two-dimensional Kolmogorov flow. Although a simple geometry and one for which pseudo-spectral methods are well suited, the flow has an analytic solution in the steady state, and given the absence of boundaries, is a good test of the convergence of our method. The domain is a doubly-periodic square with side length 2π . A forcing term of $f_x = \frac{4}{Re} \cos(2y)$ and $f_y = 0$ is added to the right hand side of the momentum equation. The analytic solution in the steady state is $u = \cos(2y)$, $v = 0$, $\rho = 1$ and

$$c = \begin{bmatrix} 1 + 8Wi^2 \sin^2(2y) & -2Wi \sin(2y) \\ -2Wi \sin(2y) & c_{yy} = 1 \end{bmatrix}. \quad (26)$$

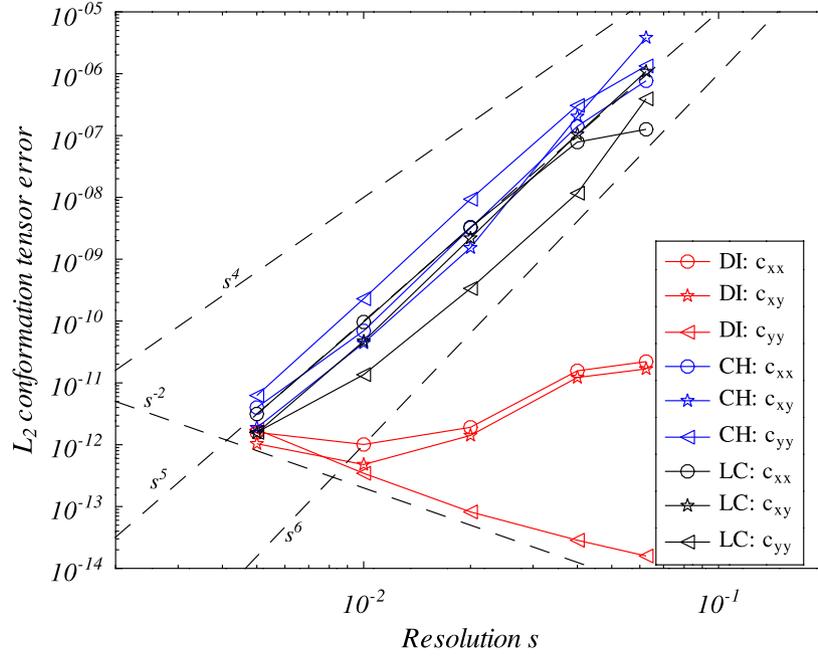


Fig. 3. Variation of the Poiseuille flow steady state L_2 error norm of the conformation tensor components with resolution for parameters $Re = 1$, $Wi = 1$, $\epsilon = 0$, $\beta = 1$ and $Ma = 0.05$, for the three formulations: DI - red lines, CH - blue lines, LC - black lines. The annotations indicate the slopes of the dashed lines.

We first set $Re = 1$, $\beta = 0.5$, $\epsilon = 0$, $Ma = 0.05$. To reduce the costs of reaching a steady state, we use a small value of $Wi = 0.1$. Note that the purpose of this test is not to demonstrate the ability of the method to reach larger Wi , but to assess the accuracy of the numerical scheme. The convergence rates observed would be unaffected if we set $Wi = 1$, but the simulation would take longer to reach a steady state. We vary the order m of the discretisation, with $m \in [4, 6, 8, 10]$. The left panel of Fig. 2 shows the convergence in the trace of the conformation tensor for all three formulations for $Re = 1$ and $Wi = 0.1$. The errors scale with s^{m-2} . At higher m , DI is slightly more accurate than LC and CH, and CH-L is least accurate, but at lower m the errors for all formulations are almost identical.

For lower Re , the time-step selection criteria is such that we have $\delta t \propto s^2$, and, by considering a steady state at a fixed time t_{end} , the total number of time-steps to reach t_{end} is $N_{steps} \propto \delta t^{-1} \propto s^{-2}$. The total error at t_{end} is given by the spatial error, dominated by terms $\propto s^m$ in the present case, multiplied by the number of time steps, resulting in a scaling of s^{m-2} , as shown in the convergence rates of the left panel of Fig. 2. The right panel of Fig. 2 shows convergence in trc at larger Re (and Wi , with $Re = 10$ and $Wi = 1$), but presenting the DI formulation only. Note that as the time step is now proportional to s (rather than s^2), given the larger Re , convergence rates tend to follow order $m - 1$, as observed in [36]. For the right panel of Fig. 2 errors are taken at dimensionless time $t = 40/\pi$, after the steady state is reached, but before the growth of any instability. The magnitudes of the errors are larger in this case because Wi is larger (with similar behaviour seen in the Poiseuille flow case to follow in Section 4.2). For a fixed Re , increasing Wi will increase the error magnitude, but leave the rate of convergence unchanged.

For both panels in Fig. 2 we see slightly lower convergence rates at very coarse resolutions because the wavelengths on which the high-order filters act are closer to those of the base solution. This behaviour has been observed and discussed in [36], and further in the context of three-dimensional turbulence simulations in [37]. If filtering for the conformation tensor equation is removed, this *does not* affect the order of convergence, but does reduce the magnitude of the errors by $\mathcal{O}(10^3)$.

However, at coarse resolutions, filtering is essential for stability with c becoming unbounded in the absence of filtering for all formulations (DI, CH, CH-L, LC).

4.2. Poiseuille flow

Poiseuille flow is an important and classical flow test admitting analytical solutions for Oldroyd B fluids in the unsteady transient start-up flow as well as at steady state, and hence provides a good test of method accuracy. For all Poiseuille flow tests considered, the domain is a unit square, periodic laterally with no-slip wall boundaries at the top and bottom. The flow is driven by a constant and uniform body force such that the channel centreline velocity in the steady state is unity.

4.2.1. $Re = 1$, $Wi = 1$, $\epsilon = 0$, $\beta = 1$, $Ma = 0.05$, Varying resolution s

This selection of parameters decouples the momentum equation from the conformation tensor, and is a good validation of the accuracy of our discretisation. Fig. 3 shows the convergence with resolution of the L_2 norm of the conformation tensor components once a steady state has been reached. For direct integration of the conformation tensor, results are much more accurate. At these low Re , the time-step is limited by the viscous constraint, and so $\delta t \propto s^2$. This explains the increasing error from machine precision with resolution, clearly visible for c_{yy} in particular — there is an accumulation of time-stepping errors of order s^{-2} . For LC and CH approaches, we see more typical convergence behaviour of order 5 with resolution s . Indeed, the DI approach would also exhibit 5th order convergence if the errors were not already so close to machine precision, and this is shown in subsequent results in the next section.

The 5th order convergence behaviour observed follows principally from the boundary conditions employed. Although the finite difference stencils used on the wall boundary are 4th order, the dominant error in this Poiseuille flow arises mainly due to errors in the nonlinear advection terms, which are larger in the near-wall region, where the order of the LABFM discretisation is reduced from $m = 8$ to $m = 6$.

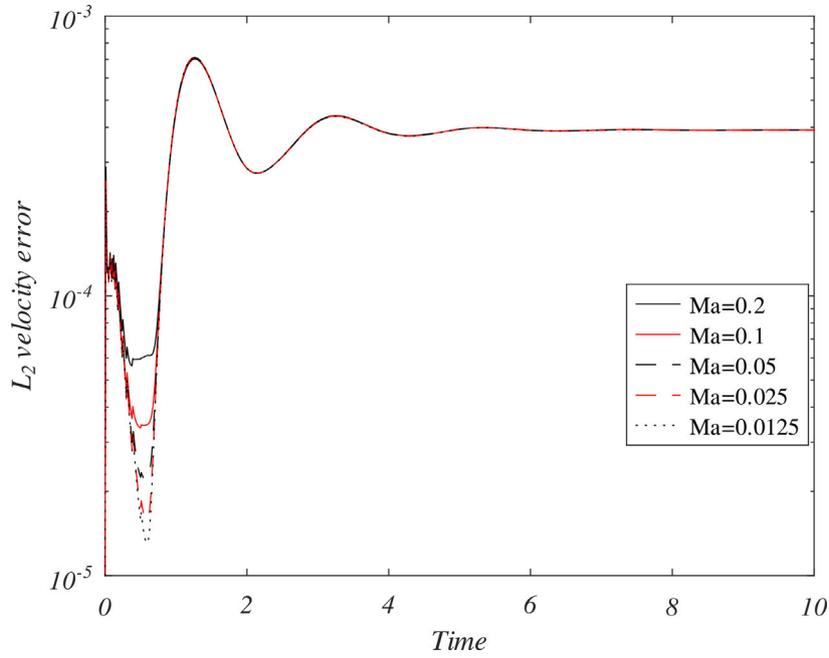


Fig. 4. Poiseuille flow: Time evolution of the L_2 error in the velocity for several values of Ma , using log-conformation formulation (denoted LC), with a resolution of $s = 1/50$.

