Journal of Computational Physics 228 (2009) 4606–4624

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

The use of PDE centres in the local RBF Hermitian method for 3D convective-diffusion problems

David Stevens, Henry Power *, Michael Lees, Herve Morvan

School of Mechanical, Materials and Manufacturing Engineering, University of Nottingham, University Park, Nottingham, Nottinghamshire NG7 2RD, UK

ARTICLE INFO

Article history:
Received 24 October 2008
Received in revised form 18 March 2009
Accepted 20 March 2009
Available online 29 March 2009

Keywords:
Numerical solution
Initial and boundary value problems
Meshless collocation approach
Radial basis function
Local collocation
Implicit time stepping scheme

ABSTRACT

In this work an extension is proposed to the Local Hermitian Interpolation (LHI) method; a meshless numerical method based on interpolation with small and heavily overlapping radial basis function (RBF) systems. This extension to the LHI method uses interpolation functions which themselves satisfy the partial differential equation (PDE) to be solved. In this way, a much improved reconstruction of partial derivatives can be obtained, resulting in significantly improved accuracy in many cases.

The implementation algorithm is described, and is validated via three convection-diffusion-reaction problems, for steady and transient situations. A Crank–Nicolson implicit time stepping technique is used for the time-dependent problems. In the proposed approach, a form of 'analytical upwinding' is implicitly implemented by the use of the partial differential operator of the governing equation in the interpolation function, which includes the desired information about the convective velocity field.

The implicit upwinding scheme intrinsic to the proposed numerical approach is tested by solving a one-dimensional travelling front problem at Péclet numbers of 500, 1000, 2000, 5000 and infinity, which corresponds to a shock front in the case of infinity. In addition, the accuracy of the numerical scheme is validated against a one-dimensional steady state solution exhibiting strong boundary layer effects, and also against a steady and a transient three-dimensional convection-diffusion problem on irregular datasets. All the test cases are validated against the corresponding analytical solutions. Finally, the effect of various interpolation stencil configurations is investigated, and some important limitations on local data-centre distribution are identified.

1. Introduction

Radial Basis Functions (RBFs) have traditionally been used to provide a continuous interpolation of scattered data sets. However, this interpolation also allows the reconstruction of the partial derivatives throughout the solution field, which can then be used to drive the solution of a partial differential equation (PDE). Since the interpolation takes place on a scattered dataset with no local connectivity, the solution is essentially meshless. RBF-based methods have been successfully used to solve a wide variety of PDEs in this fashion.

Although full-domain RBF methods are highly flexible and exhibit high order convergence rates [1], in their basic implementation (as described in Section 2) the fully-populated matrix systems produced lead to poor numerical conditioning as the size of the dataset increases. This problem is described by Schaback [2] as the "uncertainty relation"; better conditioning

* Corresponding author. Tel.: +44 115 8466232; fax: +44 115 9513800.
E-mail address: henry.power@nottingham.ac.uk (H. Power).

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.03.025

is associated with worse accuracy, and worse conditioning is associated with improved accuracy. With increasingly large datasets and increasingly flat basis functions, this problem becomes more pronounced.

It is important to note that the uncertainty relation is a consequence of ill-conditioning in the determination of RBF weighting coefficients, rather than an ill-conditioning of the RBF problem itself [3]. Recent works by Fornberg et al. [4,5] have shown how it is possible to extract data from RBF interpolations without the need to compute the ill conditioned RBF weighting coefficients, and in doing so achieve significantly improved accuracy through the use of increasingly flat basis functions. However, in their current form these formulations are relatively inflexible towards application on arbitrary geometries, and extraction of partial derivatives. Further developments in this field may eventually offer a powerful new approach to RBF interpolation for large-scale engineering applications.

Many techniques have also been developed to reduce the effect of the uncertainty relation with the traditional RBF interpolation formulation, such as RBF-specific preconditioners [6], and adaptive selection of data-centres [7]. However, at present the only reliable method of controlling numerical ill-conditioning and particularly computational cost, as problem size increases, is through domain decomposition (see, for example, [8–11]).

By taking the domain decomposition principle and applying it to very small and heavily overlapping local systems, a locally-supported RBF collocation method can be formulated. Since the individual RBF systems never grow too large, such methods can be scaled to arbitrarily large datasets without numerical conditioning issues, and with a computational cost proportional to the number of local systems used (order N complexity). Such a formulation retains the flexibility of working on a set of scattered data at the expense of introducing some localised connectivity.

The idea of using a local RBF collocation approach to solve boundary value problems can be seen as a generalization of the compact finite difference scheme in terms of scattered nodes (for more detail see Wright and Fornberg [12]). This local approach was introduced in [13–15] to solve two-dimensional Poisson, elasticity and Newtonian incompressible viscous flow problems respectively. Recent works by Sarler and Vertnik [16], Vertnik et al. [17], Kosec and Sarler [18], and Divo and Kassab [19] have shown how the local RBF collocation technique can be applied to solve transient problems. In those works several two-dimensional diffusion, metal casting and non-isothermal flow problems are solved. The approach used in [16–18] is based on an explicit time stepping scheme and the radial basis functions are employed to directly interpolate the solution at the previous time step, in order to obtain the corresponding spatial derivatives of the non-homogeneous terms of the first order ordinary equations in time, which result from the explicit formulation. On the other hand, the work in [19] is based on an implicit time stepping scheme and the Kansa (unsymmetric) RBF meshless method (see Section 2 for more details) is used at each time step to solve the pseudo boundary value problem at the local stencils. In this way, by using the obtained local approximations the global solution is found by an iterative algorithm, without constructing the corresponding global system of equations.

In [20] an explicit local RBF formulation for three-dimensional problems is described which is capable of operating on scattered datasets with arbitrary boundary conditions, and is applied to a variety of convection-diffusion problems. In [21] this method was applied to the non-linear Richards equation, modelling flow in unsaturated porous media. The major fundamental difference between the technique described in [20,21] (referred to henceforth as the Local Hermitian Interpolation, or LHI, method) and those presented in [16,19] is the inclusion of the boundary operator within the basis functions, allowing a natural description of the boundary conditions to be obtained. This formulation is further extended here, applying the PDE-operator to the basis functions in analogy to the full-domain Hermitian RBF method. The local interpolation systems resulting from such inclusion do themselves satisfy the governing PDE at some locations, and as such the partial derivatives extracted from such systems more closely represent their analytical values.

In contrast with the approach suggested in [19], in our approach the Hermitian (Symmetric) meshless method (see Section 2 for more details) is used to solve the pseudo boundary value problem at the local stencils, instead of using the Kansa method. Having obtained an approximation for the local unknown field variable in terms of its value at neighbouring data-centres, the governing equation is required to be satisfied at each local system centrepoint, in order to obtain a global system of equations. In this way, the governing equation is imposed simultaneously at the local and global levels.

The method for inclusion of PDE-operators within the LHI formulation is described in this work. The formulation of modified PDE-operators for transient problems is described using a weighted Crank–Nicolson implicit time stepping technique, which can be directly extended to other higher-order time stepping techniques. It is important to point out that having an interpolation which satisfies locally the partial differential operator at neighbouring data-centres, including the desired information about the convective velocity field, provides an implicit upwind formulation.

It is known that spurious oscillations (instabilities) in the numerical results of convective dominant problems, in the presence of discontinuities, are frequently observed when using most classical numerical techniques based on centrally defined interpolation functions. Such instabilities are due to dispersive errors. In these central schemes, each collocation point is enclosed by an interpolation stencil including points from the upstream and downstream direction of the convective velocity field, with similar weighting functions between interpolating points. To control those spurious oscillations it is necessary to utilise upwinding interpolation stencils, defined by upstream points or heavily weighted on those points. Unfortunately, the definition of upwind stencils is completely ad hoc and their implementation for scattered interpolation points when dealing with complex three-dimensional problems is not a trivial task. Besides, the use of upwinding schemes increases diffusion errors due to a fall of their order of accuracy to first order in regions of steep gradients or discontinuities.

In our approach, using interpolation functions defined in terms of a local approximation of the governing equation, is not necessary to recourse to upwinding schemes in order to obtain the solution of convective dominant problems without

spurious oscillations. Each local system is centrally defined, without the need to use interpolation stencils which are adjusted according to the magnitude and direction of the local convective flow. The required information about the convective field is analytically imposed on the local approximation of the governing equation. As in the case of the upwinding schemes our solution exhibits some diffusive errors that can be mitigated by using higher-order time integration algorithms.

The effect of the implicit upwind scheme, as introduced by the use of the PDE-operator with the local RBF interpolation, is investigated in the numerical simulation of a travelling front problem, at values of the Péclet number of 500, 1000, 2000, 5000 and infinity, which corresponds to a shock front in the case of infinity. Other one and three-dimensional benchmark convection-diffusion problems are also used to test the accuracy of the numerical method, with all problems solved in 3D domains. Finally, the effect of using different stencil configurations in the local interpolation is also investigated.

The numerical examples considered are simple but not trivial problems, in particular those dealing with the evolution of steep fronts. The development of robust numerical algorithms for this type of problem is the subject of ongoing research. Different ways of improving upwinding schemes in order to minimise both dispersive and diffusive errors have been recently proposed; see [22,23]. In the present work, we show that with our approach it is not necessary to use any upwinding scheme.

2. Basic RBF formulation

A radial basis function $\Psi(\|x-\xi_j\|)$ depends upon the separation distances of a set of functional centres, or trial points $\{\xi_j \in \Re^n;\ j=1,2,\ldots,N\}$, and exhibits spherical symmetry around these centres. The most commonly used RBFs are given in Table 1.

The Gaussian and the inverse multiquadric ($m<0$ in the generalized multiquadric function), are positive definite functions. The thin-plate spline and the multiquadric ($m>0$) are conditionally positive definite functions of order m, which require the addition of a polynomial term of order $m-1$, together with a homogeneous constraint condition, in order to obtain an invertible interpolation matrix. For the multiquadric with $m>0$, only odd values of m are permitted [3]. The 'c' term is known as a 'shape-parameter', and describes the relative width of the RBF functions about their centres. In practice, tuning of this parameter can dramatically affect the quality of the solution obtained. Increasing the value of c will lead to a flatter RBF. This will, in general, improve the rate of convergence at the expense of increased numerical ill-conditioning of the resulting linear system [24].

The meshless RBF method for solving PDEs, as described by Kansa [25,26], constructs the (continuous) solution $u(x)$ of the PDE from a distinct set of N quasi-randomly distributed functional centres ξ_j

$$u(x)=\sum_{j=1}^{N}\lambda_j\Psi(\|x-\xi_j\|)+\sum_{j=1}^{NP}\lambda_{j+N}P^j_{m-1}(x)\quad x\in\Re^n \tag{1}$$

Here P^j_{m-1} is the jth term of an order $(m-1)$ polynomial, under the constraint

$$\sum_{j=1}^{N}\lambda_j P^k_{m-1}(x_j)=0\quad k=1,\ldots,NP \tag{2}$$

with NP being the total number of terms in the polynomial (determined by the polynomial order and the number of spatial dimensions).