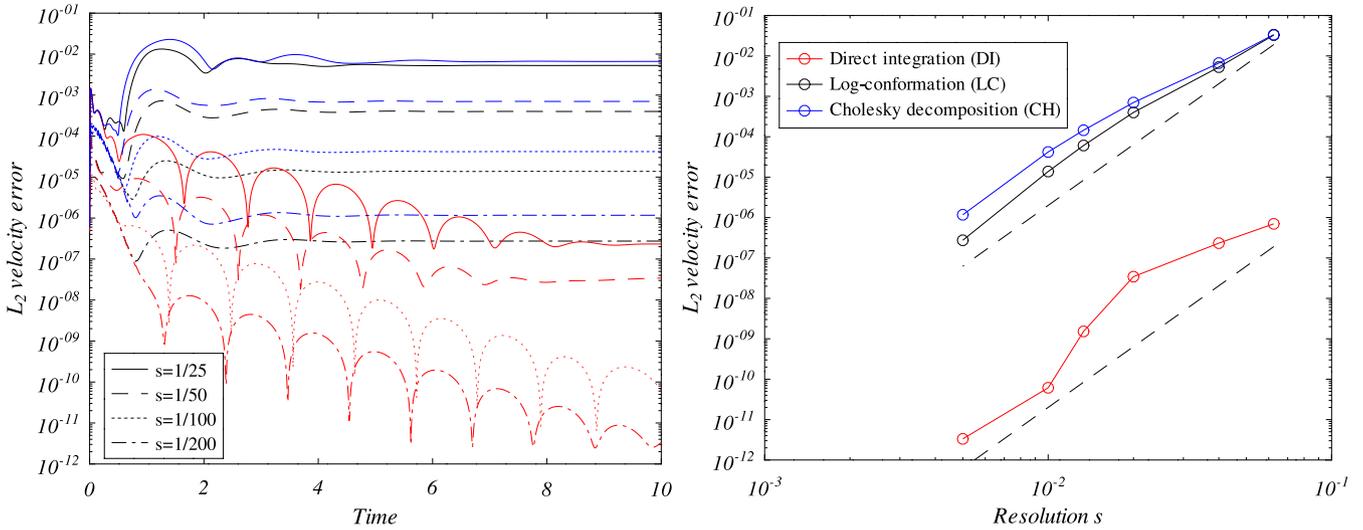


Fig. 5. Poiseuille flow. Left panel: Time evolution of the L_2 error in the velocity for all three formulations (DI - red lines, CH - blue lines, LC - black lines), for several resolutions s . Right panel: Convergence of the L_2 error in the velocity at time $t = 10$ with resolution for all three formulations. The dashed lines show convergence rates of order 5.

The elements of c have quadratic and linear forms, and therefore for DI, their derivatives are reproduced exactly. Due to the Cholesky- and log- transformations, the elements of both L and Ψ have more complex forms, and first derivatives are accurate to order m , whilst second derivatives are accurate to order $m-1$. The order $m = 6$ consistency near the walls, combined with the accumulation of time-stepping errors, results in the observed convergence rates of order 5.

4.2.2. $Re = 1, Wi = 1, \epsilon = 0, \beta = 0.1$, Varying Ma and resolution s

We now demonstrate the effect of varying Mach number on the solution. At steady state, the solution has uniform pressure and so offers a good test on the role of Mach number as compressibility should not affect the final solution. This is confirmed by Fig. 4 showing the time evolution of the L_2 error in the velocity for several values of Ma using LC and a resolution $s = 1/50$. There are small error

differences in the transient flow at early times, but the steady state is unaffected. Accordingly, for the remainder of this section, we use $Ma = 0.05$. In [36] we also showed the effects of changing Ma but in the Newtonian setting by comparing to analytical solutions for Taylor Green vortices, and also found $Ma = 0.05$ to be a suitable value. However, viscoelastic simulations are invariably more complex and some care does need to be taken at larger Wi , where the value of Ma can affect stability of the simulation, particularly if Ma is too small, and also at very low Re , where large viscous stresses can require smaller Ma to ensure negligible compressibility. This behaviour is described further below.

The left panel of Fig. 5 shows variation of L_2 error in velocity with time for several resolutions, for direct integration of the conformation tensor (DI), Cholesky-decomposition (CH) and the log-conformation

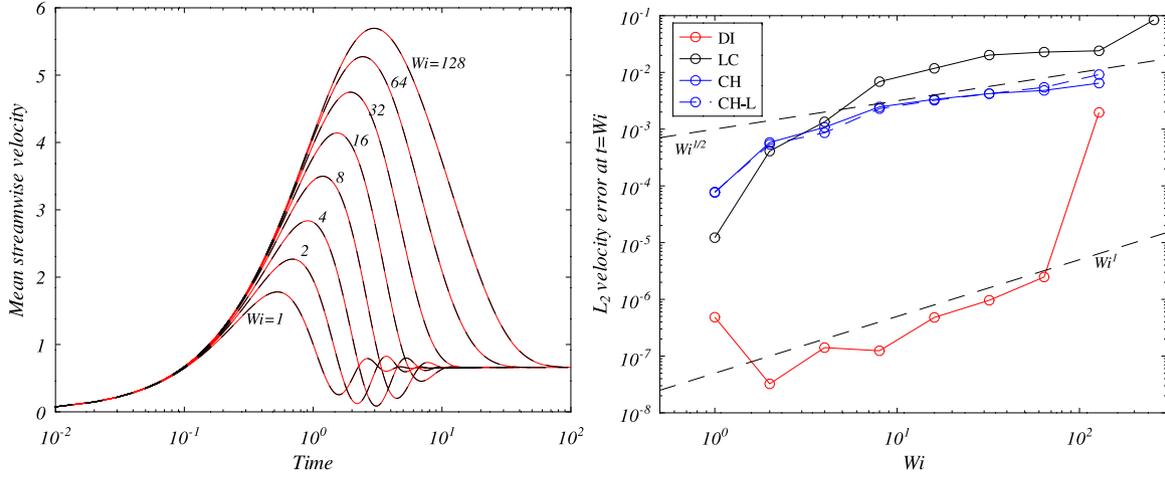


Fig. 6. Poiseuille flow. Left panel: Time evolution of the mean velocity for a range of Wi , obtained with the LC formulation. Red lines indicate the analytic solution, dashed black lines indicate the numerical results. Right panel: Variation with Wi of the L_2 error in the velocity at dimensionless $t = Wi$, for all formulations. In all cases $s = 1/100$.

formulation (LC). The right panel of Fig. 5 shows the variation of the L_2 error in velocity at $t = 10$ with resolution s for all three formulations.

LC, CH and DI all converge with approximately 5th order. The DI approach is more accurate by several orders of magnitude. The initial oscillation in the error seen in the left panel of Fig. 5 results from the interplay between the elastic stresses and small acoustic waves generated at start-up, with the size of these small amplitude oscillations decreasing with decreasing Ma . The larger error magnitude for DI here than in the previous case with $\beta = 1$ arises because, although the steady solution is quadratic in y , the transient solution is not. As a result, the advective errors described in the previous subsection for LC and CH also occur for DI in the present case where $\beta \neq 1$.