Consider a typical boundary value problem

$$\begin{aligned} L[u]&=f(x)\quad \text{on } \Omega\\ B[u]&=g(x)\quad \text{on } \partial\Omega \end{aligned} \tag{3}$$

where the operators $L[]$ and $B[]$ are linear partial differential operators on the domain Ω and on the contour $\partial\Omega$, describing the governing equation and boundary conditions respectively. The operator $B[]$ corresponds to the unit operator when Dirichlet conditions are prescribed, the normal derivative for Neumann conditions, and a combination of the two in the case of Robin or mixed conditions.

Collocating the system at N distinct locations known as test points, x_j, coinciding with the trial points ξ_j, leads to a system of equations

$$\begin{bmatrix} B[\Psi] & B[P_{m-1}]\\ L[\Psi] & L[P_{m-1}]\\ P_{m-1} & 0 \end{bmatrix}\lambda=\begin{bmatrix} g\\ f\\ 0 \end{bmatrix} \tag{4}$$

Table 1
Commonly used radial basis functions.

Generalized thin-plate spline	Generalized multiquadric	Gaussian
Radial basis functions (RBFs) $r^{2m-2}\ln(r)$ where m is an integer and $r=\|x-\xi_j\|$	$(r^2+c^2)^{m/2}$	$\exp(-r^2/c^2)$

which is fully populated and non-symmetric. This approach, known as Kansa's method or the unsymmetric method, has been applied to a wide range of problems with great success. However, for the standard formulation there is no formal guarantee that the collocation matrix will be non-singular [7]. In the special case of a numerically singular system, a small perturbation of functional centre locations (or of the shape-parameter) will in general return a non-singular collocation matrix [6]; however, the perturbed collocation matrix may suffer from numerical ill-conditioning issues.

An alternative approach proposed by Fasshauer [27] uses the $L[]$ and $B[]$ PDE-operators applied to the RBF as basis functions

$$u(x) = \sum_{j=1}^{NB} \lambda_j B_\xi \Psi(\|x - \xi_j\|) + \sum_{j=NB+1}^{N} \lambda_j L_\xi \Psi(\|x - \xi_j\|) + \sum_{j=1}^{NP} \lambda_{j+N} P_{m-1}^j(x) \tag{5}$$

Collocating in a similar way leads to the system of equations

$$\begin{bmatrix} B_x B_\xi[\Psi] & B_x L_\xi[\Psi] & B_x[P_{m-1}] \\ L_x B_\xi[\Psi] & L_x L_\xi[\Psi] & L_x[P_{m-1}] \\ B_\xi[P_{m-1}]^T & L_\xi[P_{m-1}]^T & 0 \end{bmatrix} \lambda = \begin{bmatrix} g \\ f \\ 0 \end{bmatrix} \tag{6}$$

In the above matrix equation, the operators with ξ subscript are applied at the trial points and the operators with x subscript are applied to test points.

The approach based on the interpolation formula of Eq. (5) is known as the Hermitian, or symmetric method, producing a collocation matrix which is symmetric (hence attractive for storage benefits), and was shown by Wu [28,29] to be non-singular provided that no two data-centres sharing the same operator are placed at the same location. Eq. (5) will be the base of our Local Hermitian Interpolation method, where for simplicity in the presentation the $L[]$ operator will be referred to as the PDE-operator and $B[]$ as the boundary operator.

Similar matrix conditioning difficulties as those found with the Unsymmetric approach (Kansa), are also encountered when using the Symmetric approach (Hermitian) with a large number of data-centres.

3. Mathematical formulation of the Local Hermitian Interpolation method

The solution space Ω is covered by two sets of (quasi-)scattered data-centres, as shown in Fig. 1. These two sets of data-centres represent the locations at which the solution value and the PDE-operator will be enforced within the interpolation function. The two sets of data-centres may be placed independently (as shown), or may occupy the same physical location. If no PDE-operator data-centres are present within the domain, the implementation becomes identical to that described in [20,21]. Data-centres are also placed on all boundary surfaces Γ. If a stencil is sufficiently close to the boundary, the data-centres at the intersection between the boundary and the stencil are also included in the local interpolation (see Fig. 1), i.e. including the boundary operator $B[]$. A list of definitions for locations at which data is stored or manipulated is given in the Appendix (Table A1).

Each solution data-centre has associated with it a local system, comprised of other solution centres and also PDE- and boundary centres where present. Local system sizes are generally in the region of around 10–30 data-centres in 3D.

The solution is approximated via a series of Hermitian interpolations on each of the local systems, using the functional values for solution centres, the source term values for PDE centres, and the boundary operator values at boundary centres in those cases where it is required. As with the full-domain Hermitian method, the various operators are applied to the basis functions at their respective trial points. A polynomial term is included to complete the underlying vector space, as with Eq. (5).

In this way, at each local system the field variable $u(\mathbf{x})$ is approximated by

$$u^{(k)}(x) = \sum_{j=1}^{NS} \alpha_j^{(k)} \Psi(\|x - \xi_j\|) + \sum_{j=NS+1}^{NS+NB} \alpha_j^{(k)} B_\xi[\Psi(\|x - \xi_j\|)] + \sum_{j=NS+NB+1}^{NS+NB+NPDE} \alpha_j^{(k)} \overline{L}_\xi[\Psi(\|x - \xi_j\|)] + \sum_{j=1}^{NP} \alpha_{j+N} P_{m-1}^j(x) \tag{7}$$

where k is the local system index; NS is the number of solution centres in the local system; NB is the number of boundary centres in the local system; $NPDE$ is the number of PDE centres in the local system; NP is the number of terms in a 3D polynomial of degree $m-1$.

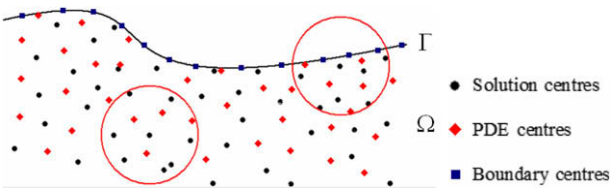

Fig. 1. Schematic representation of local systems.

Consider a Local Hermitian Interpolation of a general scalar field, $u(\mathbf{x})$, which satisfies a PDE-operator $\overline{L}$ in the domain Ω, and a boundary operator B at the boundary Γ, with values of the scalar field and operators given at a discrete set of points, as outlined in Fig. 1

$$\begin{aligned} &\overline{L}[u(\mathbf{x})] = S(\mathbf{x}) \quad \text{in } \Omega \\ &u(\mathbf{x}_i) = f_i \quad \text{in } \Omega \\ &B[u(\mathbf{x})] = g(\mathbf{x}) \quad \text{on } \Gamma \end{aligned} \tag{8}$$

In the proposed LHI meshless collocation method, the scalar values f_i will correspond to the unknown solution values within the computational domain Ω, which are to be determined by the solution of a global system of equations. This description can be collocated at each of the data-centres within the local systems, Ω^k, described by Eq. (7), applying once again the boundary operator at boundary centres and the PDE-operator at the PDE centres. This leads to a series of small symmetric linear systems which can be solved independently to interpolate the solution over the domain Ω in a piecewise fashion. Therefore, a series of local systems

$$A^{(k)}\alpha^{(k)} = d^{(k)} \tag{9}$$

are constructed for the interpolation coefficients $\alpha^{(k)}$, where

$$A^{(k)} = \begin{bmatrix} \Psi_{ij} & B_\xi[\Psi_{ij}] & \overline{L}_\xi[\Psi_{ij}] & P_{m-1} \\ B_x[\Psi_{ij}] & B_xB_\xi[\Psi_{ij}] & B_x\overline{L}_\xi[\Psi_{ij}] & B_x[P_{m-1}] \\ \overline{L}_x[\Psi_{ij}] & \overline{L}_xB_\xi[\Psi_{ij}] & \overline{L}_x\overline{L}_\xi[\Psi_{ij}] & \overline{L}_x[P_{m-1}] \\ P_{m-1}^T & B_\xi[P_{m-1}]^T & \overline{L}_\xi[P_{m-1}]^T & 0 \end{bmatrix} \quad \text{and} \quad d^{(k)} = \begin{bmatrix} f_i \\ g_i \\ S_i \\ 0 \end{bmatrix} \tag{10}$$

where $\Psi_{ij} = \Psi(\|\mathbf{x}_i - \xi_j\|)$ for data-centres contained in local system k, ordered as indicated in Eq. (7).

In this way the value of $u(\mathbf{x})$ close to point k can be written as

$$u^{(k)}(\mathbf{x}) = H(\mathbf{x})^{(k)}\alpha^{(k)}$$

where

$$H(\mathbf{x})^{(k)} = [\Psi(\mathbf{x} - \xi_j), \quad B_\xi[\Psi(\mathbf{x} - \xi_j)], \quad \overline{L}_\xi[\Psi(\mathbf{x} - \xi_j)], \quad P_{m-1}]^T \tag{11}$$

Similarly, any partial differential operator, Q, can be applied to the reconstruction vector $H^{(k)}$ in order to reconstruct the partial derivatives of the interpolated function:

$$Q[u^{(k)}(\mathbf{x})] = Q[H^{(k)}]\alpha^{(k)} = [Q_x[\Psi(\mathbf{x} - \xi_j)], \quad Q_xB_\xi[\Psi(\mathbf{x} - \xi_j)], \quad Q_x\overline{L}_\xi[\Psi(\mathbf{x} - \xi_j)], \quad Q_x[P_{m-1}]]\alpha^{(k)} \tag{12}$$

By multiplying the operator-reconstruction vector by the local matrix inverse, a vector of stencil weights can also be obtained:

$$Q[u^{(k)}(\mathbf{x})] = Q[H^{(k)}(\mathbf{x})]\alpha^{(k)} = (Q[H^{(k)}(\mathbf{x})][A^{(k)}]^{-1})d^{(k)} = W_Q^{(k)}(\mathbf{x})d^{(k)} \tag{13}$$

3.1. Solution technique for steady problems

For the solution of steady, boundary value PDEs, the solution field $u(\mathbf{x})$ is unknown. It is hence necessary to obtain an implicit relation for u at the solution centres.

For a steady PDE:

$$\begin{aligned} &\overline{L}[u] = S(\mathbf{x}) \quad \text{in } \Omega \\ &B[u(\mathbf{x})] = g(\mathbf{x}) \quad \text{on } \Gamma \end{aligned} \tag{14}$$

With $A^{(k)}$ as indicated in Eq. (10), the data-vector becomes

$$d^{(k)} = \begin{bmatrix} u_i \\ g_i \\ S_i \\ 0 \end{bmatrix} \tag{15}$$

with $f_i = u_i$, representing the unknown values of $u(\mathbf{x})$ at the local solution centres, and the rest of the data-vector composed of known quantities.

By applying the PDE-operator, $\overline{L}[u] = S(\mathbf{x})$, to a reconstruction vector at the local system centrepoint, a relation can be obtained linking the values of u_i within the local system (Eq. (14))

$$\begin{aligned} S_{centre} &= \overline{L}[u^{(k)}(\mathbf{x}_{centre})] = \overline{L}[H^{(k)}(\mathbf{x}_{centre})]\alpha^{(k)} \\ &= [\overline{L}_x[\Psi(\mathbf{x}_{centre} - \xi_j)], \quad \overline{L}_xB_\xi[\Psi(\mathbf{x}_{centre} - \xi_j)], \quad \overline{L}_x\overline{L}_\xi[\Psi(\mathbf{x}_{centre} - \xi_j)], \quad \overline{L}_x[P_{m-1}]]([A^{(k)}]^{-1}d^{(k)}) = W_{\overline{L}}^{(k)}(\mathbf{x}_{centre})d^{(k)} \end{aligned} \tag{16}$$

By applying the above reconstruction to each local system k, a series of N simultaneous equations are produced for $u_i,\ i = 1, 2, \ldots, N$, where N is the global number of solution centres. In the resulting global linear system the corresponding boundary conditions of the problem have already been imposed, at the local interpolation, in those stencils containing boundary centres. This sparse linear system of equations can be solved efficiently using standard solution techniques for sparse linear systems. The GMRES algorithm from the SPARSKIT toolbox [30] is currently used to solve such systems.

Note that the PDE-reconstruction formula, Eq. (16), can be applied on a system-by-system basis. This eliminates the need to store the entire set of local system matrices; each matrix $A^{(k)}$ can be calculated, inverted, used to form the appropriate line in the sparse global system, and then deallocated. This can significantly reduce storage requirements for implementing the solution algorithm.

3.2. Solution technique for time-dependent problems

The method of solution for time-dependent problems, using implicit time advancement, requires the creation of a modified PDE-operator via a finite difference approximation of the time-derivative. The procedure is illustrated here using a Crank–Nicolson approach, but is easily extensible to any number of finite difference time advancement formulae.

For a general initial–boundary value problem

$$\begin{aligned} &\frac{\partial u(\mathbf{x},t)}{\partial t} = L[u(\mathbf{x},t)] + S'(\mathbf{x}) \quad \text{in } \Omega \\ &u(\mathbf{x},0) = f(\mathbf{x}) \\ &B[u(\mathbf{x},t)] = g(\mathbf{x},t) \quad \text{on } \Gamma \end{aligned} \tag{17}$$

a finite difference approximation is made to the time-derivative

$$\frac{u^n - u^{n-1}}{\Delta t} = \theta L[u^n] + (1-\theta)L[u^{n-1}] + S'(\mathbf{x}) \tag{18}$$

From this approximation, a modified PDE-operator is obtained

$$\overline{L}[u^n] = \widehat{L}[u^{n-1}] + S'(\mathbf{x}) \tag{19}$$

where

$$\begin{aligned} \overline{L} &= 1 - \theta \Delta t L \\ \widehat{L} &= 1 + (1-\theta)\Delta t L \end{aligned}$$

The proposed time stepping algorithm implies that the original initial–boundary value problem reduces at each time step to the solution of a boundary value problem defined by the non-homogeneous partial differential Eq. (19), with the non-homogeneous term given in terms of the solution of the problem at the previous time step. At this stage the formulation at a given time step is identical to the formulation for steady problems, assuming that $\widehat{L}[u^{n-1}]$ is known from interpolation at the previous time step. $S(\mathbf{x}) = \widehat{L}[u^{n-1}] + S'(\mathbf{x})$ is taken as the equivalent to the non-homogeneous term of the PDE in Eq. (14). As such the solution procedure at each time step is performed in the same way as for the steady problem, using the modified operator and the corresponding non-homogeneous term.

After solution of the global linear system, and using the current interpolation matrix systems and updated data-vectors, a reconstruction of $\widehat{L}[u^n]$ at solution and PDE centres must be made, ready for the calculation of $S(\mathbf{x})$ at the next time step. This is done via the creation of further reconstruction arrays

$$\widehat{L}[u^n(\mathbf{x})] = \widehat{L}[H^{(k)}(\mathbf{x}_{centre})]\alpha^{(k)} = [\widehat{L}_x[\Psi(\mathbf{x}-\xi_j)],\quad \widehat{L}_x B_\xi[\Psi(\mathbf{x}-\xi_j)],\quad \widehat{L}_x \overline{L}_\xi[\Psi(\mathbf{x}-\xi_j)],\quad \widehat{L}_x[P_{m-1}]]([A^{(k),n}]^{-1} d^{(k),n}) = W^{(k)}_{\widehat{L}}(\mathbf{x}) d^{(k)} \tag{20}$$

As with the steady case, it is not necessary to store the entire set of local matrices. The stencil weight vectors for the reconstruction of $\widehat{L}[u]$ at the required locations can be calculated on a system-by-system basis, then stored and multiplied with the updated data-vectors after the solution of the global system.

The initial time step is always performed using the $\theta = 1$ first-order implicit time stepping formulation. Were a value of $\theta < 1$ to be used, $L[u^0]$ would be required in order to calculate the S_i which appear in the local system data-vector of Eq. (15). This quantity is unknown at the initial configuration, and in order to be calculated would require an interpolation of the initial solution field, using a different interpolation system which does not rely on a non-existent previous time step. At subsequent time steps, any value of $0 < \theta \leqslant 1$ can be chosen without altering the interpolation system.

If the modified PDE-operator $\overline{L}$ changes with time, then local system matrices must be re-formed and re-inverted at each time step, in order to calculate the new reconstruction vectors. However, if the modified PDE does not change with time it is desirable to store all of the stencil weight vectors throughout the time-loop, as this eliminates the need to re-calculate the local system matrix inverses at each time step.

3.3. Shape-parameter

It is important to note that the value of the multiquadric shape-parameter, c, has not been explicitly defined (see Table 1). The shape-parameter represents a local length-scale. As such, in datasets which include local systems that vary significantly in size, the non-dimensionalised shape-parameter, c^*, should be applied instead:

$$c^* = \frac{c}{\Delta_{\max}} \tag{21}$$

where $\Delta_{\max}$ is the separation from the local system centrepoint to the most distant point included in the system.

Determination of an optimal c^* value is a nontrivial issue, and requires further research for local RBF methods. With the current formulation it has been observed, experimentally, that a value of c^* between 2.0 and 5.0 appears to provide accurate and fairly consistent results for most convection-diffusion problems on regular data-centre distributions. On irregular data-centre distributions, values of c^* between 1.0 and 2.0 appear preferable.

An efficient and reliable procedure to determine the optimal shape-parameter value(s) is currently outstanding, for local RBF methods in general. Many authors examine the condition number of the local system matrices (see for example [31]), and adjust the shape-parameter to match a target maximum conditioning value. However, it can be demonstrated that the optimal value of the shape-parameter is not dependent only upon the conditioning of the local matrices, but also on the nature of the data to be interpolated (see [32] or [33] for demonstrations of data-dependent shape-parameter behaviour with full-domain RBF methods, or [20] for an example using spatially varying shape-parameter values to improve the LHI method).

In order to obtain a high quality optimisation of the shape-parameter it is possible to sample the value of the local PDE residual, at locations other than the data-centres, via a simple reconstruction (see [34] for an example of implementation). Minimisation of this quantity over the solution domain via alteration of c^* will represent a near-optimal solution. However, such a procedure is extremely expensive, requiring the problem to be run several times in order to obtain the optimal c^* value. Additionally, for transient problems the optimisation may need to be performed many times, as the properties of the solution field evolve. In general, problems involving sharp changes in gradient and other near-discontinuous features tend to benefit from lower values of c^*.

The authors are in the process of investigating alternative procedures for shape-parameter selection with the LHI method, which take into account the nature of the solution field at a viable computational cost. However, throughout the current work a suitable shape-parameter is selected for each example problem and is used throughout, unless otherwise specified.

4. Stencil configurations

The choice of local system stencil configurations can have a significant effect on both solution quality and computational cost. The size of the global linear system is dependent only on the number of local systems, and the number of solution centres in each local system stencil. As such, the cost of solving the global system is strongly dependent on the number of solution centres in each stencil. However, the cost of forming and inverting local system matrices is dependent on the total number of data-centres in each system.

The cost of performing each iteration in solving the global system will scale linearly with the total number of solution centres, however the cost of forming each interpolation system scales with the fourth power of the total number of centres in the stencil, when using a direct solution method. As such, choosing a larger stencil size will necessarily increase the cost of setting up the required interpolation systems (proportional to the fourth power of the number of data-centres in the local stencils). However, the larger stencil sizes may lead to more rapid convergence when utilising an iterative solver at global level. As such, for transient problems, the most computationally efficient stencil size will often be problem-dependent.

For transient problems where the $\overline{L}$ operator changes with time it is generally desirable to choose a relatively small stencil size where possible, as the local linear systems must be re-formed and solved at every time step. It is worth noting that, even with a relatively small stencil size, the computational cost of the algorithm will increase significantly when the $\overline{L}$ operator varies with time. In such a scenario it may, in certain cases, be possible to obtain an acceptable solution for a given computational cost without the PDE centres (where re-inversion is not required), by utilising a finer dataset. However, for a high-precision and stable solution, PDE centres should always be included.

4.1. Collocated stencil configurations

The Hermitian RBF interpolation procedure used by the LHI method allows the placement of multiple data-centres in the same physical location, so long as the operators applied there are unique [28,29]. This flexibility can be exploited to allow stencil configurations where the PDE centres are collocated with the solution centres. Such a configuration allows a reduction in the number of locations at which data must be stored.

It is worth noting, however, that a PDE centre cannot be placed at the local system centrepoint, where the reconstruction of the modified PDE-operator takes place (Eq. (16)) in order to form the global system of equations. In conceptual terms this is to avoid a repeated application of the PDE-operator at a single location. In practical terms, the reconstruction vector $\overline{L}[H^{(k)}(\mathbf{x}_{centre})]$ of Eq. (16) becomes identical to the equivalent row of the local system matrix described by Eq. (10). As a

consequence, the stencil weight vector does not link with other values of u_i. The row of the global linear system becomes zero, leading to a degenerate system of equations.

4.2. Staggered stencil configurations

A configuration of PDE centres which is distinct from the solution centres allows greater flexibility and control over their placement. However such a configuration increases the number of data-storage locations, and may also increase the complexity of initial dataset generation.

For time-dependent problems the modified PDE source term $\widehat{L}[u^n]$ must be calculated at every time step, as described by Eq. (19). Staggered PDE centres may appear in several local systems, and as such the predicted value of $\widehat{L}[u^n]$ at the PDE centres may differ slightly depending on the local system chosen for reconstruction. In order to avoid instabilities it is desirable to formulate a globally consistent value for each PDE centre, based on reconstructions from the local systems in which it appears.