As indicated above, at lower Ma the solution for all cases becomes less stable, with a gradual increase in error at late times, which appears to be very small (machine precision order) errors acting multiplicatively every time step. Indeed, for the resolution $s = 1/200$, the time-step is $\delta t = 7.5 \times 10^{-6}$ (dimensionless units), so a simulation up to $t = 10$ requires a significant number of time-steps (more than 10^6). Evidently, larger Ma values introduce a degree of acoustic dissipation that limits such error growth and helps to stabilise the simulation. In this regard setting an upper limit of $Ma = 0.05$ for simulations provides a good compromise between maintaining numerical stability and providing a good approximation to flow incompressibility.

4.2.3. $Re = 1$, $\beta = 0.1$, $\varepsilon = 0$, $Ma = 0.05$, $s = 1/100$, Varying Wi

In this subsection we vary the Weissenberg number, Wi , to determine the accuracy of the solution at different levels of elasticity and to explore the largest allowable Wi for a stable simulation for this test case. The left panel of Fig. 6 shows the time evolution of the mean velocity for a range of values of Wi (dashed black lines) compared with the unsteady analytical solution (red lines). Results shown here were obtained with LC, but the results for DI and CH are indistinguishable on these axes. As can be seen, there is an excellent match for $t < 100$ up to $Wi = 128$. At this resolution, the CH and DI approaches break down above $Wi = 128$, whilst LC is stable at $Wi = 256$. Beyond these Wi values, at this resolution, the three schemes fail. All three formulations are capable of reaching higher Wi if the resolution is increased (s reduced) further as key terms in the governing equations are better resolved.

In particular, the errors in evaluating advection terms are larger with increasing Wi , resulting in deviation of c_{yy} from unity for LC and CH (and CH-L), with this deviation increasing with Wi and decreasing with larger s . Typically errors of order 10^{-2} when $Wi = 16$ and $s = 1/100$ are seen in velocity and stress profiles for LC and CH/CH-L approaches, with DI errors smaller by orders of magnitude (but

still growing with increasing Wi). This behaviour can be seen in the right panel of Fig. 6 which shows the L_2 error in the velocity at (dimensionless) $t = Wi$ for a range of Wi , for all three formulations. The orders of error growth with Wi are indicated by the dashed lines, and whilst DI has lower overall error, the rate of error growth with increasing Wi is larger than the LC and CH formulations.

The difference in growth of errors with Wi between DI and LC and CH can be attributed to differences in the solution profiles across the channel. As discussed earlier, whilst in the steady state the cross-channel profiles of c_{xx} and c_{xy} are quadratic and linear (respectively) in y , the log- and Cholesky-transformations of c result in the components of Ψ and L having profiles with more complex structure. In particular, whilst the profiles of the components of Ψ and L are nearly linear and quadratic near the channel walls, they have greater curvature in the channel centre (where the stress is zero), and this curvature increases for increasing Wi . As such, at larger Wi , the dominant errors for LC and CH occur near the channel centre, whilst for DI they are more uniform across the domain (with the only variation being near the walls, where m is reduced towards $m = 4$). Finally, we note that the growth of errors with Wi is very similar for both CH and CH-L formulations, as it is the non-exact (but still $\mathcal{O}(s^m)$) advection of the non-linear profiles near zero-stress points which dominates in this case.

4.3. Periodic array of cylinders

This flow case provides a test of the method in a non-trivial geometry and allows us to assess the performance of the different formulations (LC, CH and DI) for non-parallel flows. The periodic cylinders case simulated follows that of [26] and is based on [25]. The domain is rectangular with dimension $6R \times 4R$, with a cylinder of radius R located in the centre. At the upper and lower boundaries, and the cylinder surface, no-slip wall boundary conditions are imposed, whilst the domain is periodic in the streamwise direction. The flow is driven by a body force, the magnitude of which is set by a PID controller such that the mean velocity magnitude (averaged over the domain) is unity. We take the cylinder radius R as the characteristic length-scale for non-dimensionalisation.

In all cases we set $Re = 2.4 \times 10^{-2}$, $\beta = 0.59$ and $\varepsilon = 0$, to match those parameters used in [25,26]. For this non-parallel flow case at low Re , the magnitude of the body force required to drive the flow is larger than the previous cases, and a smaller value of Ma is required to ensure density variations remain small. We set $Ma = 10^{-3}$, which results in density variations of less than 0.5%. For this test we discretise the domain with a uniform resolution $s_a = s$ for all nodes a , allowing comparison with the aforementioned previous works.

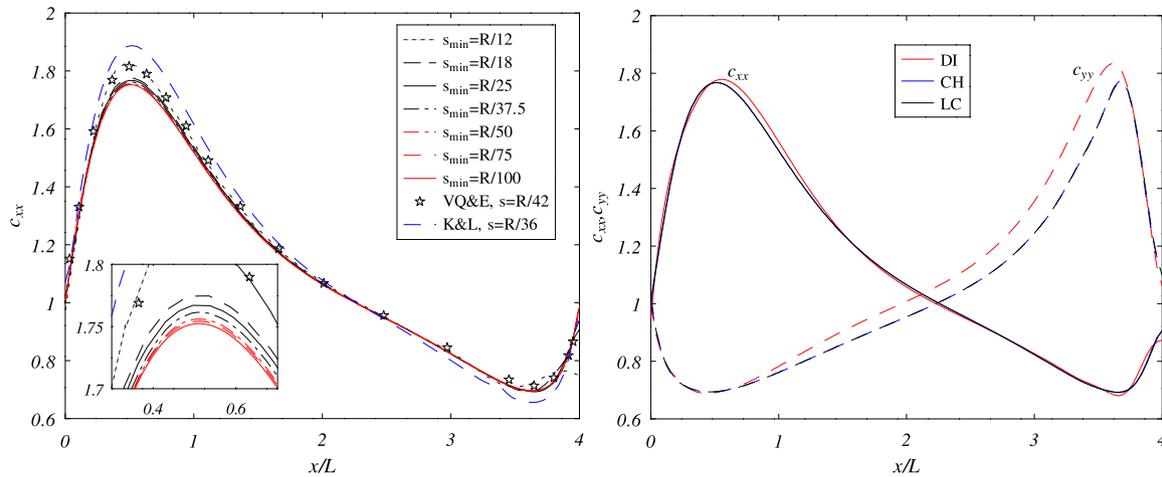


Fig. 7. Profiles of the conformation tensor components along the channel centreline for $Wi = 0.2$ at steady state (20 dimensionless time units). Left panel: c_{xx} for a range of resolutions, using the log-conformation formulation. Results taken from SPH simulations are shown with black stars ([25]) and a dashed blue line ([26]). Right panel: c_{xx} (solid lines) and c_{yy} (dashed lines) for a resolution of $s = R/25$ for the three different formulations.

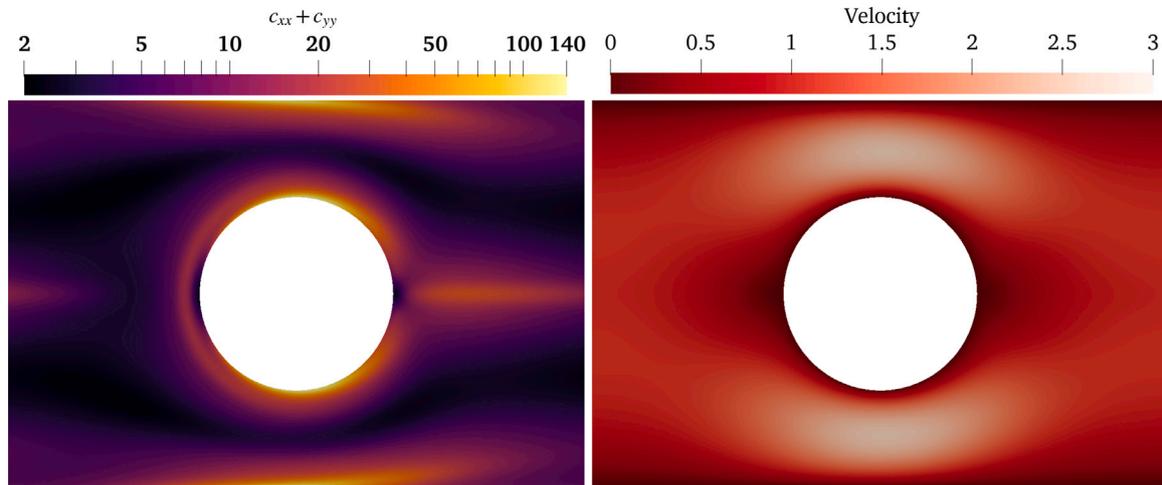


Fig. 8. Isocontours of conformation tensor trace (left) and velocity magnitude (right) for the periodic array of cylinders with $Wi = 0.8$ and $s_{min} = R/50$, obtained using the log-conformation formulation.