Stencil weight vectors are created for the reconstruction of $\widehat{L}[u^n]$ from the $[A^{(k)}]^{-1}$ matrix and stored, until updated data is retrieved from the solution of the global linear system, as described by Eq. (20). The number of vectors which must be stored increases with the number of systems used in the reconstruction of each PDE centre. If the data-centre appears in several local systems, the relative increase in storage can be significant.

5. Time-dependent numerical benchmark

The following unsteady convection-diffusion equation is solved

$$\frac{\partial \phi}{\partial t} = D\frac{\partial^2 \phi}{\partial x_i^2} - u_i \frac{\partial \phi}{\partial x_i} \tag{22}$$

over the domain

$$\begin{aligned} &0 \leqslant x \leqslant 1.0 \\ &0 \leqslant y \leqslant 0.2 \\ &0 \leqslant z \leqslant 0.2 \end{aligned}$$

with

$$\begin{aligned} &\mathbf{u} = (1, \quad 0, \quad 0)^T \\ &\phi(x = 0, t) = 2.0 \\ &\phi(x > 0, t = 0) = 1.0 \end{aligned} \tag{23}$$

Zero flux is applied at the y and z boundaries, leading to a one-dimensional solution:

$$\phi(x,t) = 1.0 + \frac{1}{2}\left(erfc\left(\frac{x-t}{2\sqrt{Dt}}\right) + e^{\frac{x}{D}} erfc\left(\frac{x+t}{2\sqrt{Dt}}\right)\right) \tag{24}$$

For $D = 0$ the analytical solution is a simple travelling shock front.

This problem is run using a regular Cartesian distribution of data-centres in 3D. To test the effect of altering local system size and configuration, a variety of stencils are examined (see Figs. 2 and 3), where for convenience only 2D projections of the stencils are presented in the figures. Besides, for simplicity, all the stencils showed in Figs. 2 and 3 are defined in a structural arrangement. However in the general case, the stencils are defined in an unstructured topology.

Two sets of stencils are examined, each set having a different number of solution centres. The first set links the system centrepoint to its immediate neighbours, creating a 7-point stencil in 3D. A 2D representation of this configuration is given in Fig. 2. The second set links each system centrepoint to every point having a Cartesian location index within one jump of itself, representing a larger local support size and creating a 27-point stencil in 3D. The corresponding 2D representation is given in Fig. 3.

For each set of stencil sizes, three configurations of PDE centres are tested. The first configuration has no PDE centres, relying only on solution and boundary operator interpolation, as was the case with the previous LHI formulation described in [20,21] (Figs. 2(a) and 3(a)). In this case, the third term on the right hand side of the interpolation function (7) is not included and in the corresponding interpolation matrix (10) the third column and row are not considered. In this way, at the local level we have a simple interpolation that takes into account the boundary conditions on those stencils containing boundary centres. The second configuration uses a collocated distribution of PDE centres, placing PDE centres at each point which lies immediately adjacent to the stencil centrepoint (Figs. 2(b) and 3(b)). No PDE centre is placed at the stencil centrepoint, leading to a formulation with six collocated PDE centres in 3D. The third configuration places the PDE centres at a staggered location. The centres are placed at each half-index immediately surrounding the stencil centrepoint, producing eight staggered PDE centres in 3D (Figs. 2(c) and 3(c)).

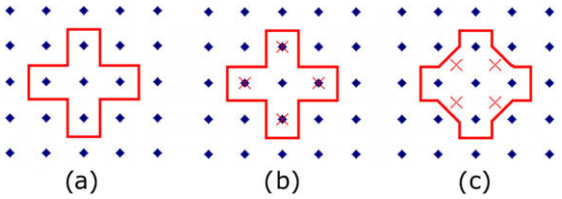

Fig. 2. 2D cross-sections of (a) 7 point, (b) 7+6c point and (c) 7+8s point stencils. The 3D configuration appears as above in the x–y, x–z and y–z cross-sections.

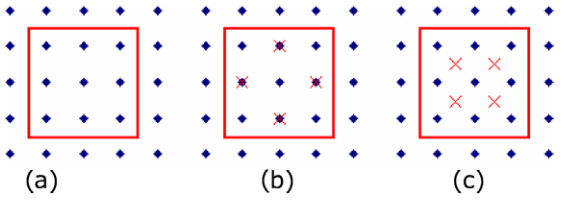

Fig. 3. 2D cross-sections of (a) 27 point, (b) 27+6c point and (c) 27+8s point stencils. The 3D configuration appears as above in the x–y, x–z and y–z cross-sections.

In the second and third configurations, using PDE centres, an implicit upwind scheme is introduced in the formulation by having an interpolation which satisfies locally the partial differential operator at neighbouring data-centres, including the desired information about the convective velocity field.

In the schematic diagrams of the stencil configuration, the diamond symbol corresponds to the locations where the value of the function is interpolated (the solution centres), and the crosses to the locations where the PDE-operator is imposed (the PDE centres). Note that when each of these stencils is placed sufficiently close to a boundary (within one Cartesian index of a boundary centre), some of the solution centres will be replaced with boundary centres.

To test stencil performance, a point separation of $\Delta = 0.025$ is used, leading to a uniform distribution of $(41 \times 9 \times 9)$ boundary and solution centres. A time step of $\Delta t = 10^{-3}$ is used, with the $\theta = 0.65$ time stepping scheme. Each test is run until $t = 0.5$. A shape-parameter of $c = 0.2$ is used throughout, to represent a non-dimensional shape-parameter (c^*) value of close to 5. It is worth noting here that the shape-parameter is chosen to provide reasonable results all around, but has not been chosen to minimise the overall error against the analytical solution. If this were the case, then a different shape-parameter value would be required for each stencil configuration, and perhaps also for each Péclet number and time step size used.

To assess the stabilising effect of the implicit upwinding introduced by the PDE-operator in the interpolation formula, cases at different Péclet number are simulated both with and without the use of the local PDE-operator. Comparisons between the results obtained with the two approaches are reported in Figs. 4–6. The figures show the performance of both approaches when using various stencil configurations at Pe = 500, Pe = 1000 and Pe = 2000 respectively. In each case, the addition of PDE centres has a stabilising effect in comparison to the basic RBF interpolation, as well as offering a more precise capture of the travelling front. The best stencil configuration in this case is the 27-point stencil with staggered PDE centres (27+8s). This stencil is completely stable in the Pe = 1000 case, and exhibits the least oscillation in the Pe = 2000 case. As will be shown later, these solutions can be improved by increasing the number of data points and reducing the time step. In this way, it is even possible to obtain stable solution at Péclet infinity.

It is observed from the numerical simulations that the 7-point stencil with staggered PDE centres (7+8s) produced unstable oscillations at high Pe values. The source of the instability appears to be inaccuracy in the reconstruction of the $\widehat{L}[\phi^{n-1}]$ operator value at the staggered PDE centres. The L_2 error norms shown in Table 2 represent the accuracy of each stencil configuration. Apart from the unsuccessful 7+8s stencil, the inclusion of PDE centres reduces the solution error by a factor of between five and ten when compared to the best stencil without PDE centres.

The effect of varying the time step is examined using the same $\Delta = 0.025$ point distribution at Pe = 1000 and Pe = 2000, and with the 27+8s stencil. The time step is varied between $\Delta t = 10^{-4}$ and $\Delta t = 10^{-2}$. For the $\theta = 1.0$ implicit time stepping scheme, Fig. 7, and the $\theta = 0.65$ mixed scheme, Fig. 8, increasing the time step appears to stabilise the solution, at the cost of introducing a diffusive error component. This does not appear to be the case with the $\theta = 0.5$ Crank–Nicolson scheme, Fig. 9, where increasing the time step instead has an adverse effect on stability. It is possible to obtain a stable solution to the problem at Pe = 1000 with any of the time stepping schemes. At Pe = 2000 it is also possible to obtain a stable solution by using a sufficiently large time step with the $\theta = 1.0$ and $\theta = 0.65$ schemes. However, with the $\theta = 0.5$ scheme it is not possible to obtain a stable solution to the Pe = 2000 problem on a dataset of the given density.

Using a refined distribution of data-centres, an improved solution can be achieved at high Pe numbers. Table 3 shows the reduction in L_2 error norm values as a progressively refined dataset is introduced, at Pe = 2000. A time step of $\Delta t = 10^{-4}$ is

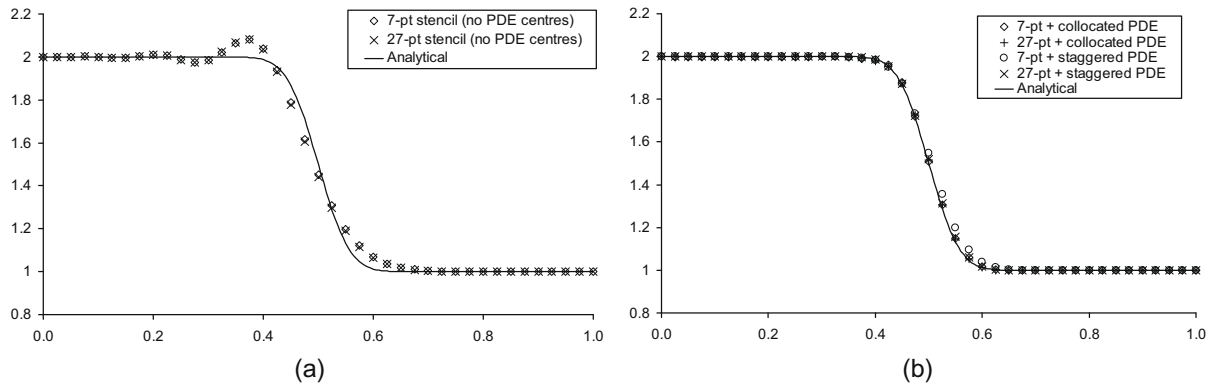

Fig. 4. Travelling shock at Pe = 500: (a) without PDE centres, (b) using PDE centres.

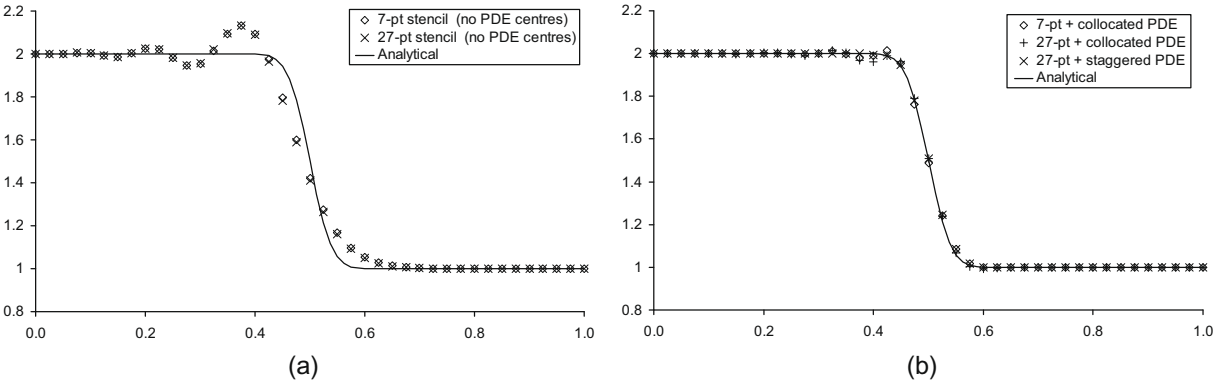

Fig. 5. Travelling shock at Pe = 1000: (a) without PDE centres, (b) using PDE centres.