We first set $Wi = 0.2$ and assess the accuracy of the method using the LC formulation. The left panel of Fig. 7 shows the profile c_{xx} along the channel centreline for a range of resolutions. We see clearly see convergence in the LABFM solution (inset). The results are compared with SPH data from [25,26], with good agreement shown despite SPH being formally low order. Indeed, both schemes in [25,26] benefit in not having to compute the advection term by their nature of being Lagrangian methods, removing a key term for error growth in the LABFM method. Furthermore, the formulation of [25] is constructed in the GENERIC framework: the symmetries of the conservation laws are matched by the discretised formulation in a thermodynamically consistent way, providing benefits for longer term dynamics and global conservation.

The right panel of Fig. 7 shows the profiles along the channel centreline of c_{xx} and c_{yy} for a fixed resolution of $s = R/25$, for the three formulations DI (red), CH (blue) and LC (black). The LC and CH formulations are almost indistinguishable, but the values from the DI approach (and c_{yy} in particular) deviate slightly.

We next increase the degree of elasticity, setting $Wi = 0.8$. Fig. 8 shows the conformation tensor trace and velocity magnitude fields in this case, with $s_{min} = R/50$, obtained with LC. The problem becomes more numerically challenging now, as the infinite polymer extensibility of the Oldroyd B model results in a singularity in the stress field in the

cylinder wake. It was found by [47] through numerical experiments that for a cylinder in a channel (no periodicity assumed), the solution was divergent for $Wi \geq 0.7$, whilst a similar result was obtained by [48], who obtained exact solutions for the wake centreline stress in the ultra-dilute case. Non-convergence with resolution for $Wi = 0.8$ was also observed in [25]. For the direct integration formulation, a catastrophic instability occurs early in the simulation at all resolutions tested. This is due to errors in the advection of the elements of c , which result in a loss of positive definiteness, leading to non-physical results. At all resolutions studied, CH and LC are in close agreement. In the remainder of this section, the results presented are obtained using the LC formulation. Fig. 9 shows the profiles of velocity (left) and conformation tensor component c_{xx} (right) along the channel centre line for a range of resolutions, alongside results from SPH simulations of [25] (black stars) and [26] (dashed blue lines).

Firstly, it is clear that the profile of the stress in the cylinder wake is diverging with resolution refinement as expected, and we note that the maximum value of c_{xx} in the cylinder wake scales linearly with R/s . Note, that although the stress field is divergent, the velocity field is converging with increasing resolution (inset of the left panel of Fig. 9). Secondly, there are clear discrepancies between the present results and the results of [25,26] using SPH. As described above, there are several contributing factors behind differences observed in Fig. 9 between the

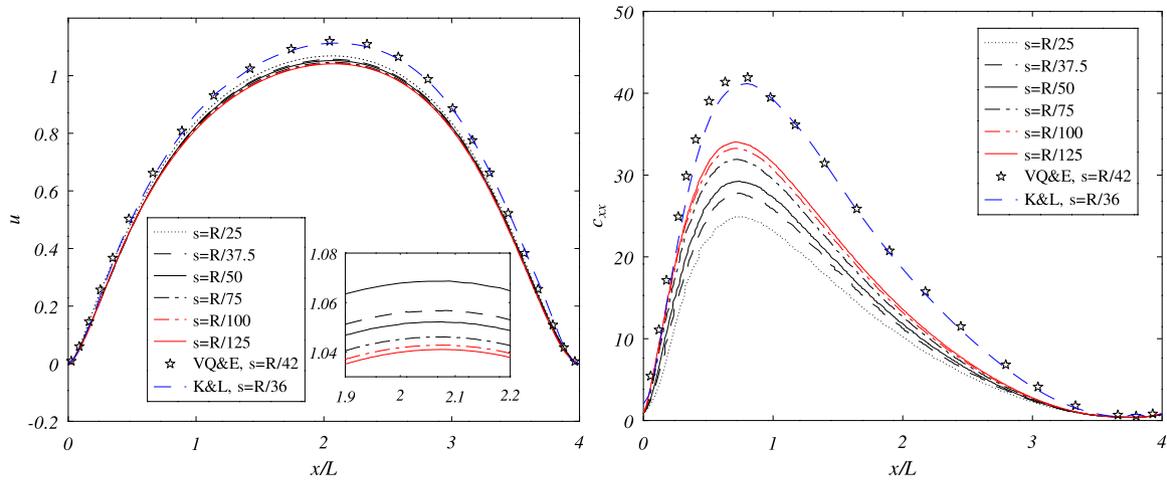


Fig. 9. Profiles along the channel centre-line of the velocity (left panel) and conformation tensor component c_{xx} (right panel) at $Wi = 0.8$, for a range of resolutions (red and black lines), compared with the SPH results of [25] (black stars) and [26] (dashed blue lines).

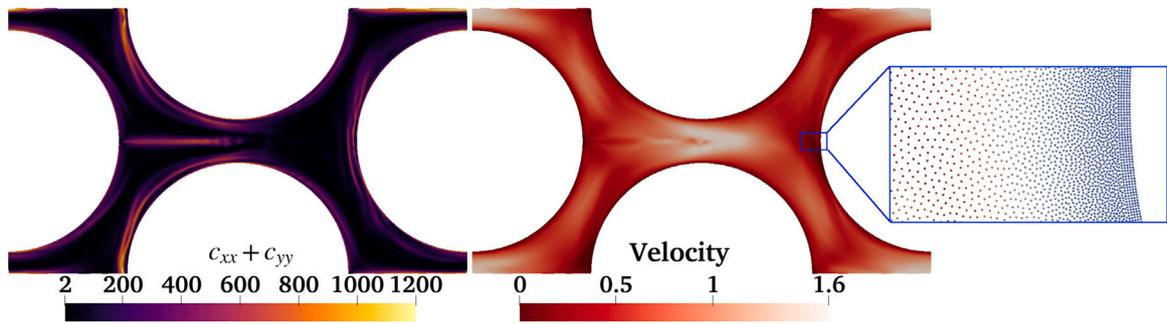


Fig. 10. Isocontours of conformation tensor trace (left) and velocity magnitude (right) for the porous geometry with $Wi = 1$ and $s_{min} = D/600$, simulated using LC. The inset on the right shows the node distribution near one of the cylinders.