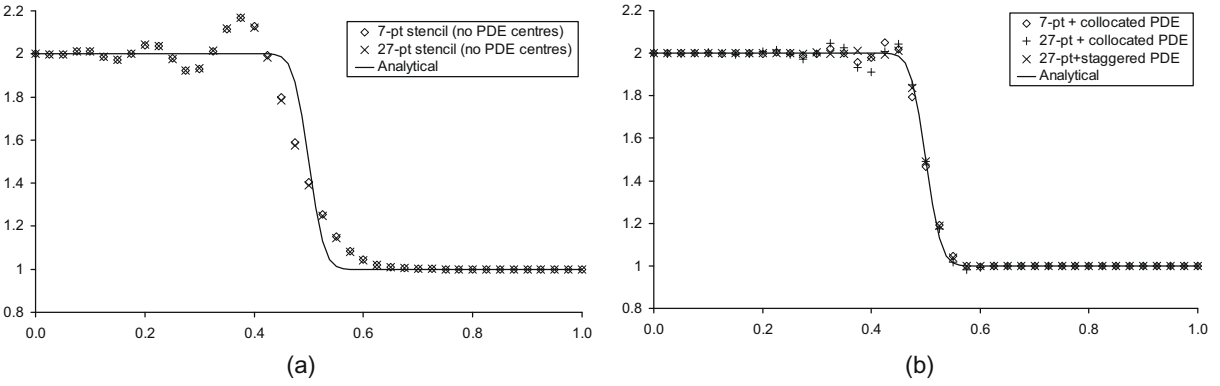

Fig. 6. Travelling shock at Pe = 2000: (a) without PDE centres, (b) using PDE centres.

Table 2
L_2 error norms for tested stencil configurations ($\theta = 0.65$ scheme, $\Delta t = 10^{-3}$).

	Stencil					
	7+0	27+0	7+6c	27+6c	7+8s	27+8s
Pe = 500	2.60×10^{-2}	2.69×10^{-2}	3.37×10^{-3}	3.34×10^{-3}	5.48×10^{-3}	3.37×10^{-3}
Pe = 1000	4.26×10^{-2}	4.39×10^{-2}	5.65×10^{-3}	5.39×10^{-3}	–	3.97×10^{-3}
Pe = 2000	5.75×10^{-2}	5.89×10^{-2}	1.27×10^{-3}	1.20×10^{-3}	–	7.14×10^{-3}

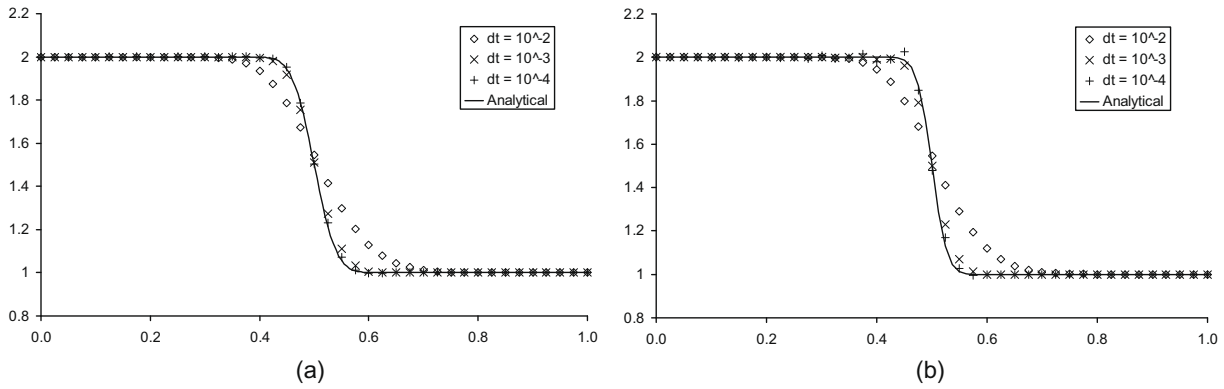

Fig. 7. Effect of time step for $\theta = 1.0$ implicit scheme: (a) Pe = 1000, (b) Pe = 2000.

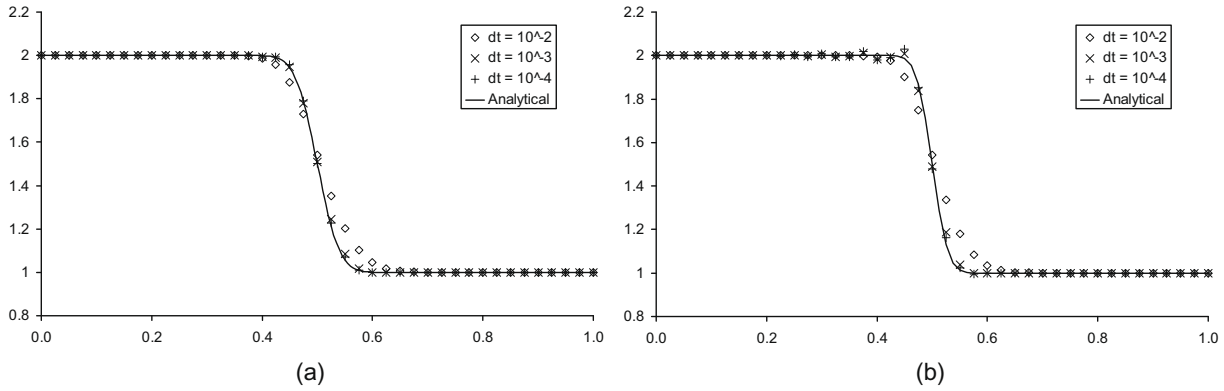

Fig. 8. Effect of time step for $\theta = 0.65$ mixed scheme: (a) Pe = 1000, (b) Pe = 2000.

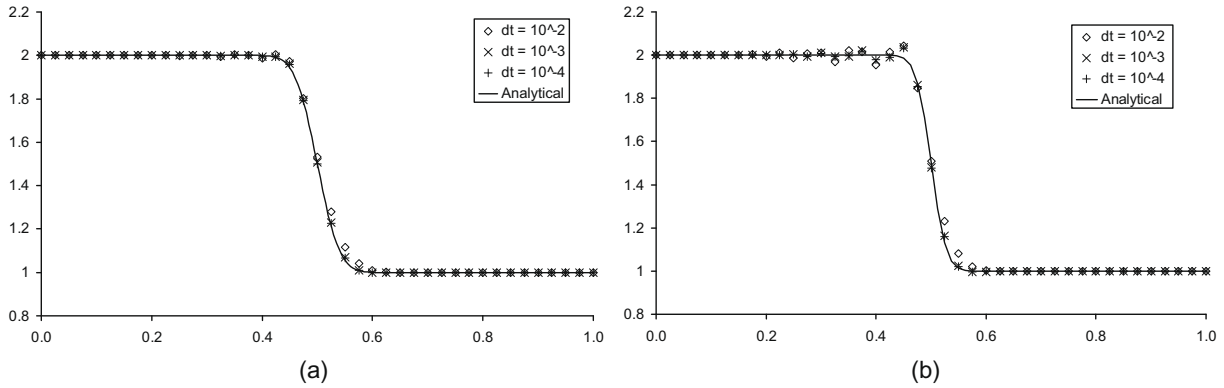

Fig. 9. Effect of time step for $\theta = 0.5$ Crank–Nicolson scheme: (a) Pe = 1000, (b) Pe = 2000.

Table 3
Solution convergence at Pe = 2000 with 27+8s stencil ($\theta = 0.65$ scheme, $\Delta t = 10^{-4}$).

Point separation	0.05	0.0333...	0.025	0.01666...	0.0125
L_2 solution error	1.62×10^{-2}	1.17×10^{-2}	6.57×10^{-3}	2.02×10^{-3}	1.30×10^{-3}

used, along with the $\theta = 0.65$ scheme. The shape-parameter is scaled with the local-system length-scale (see Eq. (21)), in order to maintain a consistent value of c^*. The solution convergence is roughly second order in this case.

Using a dataset with a point separation of $\Delta = 0.0125$ allows improved capturing of sharp fronts. Fig. 10 shows the performance at Pe = 2000, Pe = 5000 and infinite Péclet number ($D = 0$). At this time step the $\theta = 0.65$ scheme was required for stable capture of the Pe = 5000 front, and a $\theta = 1.0$ scheme was required at infinite Péclet number.

If in the present case, our previous explicit approach [20] is used, i.e. without introducing any kind of upwinding scheme on the selection of the interpolation stencils, the obtained numerical result is similar to the one predicted by the no PDE centres configuration, i.e. the solutions reported in Figs. 4(a), 5(a) and 6(a), with the corresponding dispersive error or instability.

6. Steady-state numerical benchmark

The steady convection-diffusion-reaction equation is solved

$$\frac{\partial^2 \phi}{\partial x_i^2} - (A_i + k x_i)\frac{\partial \phi}{\partial x_i} - k\phi = 0 \tag{25}$$

on the domain

$$0 \leqslant x \leqslant 1.0$$
$$0 \leqslant y \leqslant 0.2$$
$$0 \leqslant z \leqslant 0.2$$

with

$$\phi(0) = \phi_0 = 300$$
$$\phi(1) = \phi_1 = 100$$
$$A_i = \left[\ln\left(\frac{\phi_1}{\phi_0}\right) - \frac{k}{2}\right]\delta_{1i} \tag{26}$$

Zero-flux conditions are applied to the y and z boundaries. As such the problem has a one-dimensional analytical solution given by

$$\phi(x_i) = \phi_0 e^{\left(\frac{kx_1^2}{2} + Ax_1\right)} \tag{27}$$

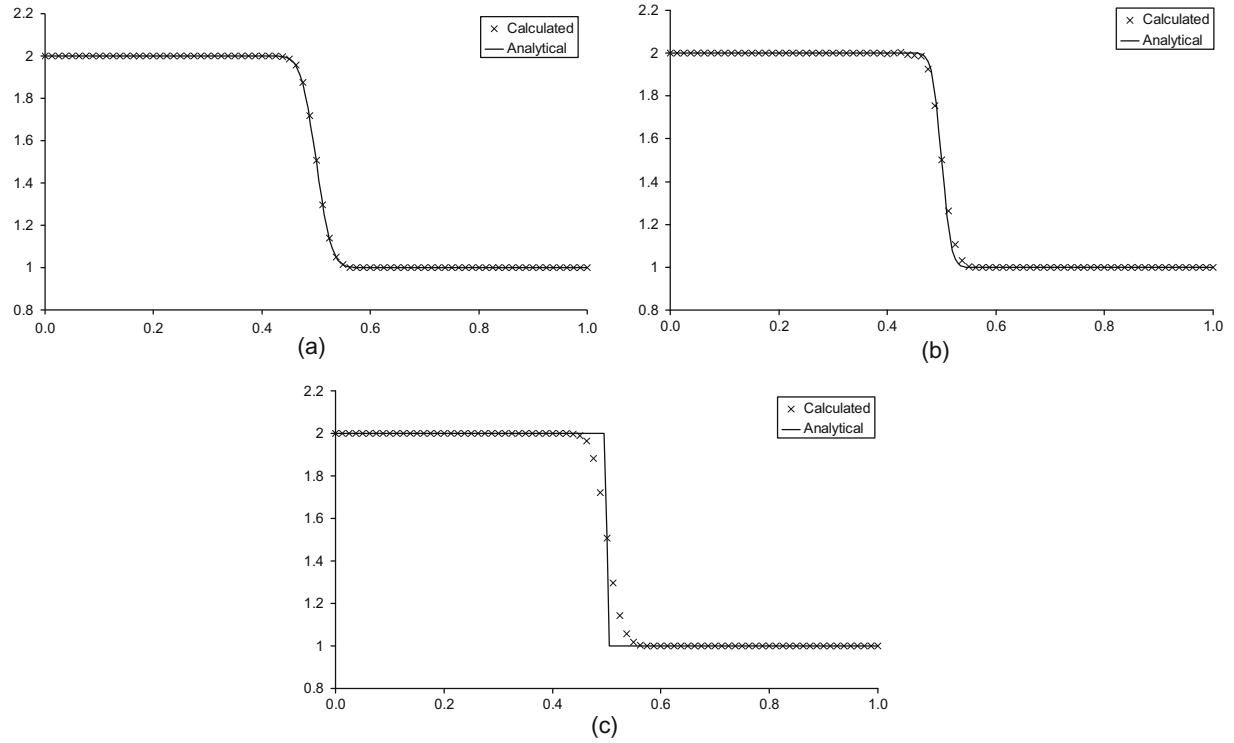

Fig. 10. Front capturing with $\Delta t = 10^{-3}$ and $(81 \times 17 \times 17)$ dataset: (a) Pe = 2000 with $\theta = 0.5$, (b) Pe = 5000 with $\theta = 0.65$, (c) infinite Péclet number with $\theta = 1.0$.

For $k > 2\ln(3)$, the velocity field changes sign within the domain. As such, the scalar field is pushed towards either end of the domain, with the central region left relatively empty. This effect is magnified as k is increased, with the minimum value decreasing exponentially with k (see Eq. (28))

$$\phi_{\min} = \frac{\phi_0}{\sqrt{3}} e^{-\left(\frac{\ln(3)^2}{2k} + \frac{k}{8}\right)} \approx \frac{\phi_0}{\sqrt{3}} e^{-\frac{k}{8}} \quad \text{for large values of } k \tag{28}$$

This presents a numerical difficulty, since the large values of ϕ at the domain ends and the small values towards the domain centre must be modelled accurately. This test case has proven very difficult to reproduce at values of $k \geqslant 40$, when using full-domain RBF methods [35].

As with the transient test case, the 27+8s stencil performs the best in most circumstances. However, the 27+6c stencil performs very poorly, often worse than the stencils without PDE centres. Apart from this particular stencil, the inclusion of PDE centres significantly improves solution quality. Table 4 shows the performance of the various stencils for $k = 120$, with a shape-parameter of $c = 0.08$ applied to each stencil. Computational cost is also recorded, excluding the time taken to read data files or to post-process data. Calculations were performed on a 2.83Ghz Intel core2 CPU, with the FORTRAN 90 code used throughout this work. The computational times broadly follow the expected trends, with the cost of forming the interpolation systems significantly outweighing the cost of solving the global linear system. It is worth noting that the 7+6c and 7+8s stencils produce results to a higher degree of accuracy than the 27+0 stencil, at a reduced computational cost.

In addition to an improvement in absolute error, the small values close to the centre of the domain can be better captured by using PDE centres. This is demonstrated in Figs. 11–13, which offer a comparison of the obtained results with the 27+8s stencil, series 1 in the Figures, against those obtained with the 27-point stencil without PDE centres, series 2 in the figures. A constant shape-parameter of value $c = 0.1$ is used for the $k = 80$ case, Fig. 11. Owing to the increasing sharpness of the solution profile as k is increased, the shape-parameter is reduced slightly to $c = 0.08$ for the $k = 120$ case, Fig. 12, and $c = 0.06$ for $k = 160$, Fig. 13. More information on the data-dependent variation of shape-parameters for the LHI method with this test case can be found in [20].

In each of the three cases tested, the configuration with the PDE centres performs significantly better throughout the domain. The addition of the PDE centres allows a good solution to be obtained at $k = 160$ on this relatively coarse dataset. It is important to point out that the apparent high relative error obtained at the centre of the domain is due to the dimensionless form used in the definition of the relative error (the difference between the numerical and analytical solutions divided by the absolute value of the analytical solution), which results in division by a very small value of the potential in the central region of the domain. It should be noted that, even in this case, the relative error obtained with the use of the PDE-operator is very small, almost predicting the near-zero values in the central region.

In cases such as this, it may be desirable to improve solution accuracy via local refinement of the dataset, for example in the end regions in order to more accurately capture the sharply varying solution profile. Such a refinement is performed in

Table 4
L_2 error norm and computational cost for tested stencil configurations. $k = 120$, $c = 0.08$.

Stencil configuration	7+0	27+0	7+6c	27+6c	7+8s	27+8s
L_2 solution error	1.7104	2.62798	0.301	2.678	0.289	0.074
CPU time (s)	0.312	0.866	0.416	1.185	0.425	1.378

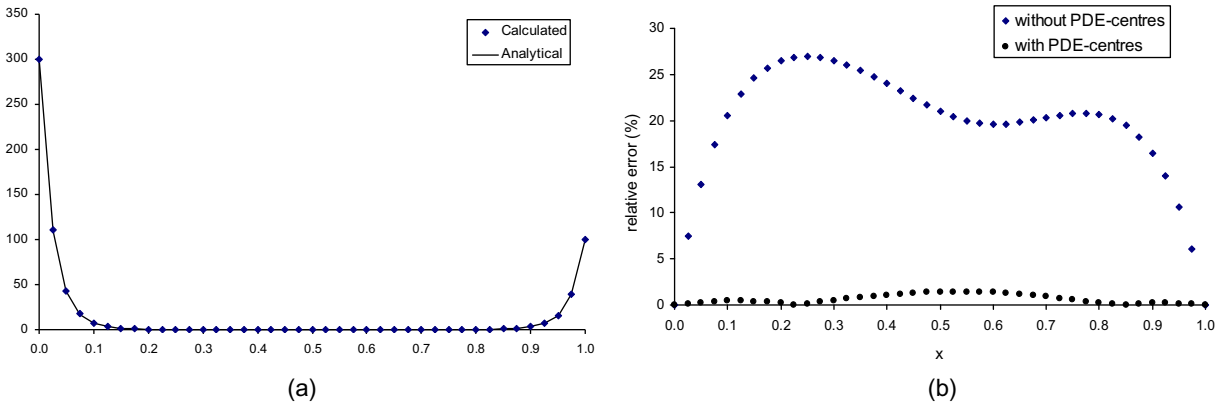

Fig. 11. Solution at $k = 80$ ($\phi_{\min} = 7.86 \times 10^{-3}$): (a) solution profile, (b) relative error profile.

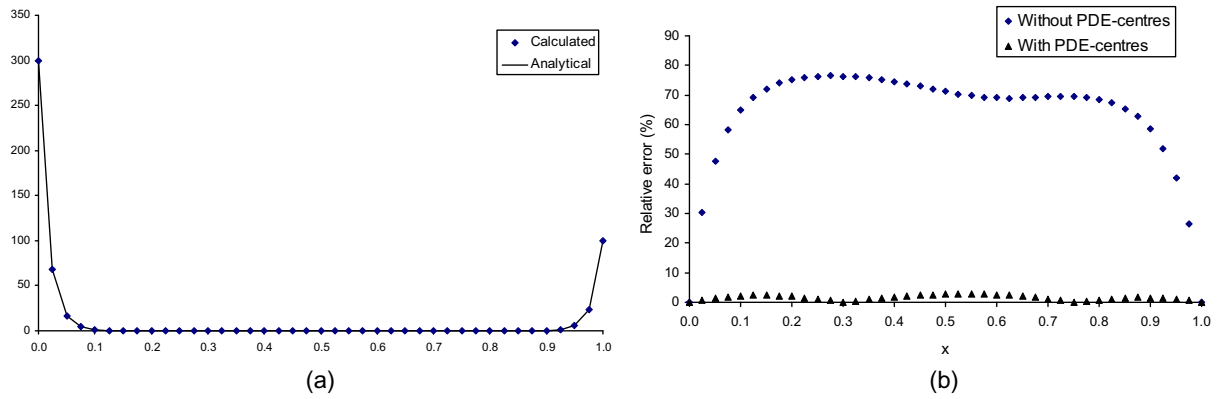

Fig. 12. Solution at $k = 120$ ($\phi_{\min} = 5.3 \times 10^{-5}$): (a) solution profile, (b) relative error profile.

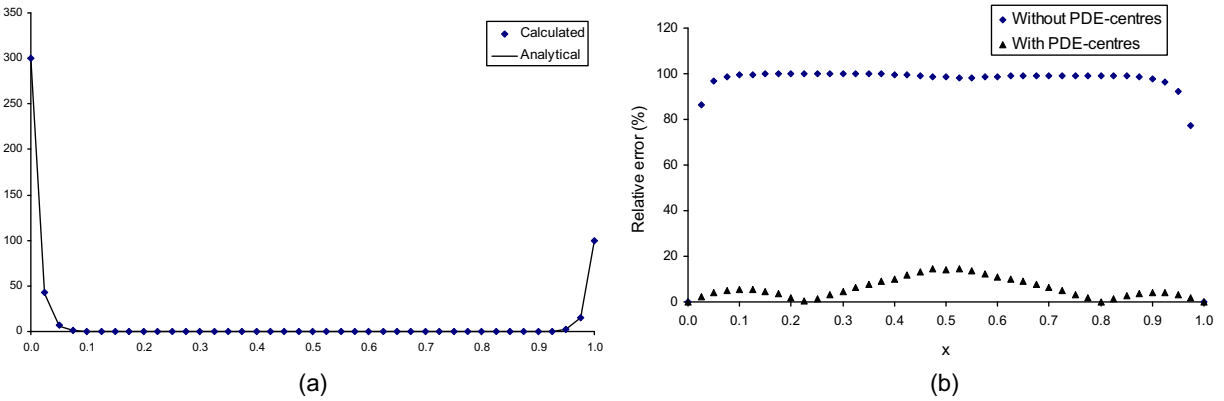

Fig. 13. Solution at $k = 160$ ($\phi_{\min} = 3.57 \times 10^{-7}$): (a) solution profile, (b) relative error profile.