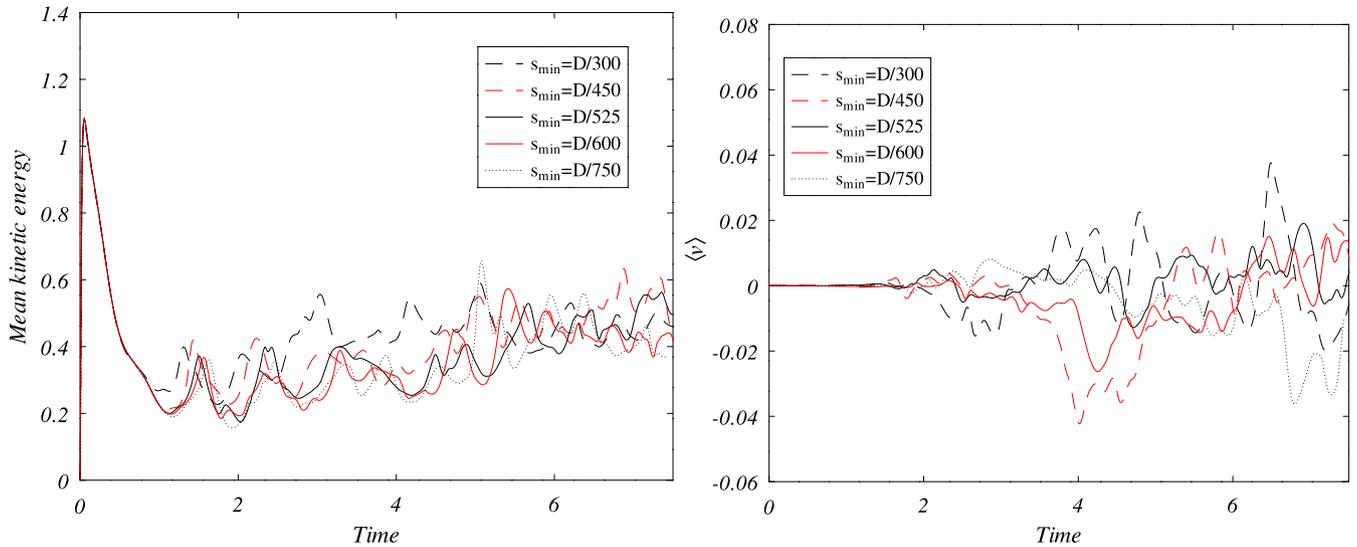


Fig. 11. The time evolution of the volume averaged value kinetic energy (left panel) and transverse velocity $\langle v \rangle$ (right panel), for a range of resolutions.

SPH simulation results of [25,26] and the present work - not least of which being that SPH is formally low order, with such differences in method accuracy exacerbated in a parameter regime with a divergent stress field incorporating steep gradients. However, the exact cause of the discrepancy is not clear, and we also note that in the cited SPH results, the method is Lagrangian, and the non-linear advection

terms are implicitly included in the temporal evolution of the particle positions.

4.4. Representative porous geometry

We next consider a repeating unit of a representative porous geometry, consisting of cylinders with diameter D and spacing S . Several

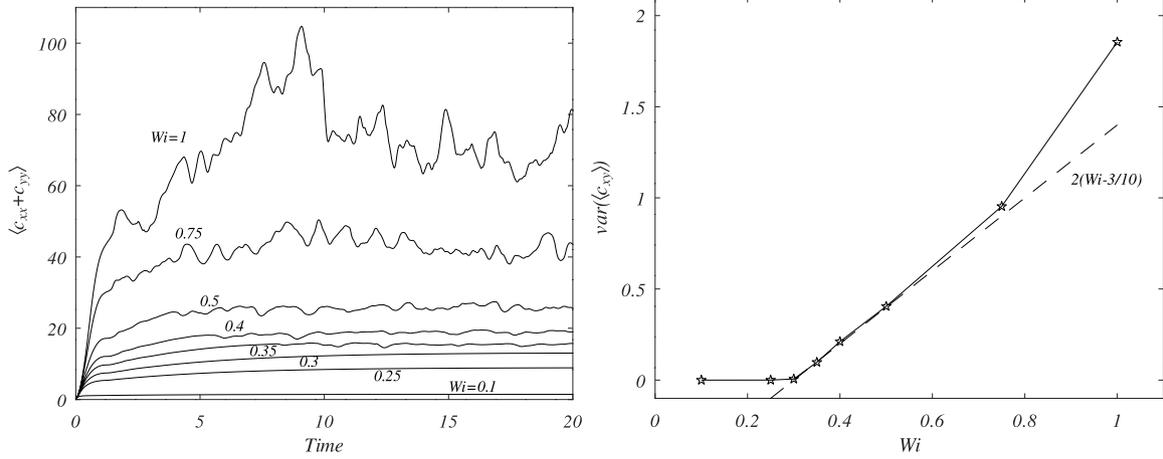


Fig. 12. Left panel: The time evolution of the volume averaged value of the conformation tensor trace $c_{xx} + c_{yy}$. Right panel: The variation with Wi of the variance of the volume averaged value of c_{xy} . The dashed line illustrates a linear dependence on Wi .

authors have studied similar configurations, with Lattice-Boltzmann methods [49,50] and finite volume methods [51,52]. All these works study the flow at negligible Re , whilst we use $Re = 1$. Whilst there are differences in the exact geometries between these works, they all contain the same essential features, allowing for qualitative comparisons to be made. The computational domain has size $\sqrt{3}S \times S$, and is periodic in both directions. A cylinder of diameter $D = S/1.2$ is centred on the midpoint of each boundary. The domain therefore represents a minimal repeating unit of a hexagonal lattice of cylinders. The geometry can be seen in Fig. 10. The system is non-dimensionalised by the cylinder diameter D and the mean velocity magnitude U . The flow is driven by a body force in the x direction, which is set by a PID controller to track $U = 1$. The domain is discretised with a non-uniform resolution s_{min} at the cylinder surfaces, and s increases smoothly away from the cylinders to $s_{max} = 3s_{min}$ at distances greater than $25s_{max}$ from the cylinders. The node distribution near the cylinder is shown in the inset on the right of Fig. 10. With the finest flow structures located near the cylinder walls, the accuracy of the simulations is largely controlled by s_{min} , which we use to characterise the resolution of each simulation.

In all cases, we set $Re = 1$, $\beta = 0.5$, $\epsilon = 10^{-3}$, $Ma = 10^{-2}$. We vary Wi and the resolution. The inclusion of non-zero ϵ (thus representing a PTT fluid rather than an Oldroyd B fluid) avoids the singularity present in the previous test case. We note that the values of Wi studied are small, and based on U and D . An effective Weissenberg number Wi_{eff} for the flow within the pore space may be a more pertinent measure, and could be defined based on the pore size $S - D = D/5$, giving $Wi_{eff} = 5Wi$.

4.4.1. Effect of formulation and resolution at fixed $Wi = 1$

In the first instance we run the simulation for all three formulations. For DI, the simulation quickly becomes unstable, as the thin regions where c_{yy} are large just upstream of each cylinder (see left panel of Fig. 10) cannot be accurately advected. Local oscillations occur, resulting in negative values of c_{yy} , and loss of positive definiteness of the conformation tensor. For CH, the simulation exhibits numerical artefacts at lower resolutions than LC. It appears that at higher resolutions, CH is capable of handling this problem, but LC can achieve accurate solutions at lower resolution, and hence lower cost.

With LC identified as the best formulation for this problem, we focus in more detail on the effect of resolution. In all cases hereafter, we use the LC formulation. Fig. 11 shows the time variation of volume averaged kinetic energy and transverse velocity for resolutions $s_{min}/D \in [1/300, 1/450, 1/525, 1/600, 1/750]$, when using the log-conformation formulation. For resolutions finer than $s_{min} = D/525$ the kinetic energy is approximately converged, and the global drag in the system (as measured by the body force required to drive the flow) is converged

to within 2.4%. The entire system has a chaotic/sensitive dependence on initial conditions, and hence we do not see convergence in the exact trajectory of these global statistics. This behaviour is especially obvious in the right panel of Fig. 11 showing the symmetry breaking given the considerable variation in the volume averaged transverse velocity, $\langle v \rangle$.

A particular computational challenge is that for low Re , we require very small time-steps due to the viscous time-step constraint, with the finest resolution $s_{min} = D/750$ requiring $> 10^7$ time steps to simulate 8 dimensionless time units. Conversely, at higher Re we need exceptionally fine resolution to stably resolve the steep stress gradients and transients leading to the onset of elastic instability. Indeed, whilst high-order discretisations are invaluable for this problem, there is significant benefit to be had from variable and potentially adaptive resolution (in addition to high-order interpolants) for simulations of elastic instabilities. A fully implicit method utilising the present high-order interpolants and discretisation scheme would permit larger time-steps, and enable these simulations at reduced costs. Such an approach is an avenue we are interested in pursuing for future work.

4.4.2. Symmetry breaking with increasing Wi

As a precursor to the complete study and direct numerical simulation of elastic instability in this complex geometry, we consider in more detail the case of symmetry breaking in the flow with increasing Wi up to $Wi = 1$. Note that whilst we show and quantify symmetry breaking, this is a preliminary study and our main focus is on the numerical method. All results in this section are obtained using the LC approach with $s_{min} = D/600$. The Weissenberg numbers considered are $Wi \in [0, 0.1, 0.25, 0.3, 0.35, 0.4, 0.5, 0.75, 1]$. Beyond these values (approximately $Wi = 1.5$) we expect transition to three-dimensional flow, as reported in [53] (for example), and hence extension to 3D simulations remains an area for future work.

We define the instantaneous volume averaged conformation tensor elements as

$$\langle c_{ij}(t) \rangle = \int_V c_{ij}(t) dV, \quad (27)$$

which corresponds to the volume averaged trace if $j = i$, and the volume average of c_{xy} if $j \neq i$. The left panel of Fig. 12 shows the time evolution of the volume averaged conformation tensor trace. As expected, fluctuations of increasing magnitude are seen with increasing Wi , but with values of the volume averaged conformation tensor trace levelling out (on average) with time, indicative of a statistically steady state.

We evaluate the variance of $\langle c_{xy} \rangle$ once the statistically steady state has been reached, over the interval $t \in [10, 20]$. This is a proxy for the measure of asymmetry in the polymeric deformation field. The right

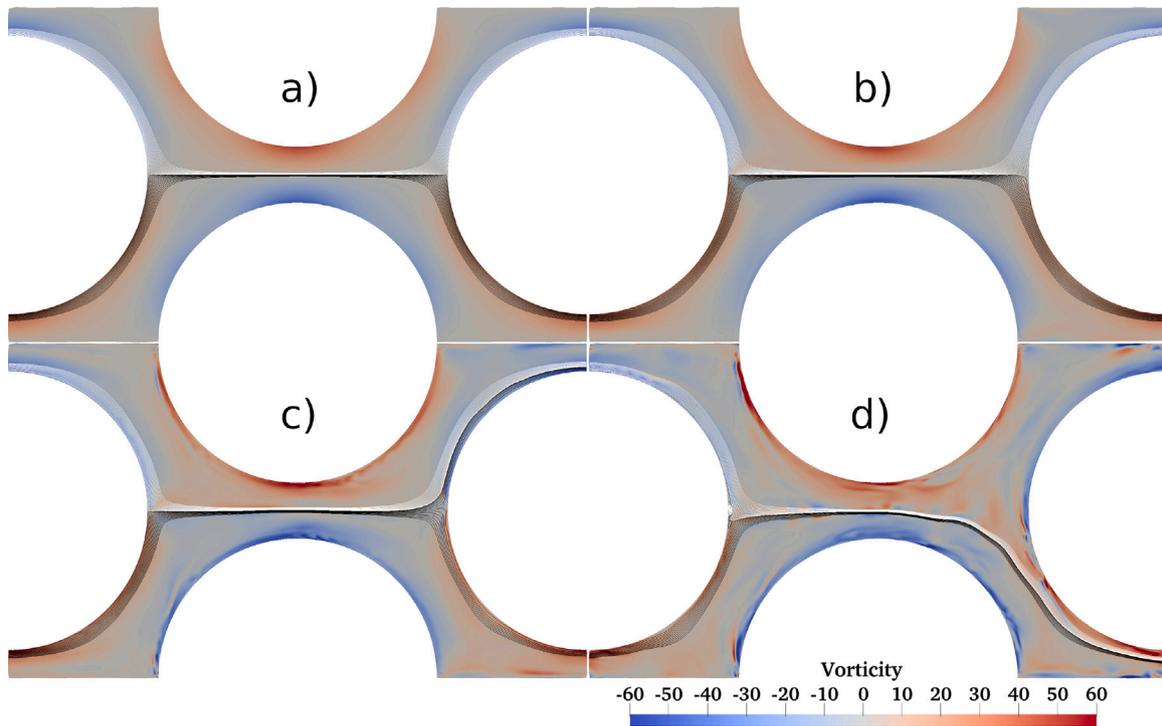


Fig. 13. Isocontours of vorticity (red-blue), with streamlines superimposed, showing the symmetry breaking with increasing Wi . Streamlines originating in the upper half of the domain are coloured white, and those originating in the lower half are coloured black. Panels: (a) $Wi = 0.0$, (b) $Wi = 0.25$, (c) $Wi = 0.5$, and (d) $Wi = 1$.

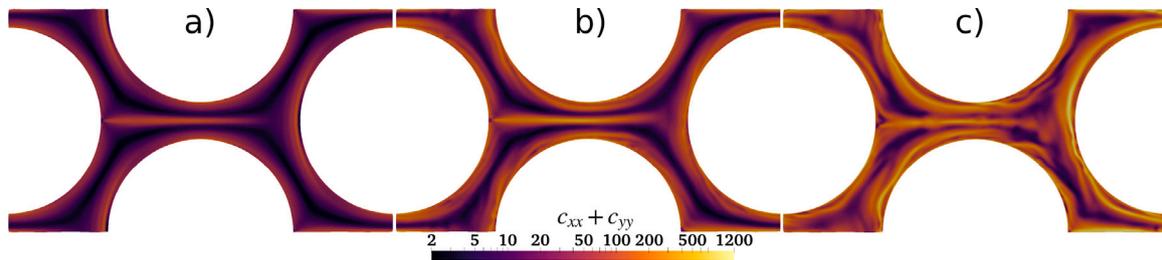


Fig. 14. Isocontours of conformation tensor trace for (a) $Wi = 0.25$, (b) $Wi = 0.5$, and (c) $Wi = 1$.

panel of Fig. 12 shows the dependence of $var \langle c_{xy} \rangle$ on Wi . We see at small Wi where the flow is steady, the variation is negligible. At $Wi \geq 0.3$ the flow is unsteady and symmetry is broken, with the extent of the symmetry breaking having a linear dependence on Wi (dashed line) with slope 2. Note that for $Wi > 0.75$ this relation ceases, likely as the flow enters a different, more elastic, regime.

Fig. 13 shows isocontours of the vorticity field (red-blue) with streamlines showing flow crossing from the upper to lower halves of the domain at increasing Wi . Note that between $Wi = 0$ and $Wi = 0.25$ (panels (a) and (b) respectively) the vorticity field develops a streamwise asymmetry as expected, and observed in a similar configuration by [51]. By $Wi = 0.5$ the instability has developed and the symmetry is broken in the transverse direction, as shown by streamlines crossing the domain centreline. Similarly, Fig. 14 shows isocontours of the conformation tensor trace $c_{xx} + c_{yy}$ for (a) $Wi = 0.25$, (b) $Wi = 0.5$, and (c) $Wi = 1$. As above, by $Wi = 0.25$ the field has developed a streamwise asymmetry, which clearly breaks in the transverse direction by $Wi = 0.5$. Beyond this ($Wi = 1$), symmetry breaking in the flow is clear with unsteady elastic flow structures larger in magnitude - a precursor stage before the flow evolves fully 3D flow structures. The unsteady thin near-wall structures are qualitatively similar to those found in [49], who studied a similar geometry but in the creeping flow regime.

5. Conclusions

In this work a new high-order meshless method for the solution of viscoelastic flow in two-dimensional, non-trivial geometries has been presented. Three different approaches to treating the viscoelastic stresses are considered for assessment in this new high-order meshless framework — direct integration, Cholesky decomposition (both with and without a logarithmic transform of the diagonal elements), and the log-conformation formulation. Direct integration provides notably more accurate solutions for parallel flows but the log-conformation approach provides enhanced stability across all test cases considered. Highly accurate results can be obtained with convergence up to 9th order, depending on the test case. For parallel flows, the attainable Weissenberg numbers are large, up to $Wi = 128$. For non-trivial geometries, the attainable Weissenberg numbers are more moderate, up to $\mathcal{O}(1)$, but we find that the limiting factor is the requirement to resolve the increasingly fine flow features present with increasing Wi , suggesting that our method can handle higher Wi given sufficient resolutions. The meshless nature of the method enables non-trivial geometries to be discretised straightforwardly, with variable resolution easily included. Accordingly, an initial study of a symmetry breaking elastic instability at moderate Wi is considered in a non-trivial representative porous media geometry. The results are promising and demonstrate the potential of this method for the high-fidelity study of fully 3-D elastic instabilities

in realistic, industrially relevant geometries in the longer term. The explicit nature of the present formulation renders simulations in the limit of vanishing Re impractical, and thus the method is well suited to inertial flows. An implicit formulation would allow us to simulate flows with negligible inertia, and explore purely elastic instabilities in complex geometries. These are the main goals of our future work, with any 3-D method also requiring adaptivity of resolution, both in polynomial reconstruction and spatial resolution, to enable the capture of thin, unsteady elastic flow structures in a computationally efficient manner.