[20] for a similar 1D test case. In the case that heavy refinement is required and local system separations reduce significantly, the local shape-parameter must be adjusted as described by Eq. (21).

7. Three-dimensional transient problem on an unstructured dataset

The unsteady convection-diffusion equation with a homogeneous and anisotropic diffusivity is solved

$$\frac{\partial \phi}{\partial t} = D_{ij}\frac{\partial^2 \phi}{\partial x_i \partial x_j} - u_i \frac{\partial \phi}{\partial x_i} \tag{29}$$

$$\text{where} \quad D_{ij} = \begin{pmatrix} D_1 & 0 & 0 \\ 0 & D_2 & 0 \\ 0 & 0 & D_3 \end{pmatrix} \tag{30}$$

An instantaneous point source of strength M, located at $\hat{x}_i$, is generated at $t = 0$ in an infinite domain. This leads to an analytical solution, given by Eq. (31) (see [36])

$$\phi(x_i, t) = \frac{M}{(4\pi t)^{3/2}(D_1 D_2 D_3)^{1/2}} \exp\left(\frac{[(x_1 - \hat{x}_1) - u_1 t]^2}{4D_1 t} + \frac{[(x_2 - \hat{x}_2) - u_3 t]^2}{4D_2 t} + \frac{[(x_3 - \hat{x}_3) - u_3 t]^2}{4D_3 t}\right) \tag{31}$$

Taking

$$\begin{aligned} &M = 1 \\ &D_1 = 0.1, \quad D_2 = 0.1, \quad D_3 = 0.05 \\ &\hat{\mathbf{x}} = (-0.1, \quad 0.25, \quad 0.125) \\ &\mathbf{u} = (1.0, \quad 0.0, \quad 0.0) \end{aligned} \tag{32}$$

The PDE of Eq. (29) is solved over the domain:

$$0 \leqslant x_1 \leqslant 1.0$$
$$0 \leqslant x_2 \leqslant 0.5$$
$$0 \leqslant x_3 \leqslant 0.25$$

A Dirichlet boundary condition is prescribed at all boundary faces, defined by the potential (31). The problem now represents the time evolution of the singular solution (31) inside the parallelepiped [1, 0.5, 0.25], where in order to avoid the singularity the source is located outside the computational domain, at (−0.1, 0.25, 0.125).

A dataset is generated over the solution domain through the use of a 3D tetrahedral mesh generator (GAMBIT), resulting in an unstructured and irregular distribution of points. The relatively coarse dataset consists of 1002 boundary centres and 1539 solution centres. This distribution, and the solution at $t = 0.5$, are represented in Fig. 14.

The basic local system stencils are created by linking each solution centre to every solution or boundary centre which is directly connected to it within the tetrahedral mesh. Stencil sizes vary, with a modal size of 14 connected data-centres. In addition, two stencils with PDE centres are included. The collocated configuration places PDE centres at every node in the local system, except for the system centrepoint. The staggered configuration places PDE centres at the centre of each cell which surrounds the node in the tetrahedral mesh.

Table 5 shows the L_2 solution error at the solution data-centres for a series of runtimes, created using the $\theta = 0.5$ Crank–Nicolson scheme with $\Delta t = 10^{-3}$ and a shape-parameter $c = 0.15$. The solution quality improves significantly at later runtimes as the solution is smoothed by diffusion. However, at each runtime the solution produced by the stencils using PDE centres is significantly better than that produced without PDE centres, providing at least a factor of ten reduction in the L_2 error norm in most cases. In this test, the collocated PDE stencil performs significantly better than the staggered PDE stencil at times $t = 0.2$ and beyond.

Fig. 15 plots the value of the solution, as reconstructed from the closest local system, at 50 evenly spaced locations along the problem centreline (x, 0.25, 0.125). It is clear that all three stencils are qualitatively tracking the analytical solution, but that in each case the addition of PDE centres improves the solution noticeably.

The results shown here illustrate that the LHI technique is capable of reproducing solutions with a fully 3D problem, on an irregular dataset, with anisotropic diffusivity. The addition of PDE centres continues to offer significant improvement to the solution on the scattered dataset.

8. Three-dimensional steady problem on an unstructured dataset

The final test case examines the performance of the method with a steady 3D problem on an irregular geometry. The steady convection-diffusion-reaction equation is solved with a homogeneous and isotropic diffusivity, an outward radial velocity field, and a spherically symmetric reaction coefficient (Eq. (33)).

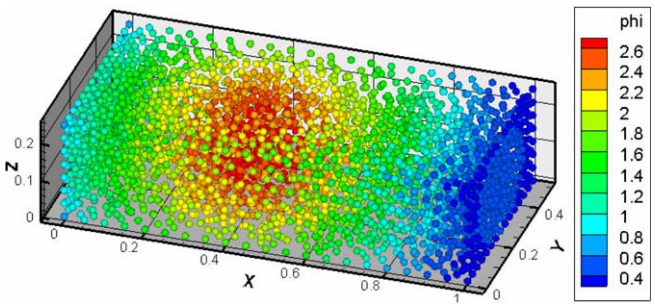

Fig. 14. Representation of dataset, and of solution at $t = 0.5$.

Table 5
L_2 solution error at solution centres for various runtimes.

	$t = 0.1$	$t = 0.2$	$t = 0.5$	$t = 1.0$	$t = 1.5$
No PDE	6.07×10^{-2}	4.99×10^{-2}	1.93×10^{-2}	4.07×10^{-3}	8.24×10^{-4}
PDE collocated	8.82×10^{-3}	3.43×10^{-3}	6.17×10^{-4}	8.19×10^{-5}	1.90×10^{-5}
PDE staggered	5.33×10^{-3}	5.79×10^{-3}	1.90×10^{-3}	3.19×10^{-4}	5.43×10^{-5}

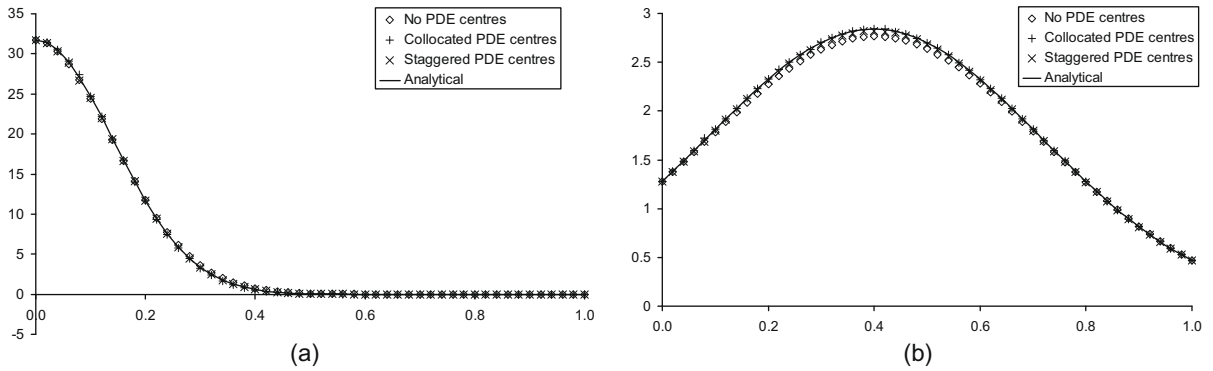

Fig. 15. Reconstructed solution profiles along (x, 0.25, 0.125): (a) $t = 0.1$, (b) $t = 0.5$.

$$\frac{\partial^2 \phi}{\partial x_i^2} - \frac{1}{x_i}\frac{\partial \phi}{\partial x_i} - x_i^2 \phi = 0 \tag{33}$$

Eq. (33) is solved over a domain formed by removing a quarter-sphere from a slightly larger cube. By applying a solution value of zero on the surface of the sphere, an analytical solution can be found (Eq. (34))

$$\phi(x_i) = \sinh\left(\frac{x_i^2 - r_1^2}{2}\right) \tag{34}$$

where r_1 is the radius of the sphere.

The radius of the sphere is chosen as $r_1 = 1.0$, and the cube is chosen to be of edge-length $\sqrt{2}$ (see Fig. 16). Application of the analytical solution over the surface of the geometry completes the problem by providing Dirichlet boundary conditions.

As with the 3D transient problem examined in Section 7, an irregular distribution of data-centres is used. The dataset is generated using an unstructured tetrahedral mesh generator (GAMBIT). In this case the dataset consists of 1560 boundary centres and 1925 solution centres. By using the same procedure as described in Section 7, three stencil configurations are generated from the tetrahedral mesh.

Table 6 gives the L_2 solution error, calculated at the solution centres, for each of the three stencil configurations. In both cases the addition of PDE centres improves the solution quality significantly, however the stencil utilising staggered PDE centres shows the biggest improvements. To provide a qualitative representation of performance, the value of the solution is reconstructed at 50 equidistant intervals along the line $x = y = z$, from the surface of the embedded sphere to the far corner of the domain (see Fig. 17). The closest local system is used for each reconstruction. The stencil configurations utilising PDE centres show a noticeable improvement over the basic interpolation stencil, whereas the two stencils employing PDE centres are virtually indistinguishable.

In this case, an attempt was made to compare the quality and computational cost of the LHI solutions with an equivalent full-domain Hermitian method. However, it was not possible to find a numerical solution with the full-domain Hermitian

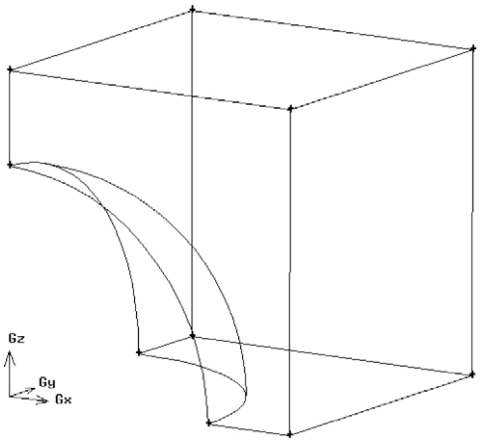

Fig. 16. Representation of the problem geometry.

Table 6
L_2 solution error at solution centres.