CRedit authorship contribution statement

J.R.C. King: Writing – review & editing, Writing – original draft, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization. **S.J. Lind:** Writing – review & editing, Writing – original draft, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

JK is funded by the Royal Society via a University Research Fellowship, UK (UR/R1\221290). We would like to acknowledge the assistance given by Research IT and the use of the Computational Shared Facility at the University of Manchester.

References

- A. Beris, R. Armstrong, R. Brown, Finite element calculation of viscoelastic flow in a journal bearing: I. small eccentricities, *J. Non-Newton. Fluid Mech.* 16 (1) (1984) 141–172, [http://dx.doi.org/10.1016/0377-0257\(84\)85008-9](http://dx.doi.org/10.1016/0377-0257(84)85008-9).
- D. Rajagopalan, R.C. Armstrong, R.A. Brown, Finite element methods for calculation of steady, viscoelastic flow using constitutive equations with a Newtonian viscosity, *J. Non-Newton. Fluid Mech.* 36 (1990) 159–192, [http://dx.doi.org/10.1016/0377-0257\(90\)85008-M](http://dx.doi.org/10.1016/0377-0257(90)85008-M), URL <https://www.sciencedirect.com/science/article/pii/037702579085008M>.
- J. Sun, N. Phan-Thien, R.I. Tanner, An adaptive viscoelastic stress splitting scheme and its applications: AVSS/SI and AVSS/SUPG, *J. Non-Newton. Fluid Mech.* 65 (1) (1996) 75–91, [http://dx.doi.org/10.1016/0377-0257\(96\)01448-6](http://dx.doi.org/10.1016/0377-0257(96)01448-6).
- J. Sun, M. Smith, R. Armstrong, R. Brown, Finite element method for viscoelastic flows based on the discrete adaptive viscoelastic stress splitting and the discontinuous Galerkin method: DAVSS-G/DG, *J. Non-Newton. Fluid Mech.* 86 (3) (1999) 281–307, [http://dx.doi.org/10.1016/S0377-0257\(98\)00176-1](http://dx.doi.org/10.1016/S0377-0257(98)00176-1).
- S.-C. Xue, R. Tanner, N. Phan-Thien, Numerical modelling of transient viscoelastic flows, *J. Non-Newton. Fluid Mech.* 123 (1) (2004) 33–58, <http://dx.doi.org/10.1016/j.jnnfm.2004.06.009>.
- X. Chen, H. Marschall, M. Schäfer, D. Bothe, A comparison of stabilisation approaches for finite-volume simulation of viscoelastic fluid flow, *Int. J. Comput. Fluid Dyn.* 27 (6–7) (2013) 229–250, <http://dx.doi.org/10.1080/10618562.2013.829916>.
- A. Morozov, Coherent structures in plane channel flow of dilute polymer solutions with vanishing inertia, *Phys. Rev. Lett.* 129 (2022) 017801, <http://dx.doi.org/10.1103/PhysRevLett.129.017801>.
- J.J.J. Gillissen, Two-dimensional decaying elastoinertial turbulence, *Phys. Rev. Lett.* 123 (2019) 144502, <http://dx.doi.org/10.1103/PhysRevLett.123.144502>.
- R. Fattal, R. Kupferman, Constitutive laws for the matrix-logarithm of the conformation tensor, *J. Non-Newton. Fluid Mech.* 123 (2) (2004) 281–285, <http://dx.doi.org/10.1016/j.jnnfm.2004.08.008>.
- R. Fattal, R. Kupferman, Time-dependent simulation of viscoelastic flows at high Weissenberg number using the log-conformation representation, *J. Non-Newton. Fluid Mech.* 126 (1) (2005) 23–37, <http://dx.doi.org/10.1016/j.jnnfm.2004.12.003>.
- T. Vaithianathan, L.R. Collins, Numerical approach to simulating turbulent flow of a viscoelastic polymer solution, *J. Comput. Phys.* 187 (1) (2003) 1–21, [http://dx.doi.org/10.1016/S0021-9991\(03\)00288-7](http://dx.doi.org/10.1016/S0021-9991(03)00288-7).
- V. Steinberg, Elastic turbulence: An experimental view on inertialess random flow, *Annu. Rev. Fluid Mech.* 53 (1) (2021) 27–58, <http://dx.doi.org/10.1146/annurev-fluid-010719-060129>.
- Y. Dubief, V.E. Terrapon, B. Hof, Elasto-inertial turbulence, *Annu. Rev. Fluid Mech.* 55 (1) (2023) 675–705, <http://dx.doi.org/10.1146/annurev-fluid-032822-025933>.
- R.G. Owens, T.N. Phillips, *Computational rheology*, World Scientific, 2002.
- S. Piliitsis, A.N. Beris, Calculations of steady-state viscoelastic flow in an undulating tube, *J. Non-Newton. Fluid Mech.* 31 (3) (1989) 231–287.
- R.G. Owens, T.N. Phillips, Compatible pseudospectral approximations for incompressible flow in an undulating tube, *J. Rheol.* 37 (6) (1993) 1181–1199.
- S. Momeni-Masuleh, T. Phillips, Viscoelastic flow in an undulating tube using spectral methods, *Comput. & Fluids* 33 (8) (2004) 1075–1095, <http://dx.doi.org/10.1016/j.compfluid.2003.09.002>.
- S. Berti, G. Boffetta, Elastic waves and transition to elastic turbulence in a two-dimensional viscoelastic Kolmogorov flow, *Phys. Rev. E* 82 (2010) 036314, <http://dx.doi.org/10.1103/PhysRevE.82.036314>.
- S. Claus, T. Phillips, Viscoelastic flow around a confined cylinder using spectral/hp element methods, *J. Non-Newton. Fluid Mech.* 200 (2013) 131–146, <http://dx.doi.org/10.1016/j.jnnfm.2013.03.004>, Special Issue: Advances in Numerical Methods for Non-Newtonian Flows.
- R. Kynch, T. Phillips, A high resolution spectral element approximation of viscoelastic flows in axisymmetric geometries using a DEVSS-G/DG formulation, *J. Non-Newton. Fluid Mech.* 240 (2017) 15–33, <http://dx.doi.org/10.1016/j.jnnfm.2016.12.008>.
- H. Garg, E. Calzavarini, S. Berti, Statistical properties of two-dimensional elastic turbulence, *Phys. Rev. E* 104 (2021) 035103, <http://dx.doi.org/10.1103/PhysRevE.104.035103>.
- J. Ramsay, M. Simmons, A. Ingram, E. Stitt, Mixing of Newtonian and viscoelastic fluids using “butterfly” impellers, *Chem. Eng. Sci.* 139 (2016) 125–141, <http://dx.doi.org/10.1016/j.ces.2015.09.026>.
- M. Ellero, M. Kröger, S. Hess, Viscoelastic flows studied by smoothed particle dynamics, *J. Non-Newton. Fluid Mech.* 105 (1) (2002) 35–51, [http://dx.doi.org/10.1016/S0377-0257\(02\)00059-9](http://dx.doi.org/10.1016/S0377-0257(02)00059-9).
- J. Fang, R.G. Owens, L. Tacher, A. Parriaux, A numerical study of the SPH method for simulating transient viscoelastic free surface flows, *J. Non-Newton. Fluid Mech.* 139 (1–2) (2006) 68–84.
- A. Vázquez-Quesada, M. Ellero, SPH simulations of a viscoelastic flow around a periodic array of cylinders confined in a channel, *J. Non-Newton. Fluid Mech.* 167–168 (2012) 1–8, <http://dx.doi.org/10.1016/j.jnnfm.2011.09.002>.
- J. King, S. Lind, High Weissenberg number simulations with incompressible smoothed particle hydrodynamics and the log-conformation formulation, *J. Non-Newton. Fluid Mech.* 293 (2021) 104556, <http://dx.doi.org/10.1016/j.jnnfm.2021.104556>.
- B. ten Bosch, On an extension of dissipative particle dynamics for viscoelastic flow modelling, *J. Non-Newton. Fluid Mech.* 83 (3) (1999) 231–248, [http://dx.doi.org/10.1016/S0377-0257\(98\)00092-5](http://dx.doi.org/10.1016/S0377-0257(98)00092-5).
- N. Phan-Thien, N. Mai-Duy, T. Nguyen, A note on dissipative particle dynamics (DPD) modelling of simple fluids, *Comput. & Fluids* 176 (2018) 97–108, <http://dx.doi.org/10.1016/j.compfluid.2018.08.030>.
- S. Litvinov, M. Ellero, X. Hu, N.A. Adams, Smoothed dissipative particle dynamics model for polymer molecules in suspension, *Phys. Rev. E* 77 (2008) 066703, <http://dx.doi.org/10.1103/PhysRevE.77.066703>.
- N. Moreno, M. Ellero, Arbitrary flow boundary conditions in smoothed dissipative particle dynamics: A generalized virtual rheometer, *Phys. Fluids* 33 (1) (2021).
- D. Nieto Simavilla, M. Ellero, Mesoscopic simulations of inertial drag enhancement and polymer migration in viscoelastic solutions flowing around a confined array of cylinders, *J. Non-Newton. Fluid Mech.* 305 (2022) 104811, <http://dx.doi.org/10.1016/j.jnnfm.2022.104811>.
- N.J. Quinlan, M. Basa, M. Lastiwka, Truncation error in mesh-free particle methods, *Internat. J. Numer. Methods Engrg.* 66 (13) (2006) 2064–2085, <http://dx.doi.org/10.1002/nme.1617>.
- S. Lind, P. Stansby, High-order Eulerian incompressible smoothed particle hydrodynamics with transition to Lagrangian free-surface motion, *J. Comput. Phys.* 326 (2016) 290–311, <http://dx.doi.org/10.1016/j.jcp.2016.08.047>.
- A.M.A. Nasar, G. Fourtakas, S.J. Lind, B.D. Rogers, P.K. Stansby, J.R.C. King, High-order velocity and pressure wall boundary conditions in Eulerian incompressible SPH, *J. Comput. Phys.* 434 (2021) 109793, <http://dx.doi.org/10.1016/j.jcp.2020.109793>.
- J. King, S. Lind, A. Nasar, High order difference schemes using the local anisotropic basis function method, *J. Comput. Phys.* 415 (2020) 109549, <http://dx.doi.org/10.1016/j.jcp.2020.109549>.
- J. King, S. Lind, High-order simulations of isothermal flows using the local anisotropic basis function method (LABFM), *J. Comput. Phys.* 449 (2022) 110760, <http://dx.doi.org/10.1016/j.jcp.2021.110760>.
- J. King, A mesh-free framework for high-order direct numerical simulations of combustion in complex geometries, *Comput. Methods Appl. Mech. Engrg.* 421 (2024) 116762, <http://dx.doi.org/10.1016/j.cma.2024.116762>.

- [38] S.J. Lind, A numerical study of the effect of viscoelasticity on cavitation and bubble dynamics (Ph.D. thesis), Cardiff University, 2010.
- [39] A.T. Mackay, T.N. Phillips, On the derivation of macroscopic models for compressible viscoelastic fluids using the generalized bracket framework, *J. Non-Newton. Fluid Mech.* 266 (2019) 59–71.
- [40] B. Fornberg, N. Flyer, Fast generation of 2-D node distributions for mesh-free PDE discretizations, *Comput. Math. Appl.* 69 (7) (2015) 531–544, <http://dx.doi.org/10.1016/j.camwa.2015.01.009>.
- [41] W. Dehnen, H. Aly, Improving convergence in smoothed particle hydrodynamics simulations without pairing instability, *Mon. Not. R. Astron. Soc.* 425 (2) (2012) 1068–1082, <http://dx.doi.org/10.1111/j.1365-2966.2012.21439.x>.
- [42] C.A. Kennedy, M.H. Carpenter, R. Lewis, Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations, *Appl. Numer. Math.* 35 (3) (2000) 177–219, [http://dx.doi.org/10.1016/S0168-9274\(99\)00141-5](http://dx.doi.org/10.1016/S0168-9274(99)00141-5).
- [43] J.R.C. King, S.J. Lind, B.D. Rogers, P.K. Stansby, R. Vacondio, Large Eddy simulations of bubbly flows and breaking waves with smoothed particle hydrodynamics, *J. Fluid Mech.* 972 (2023) A24, <http://dx.doi.org/10.1017/jfm.2023.649>.
- [44] X. Guo, B.D. Rogers, S. Lind, P.K. Stansby, New massively parallel scheme for incompressible smoothed particle hydrodynamics (ISPH) for highly nonlinear and distorted flow, *Comput. Phys. Comm.* 233 (2018) 16–28, <http://dx.doi.org/10.1016/j.cpc.2018.06.006>.
- [45] J. O'Connor, J.M. Domínguez, B.D. Rogers, S.J. Lind, P.K. Stansby, Eulerian incompressible smoothed particle hydrodynamics on multiple GPUs, *Comput. Phys. Comm.* 273 (2022) 108263, <http://dx.doi.org/10.1016/j.cpc.2021.108263>.
- [46] J.C. Sutherland, C.A. Kennedy, Improved boundary conditions for viscous, reacting, compressible flows, *J. Comput. Phys.* 191 (2) (2003) 502–524, [http://dx.doi.org/10.1016/S0021-9991\(03\)00328-0](http://dx.doi.org/10.1016/S0021-9991(03)00328-0).
- [47] M. Alves, F. Pinho, P. Oliveira, The flow of viscoelastic fluids past a cylinder: finite-volume high-resolution methods, *J. Non-Newton. Fluid Mech.* 97 (2) (2001) 207–232, [http://dx.doi.org/10.1016/S0377-0257\(00\)00198-1](http://dx.doi.org/10.1016/S0377-0257(00)00198-1).
- [48] M. Bajaj, M. Pasquali, J.R. Prakash, Coil-stretch transition and the breakdown of computations for viscoelastic fluid flow around a confined cylinder, *J. Rheol.* 52 (1) (2008) 197–223, <http://dx.doi.org/10.1122/1.2807444>.
- [49] V. Dzanic, C.S. From, A. Gupta, C. Xie, E. Sauret, Geometry dependence of viscoelastic instabilities through porous media, *Phys. Fluids* 35 (2) (2023) 023105, <http://dx.doi.org/10.1063/5.0138184>.
- [50] J.J.J. Gillissen, Viscoelastic flow simulations through an array of cylinders, *Phys. Rev. E* 87 (2013) 023003, <http://dx.doi.org/10.1103/PhysRevE.87.023003>.
- [51] S. De, J.A.M. Kuipers, E.A.J.F. Peters, J.T. Padding, Viscoelastic flow simulations in model porous media, *Phys. Rev. Fluids* 2 (2017) 053303, <http://dx.doi.org/10.1103/PhysRevFluids.2.053303>.
- [52] M. Kumar, S. Aramideh, C.A. Browne, S.S. Datta, A.M. Ardekani, Numerical investigation of multistability in the unstable flow of a polymer solution through porous media, *Phys. Rev. Fluids* 6 (2021) 033304, <http://dx.doi.org/10.1103/PhysRevFluids.6.033304>.
- [53] M. Grilli, A. Vázquez-Quesada, M. Ellero, Transition to turbulence and mixing in a viscoelastic fluid flowing inside a channel with a periodic array of cylindrical obstacles, *Phys. Rev. Lett.* 110 (2013) 174501, <http://dx.doi.org/10.1103/PhysRevLett.110.174501>.