Stencil configuration	No PDE	PDE collocated	PDE staggered
L_2 solution error	3.243×10^{-2}	6.244×10^{-3}	2.903×10^{-3}

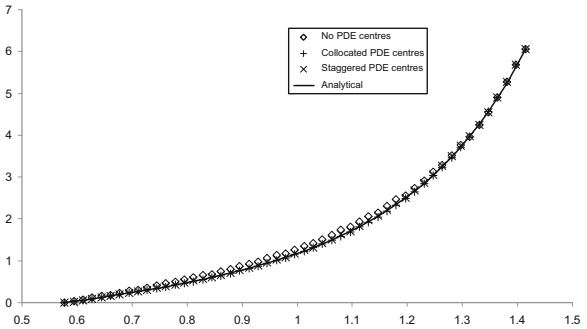

Fig. 17. Reconstructed solution profile along the line $x = y = z$.

method using the same solver, as the non-preconditioned GMRES solver which is utilised for the LHI method would not converge with any tested c-parameter value. In any case, the solution of a globally populated matrix of size 3485, as produced by the full-domain Hermitian method, will be significantly more expensive than the LHI solutions, which complete in around 1 s.

9. Conclusions

In this paper an extension to the meshless Local Hermitian Interpolation (LHI) method has been presented. The existing formulation, based on Hermitian RBF interpolation of solution values and problem boundary operators, has been extended to include interpolation with PDE-operators of the governing equation. Using an interpolation system which itself satisfies the governing equation improves the accuracy of the reconstructed partial derivatives, and in turn leads to more stable and accurate solutions to the PDE. The formulation of the method is described for both steady and transient problems, with a weighted Crank–Nicolson time advancement method utilised for the time stepping scheme. By the using the partial differential operator of the governing equation in the interpolation function, including the desired information about the convective velocity field, a kind of analytical upwinding is implicitly implemented into the numerical scheme.

Six identifying stencil configurations, including collocation of field variable, PDE- and boundary operators, which are suitable for structured datasets have been produced and applied to a benchmark steady and transient problem. In addition, steady and transient 3D problems have been solved on an unstructured dataset using a further three stencil configurations. The addition of PDE centres appears to significantly improve the performance of the LHI method in the four numerical examples tested, often providing an error reduction of a factor of ten when compared to equivalent stencils without PDE centres. The nature of the PDE centres can also have a significant effect on the solution quality. The staggered formulations appear preferable for steady problems, where no reconstruction is required at the PDE centres. However, for time-dependent problems the staggered formulation incurs an additional storage cost, and for small stencil sizes may also lead to instability in the reconstructions which are required at the PDE centres.

Further development of the proposed approach is in progress, in particular dealing with the solution of non-linear problems, as is the case with the Navier–Stokes system of equations, i.e. viscous flow problems. The task of obtaining a local approximation that satisfies the Navier–Stokes equations within the interpolation is a topic of ongoing research by several research groups. In particular, it is not clear how the non-linearity is to be considered at the local level. From previous experience, it is expected that a robust non-linear solver needs to be implemented at the global level, but most likely a simple Picard iteration can be used at the local level. In our previous work dealing with the non-linear Richards equation, it was possible to transform the strong non-linear problem into a weak non-linear problem by the use of the Kirchoff transformation. The results reported in the present article encourage us to progress with development of the proposed approach, in particular when dealing with convective dominant problems. We expect in the near future to be able to solve, with our approach, complex problems such as those governed by the Navier–Stokes system of equations. However, this is not the main objective of the present manuscript, which aims to show that with our approach it is not necessary to define an ad hoc upwinding scheme.

Appendix

See Table A1.

Table A1
List of common terms describing locations at which data is manipulated.

Trial point	Location about which a RBF is centred
Test point	Location at which a known constraint is applied
Data-centre	A location consisting of both a test- and a trial point
Solution centre	A data-centre at which the constraint is the value of the solution field
Boundary centre	A data-centre at which the constraint is the value of the PDE boundary operator
PDE centre	A data-centre at which the constraint is the value of the PDE governing operator
Stencil	A collection of data-centres over which a RBF interpolation can be performed
Local system	A RBF interpolation performed using a particular stencil
Local system centrepoint	The location around which a local system is formed. In the current implementation the local system centrepoint is always a solution centre

References

[1] W. Madych, S. Nelson, Multivariate interpolation and conditionally positive definite functions II, Mathematics of Computation 44 (189) (1990) 211–230.
[2] R. Schaback, Comparisons of radial basis function interpolants, in: Multivariate Approximation: From CAGD to Wavelets, World Scientific, Singapore, 1993.
[3] T.A. Driscoll, B. Fornberg, Interpolation in the limit of increasingly flat radial basis functions, Computers and Mathematics with Applications 43 (2002) 413–422.
[4] B. Fornberg, G. Wright, Stable computation of multiquadric interpolants for all values of the shape parameter, Computers and Mathematics with Applications 48 (2004) 853–867.
[5] B. Fornberg, C. Piret, A stable algorithm for flat radial basis functions on a sphere, SIAM Journal of Scientific Computing 30 (2007) 60–80.
[6] D. Brown, On approximate cardinal preconditioning methods for solving PDEs with radial basis functions, Engineering Analysis with Boundary Elements 29 (4) (2005) 343–353.
[7] L. Ling, R. Opfer, R. Schaback, Results on meshless collocation techniques, Engineering Analysis with Boundary Elements 30 (4) (2006) 247–253.
[8] J. Munoz-Gomez, P. Gonzalez-Casanova, G. Rodriguez-Gomez, Domain decomposition by radial basis functions for time dependent partial differential equations, in: Proceedings of the Second IASTED International Conference on Advances in Computer Science and Technology, Puerto Vallarta, Mexico, 2006.
[9] L. Ling, E.J. Kansa, Preconditioning for radial basis functions with domain decomposition methods, Mathematical and Computer Modelling 38 (5) (2004) 320–327.
[10] A. Hernandez Rosales, H. Power, Non-overlapping domain decomposition algorithm for the Hermite radial basis function meshless collocation approach: applications to convection diffusion problems, Journal of Algorithms and Technology 1 (1) (2007) 127–159.
[11] M. Ingber, C. Chen, J. Tanski, A mesh free approach using radial basis functions and parallel domain decomposition for solving three-dimensional diffusion equations, International Journal for Numerical Methods in Engineering 60 (2004) 2183–2201.
[12] G. Wright, B. Fornberg, Scattered node compact finite difference-type formulas generated from radial basis functions, Journal of Computational Physics 212 (1) (2006) 99–123.
[13] C. Lee, X. Liu, S. Fan, Local multiquadric approximation for solving boundary value problems, Computational Mechanics 30 (2003) 396–409.
[14] A. Tolstykh, D. Shirobokov, On using radial basis functions in a finite difference mode with applications to elasticity problems, Computational Mechanics 33 (2003) 68–79.
[15] C. Shu, H. Ding, K. Yeo, Local radial basis function-based differential quadrature method and its application to solve two-dimensional incompressible Navier–Stokes equations, Computer Methods in Applied Mechanics and Engineering 192 (2003) 941–954.
[16] B. Sarler, R. Vertnik, Meshless explicit local radial basis function collocation methods for diffusion problems, Computers and Mathematics with Applications 51 (2006) 1269–1282.
[17] R. Vertnik, M. Zaloznik, B. Sarler, Solution of transient direct-chill aluminium billet casting problem with simultaneous material and interphase moving boundaries by a meshless method, Engineering Analysis with Boundary Elements 30 (2006) 847–855.
[18] G. Kosec, B. Sarler, Solution of thermo-fluid problems by collocation with local pressure correction, International Journal of Numerical Methods for Heat and Fluid Flow 18 (2008) 868–882.
[19] E. Divo, A. Kassab, An efficient localised radial basis function meshless method for fluid flow and conjugate heat transfer, Journal of Heat Transfer 129 (2007) 124–136.
[20] D. Stevens, H. Power, H. Morvan, An order-*N* complexity meshless algorithm for transport-type PDEs based on local hermitian interpolation, Engineering Analysis with Boundary Elements 33 (4) (2009) 425–441.
[21] D. Stevens et al., A meshless solution technique for the solution of 3D unsaturated zone problems, based on local hermitian interpolation with radial basis functions, Transport in Porous Media, in press, doi:10.1007/s11242.008.9303.z.
[22] R. Kissmann, R. Grauer, A low dissipation essentially non-oscillatory central scheme, Computer Physics Communications 176 (2007) 522–530.
[23] S. Ii, F. Xiao, CIP/multi-moment finite volume method for Euler equations: a semi-Lagrangian characteristic formulation, Journal of Computational Physics 222 (2007) 849–871.
[24] R. Schaback, Convergence of unsymmetric kernal-based meshless collocation methods, SIAM Journal of Numerical Analysis 45 (1) (2007) 333–351.
[25] E.J. Kansa, Multiquadrics – a scattered data approximation scheme with applications to computational fluid-dynamics – I: surface approximations and partial derivatives estimates, Computers and Mathematics with Applications 19 (1990) 127–145.
[26] E.J. Kansa, Multiquadrics – a scattered data approximation scheme with applications to computational fluid dynamics – II: solution to parabolic, hyperbolic and elliptic partial differential equations, Computers and Mathematics with Applications 19 (1990) 147–161.
[27] G.E. Fasshauer, Solving partial differential equations by collocation with radial basis functions, Surface Fitting and Multiresolution Methods (1997).
[28] Z. Wu, Hermite–Birkhoff interpolation of scattered data by radial basis functions, Approximation Theory 8 (2) (1992) 1–11.
[29] Z. Wu, Solving PDEs with radial basis functions and the error estimation, in: Z. Chen et al. (Eds.), Advances in Computational Mathematics, 1998.
[30] Y. Saad, SPARSKIT: A Basic Tool-kit for Sparse Matrix Computations, 1988–2000, <http://www-users.cs.umn.edu/~saad/software/SPARSKIT/sparskit.html>.
[31] T. Cecil, J. Qian, S. Osher, Numerical methods for high dimensional Hamilton–Jacobi equations using radial basis functions, Journal of Computational Physics 196 (2004) 327–347.
[32] R.E. Carlson, T.A. Foley, The parameter R^2 in multiquadric interpolation, Computers and Mathematics with Applications 21 (9) (1991) 29–42.

[33] S. Rippa, An algorithm for selecting a good value for the parameter c in radial basis function interpolation, Advances in Computational Mathematics 11 (1999) 193–210.
[34] P. Orsini et al., An implicit volume element method based on meshless radial basis function techniques for modelling transport phenomena. International Journal of Numerical Methods in Engineering (submitted for publication).
[35] M. Portapila, H. Power, Iterative schemes for the solution of a system of equations arising from the DRM in multi domain approach, and a comparative analysis of the performance of two different radial basis functions used in the interpolation, Engineering Analysis with Boundary Elements 29 (2005) 107–125.
[36] H. Carlslaw, J. Jaeger, Conduction of Heat in Solids, Oxford at the Clarendon press, Oxford, 1959.