

1 Fast three dimensional r-adaptive mesh  
2 redistribution

3 P.A. Browne<sup>1,\*</sup>, C.J. Budd<sup>2</sup>, C. Piccolo<sup>3</sup>, and M. Cullen<sup>3</sup>

4 <sup>1</sup>Department of Meteorology, University of Reading, UK

5 <sup>2</sup>Department of Mathematical Sciences, University of Bath, UK

6 <sup>3</sup>Met Office, Fitzroy Road, Exeter, EX1 3PB, UK

7 \*Correspondence to p.browne@reading.ac.uk

8 August 6, 2013

9 **Abstract**

10 This paper describes a fast and reliable method for redistributing a  
11 computational mesh in three dimensions which can generate a complex  
12 three dimensional mesh without any problems due to mesh tangling. The  
13 method relies on a three dimensional implementation of the parabolic  
14 Monge-Ampère (PMA) technique, for finding an optimally transported  
15 mesh. The method for implementing PMA is described in detail and ap-  
16 plied to both static and dynamic mesh redistribution problems, studying  
17 both the convergence and the computational cost of the algorithm. The  
18 algorithm is applied to a series of problems of increasing complexity. In  
19 particular very regular meshes are generated to resolve real meteorolog-  
20 ical features (derived from a weather forecasting model covering the UK  
21 area) in grids with over  $2 \times 10^7$  degrees of freedom. The PMA method  
22 computes these grids in times commensurate with those required for op-  
23 erational weather forecasting.

24 This work was funded by EPSRC Knowledge Transfer Grant XXX-XXXX-  
25 XXXX.

26 **1 Introduction**

27 **1.1 Overview**

28 Many physical problems exhibit variety of different spatial scales and feature  
29 localised small scale structures embedded within a much larger scale geometry.  
30 Examples include the boundary layers frequently encountered in fluid mechanics  
31 and gas dynamics, meteorological inversion layers [1], weather fronts, combustion  
32 tion layers and shock waves. Computations on such problems using a uniform

1 computational mesh may encounter problems when the computational mesh size  
2 is too large to resolve the small scale structures. When such a computation is  
3 part of a computational fluid dynamics calculation then this may lead to large  
4 truncation errors [2]. In the data assimilation context, an adaptive mesh is a  
5 convenient way of representing spatially varying correlation structures which  
6 would otherwise have to be represented by an unmanageably large correlation  
7 matrix. It is thus often important, both for accuracy and for computational  
8 efficiency, to use a computational mesh which is adapted in some manner to  
9 the small scales in the underlying problem. This is relatively easy in one spa-  
10 tial dimension with many excellent examples of successful implementations both  
11 in PDE calculations [3] and in data assimilation, [4] leading to significant in-  
12 creases in accuracy and computational efficiency. However, the computational  
13 difficulties of (dynamically) adapting a mesh for a three dimensional problem  
14 and coupling it to a solver, are considerable [5]. Furthermore, fully three di-  
15 mensional adapted meshes can take a significant time to generate [6]. In this  
16 paper, we will describe an algorithm for *adaptive mesh redistribution* based on  
17 optimal transport ideas, which is both fast to implement, avoids mesh tangling  
18 and gives excellent three dimensional meshes for some large and challenging  
19 problems. We demonstrate the effectiveness of this procedure on a number of  
20 problems, including large meteorological calculations based on real data. These  
21 methods have the potential for relatively easy coupling to both CFD codes and  
22 data assimilation procedures.

## 23 1.2 An outline of Adaptive mesh redistribution

24 Broadly speaking adaptive meshes fall into three types. The most commonly  
25 used is *Adaptive Mesh Refinement*, AMR or h-adaptivity, in which a structured  
26 mesh is locally refined (or possibly de-refined) by the addition (or subtraction)  
27 of new mesh points [7] when some local refinement condition is satisfied [8]. This  
28 is closely related to p-adaptive methods [9] in which the order of the elements  
29 used in the computation is locally increased, again prompted by some local re-  
30 finement condition. Both of these methods have the advantages of a degree of  
31 maturity in implementation and flexibility of use. However they also suffer from  
32 various disadvantages. The complex and evolving data structures needed to de-  
33 scribe the mesh and its changing connectivity [10] can make it difficult to couple  
34 them to other software. Furthermore the very local nature of the mesh refine-  
35 ment, can lead to meshes with poor global structures, without good alignment  
36 or regularity. An alternative procedure, described in this paper, is *Adaptive*  
37 *Mesh Redistribution*, also known as r-adaptivity (or more simply as a moving  
38 mesh method). In this procedure a *fixed number* of mesh points in a *constant*  
39 *connectivity structure* is redistributed so that the points are optimally placed  
40 to resolve the fine-scale features of interest. A powerful method for doing this  
41 is to move the points so that the *point density* is controlled by equidistributing  
42 an appropriate scalar or matrix *monitor function*. This procedure has certain  
43 similarities to Lagrangian methods in which the velocity of the mesh points is  
44 coupled to convective features of the underlying solution. However, it avoids the

1 mesh tangling problems often associated with such methods <sup>Budd2009a</sup> [11]. Whilst less  
 2 mature than AMR type methods, adaptive mesh redistribution offers potential  
 3 advantages. Firstly, the constant data structure makes them straightforward  
 4 both to use in their own right and to couple to existing software. Secondly,  
 5 the fact that all of the points in the mesh are calculated together means that  
 6 both local refinement and global regularity of the mesh can be treated together,  
 7 leading to potentially very regular meshes. (Indeed it is possible to build a de-  
 8 gree of global regularity directly into the implementation of the method <sup>Budd2009a</sup> [11].)  
 9 Thirdly, the mesh points can inherit underlying dynamical features of the prob-  
 10 lem such as symmetries and self-similarity. Various methods for implementing  
 11 adaptive mesh redistribution of varying levels of complexity include Geometric  
 12 Conservation Law methods, Harmonic maps, and variational methods. See the  
 13 reviews in <sup>Budd2009</sup> [12], <sup>huang2011</sup> and [13]. All of these methods consider adaptivity in at most  
 14 two-dimensions. An alternative method based on Optimal Transport ideas is  
 15 described in <sup>Budd2009,anno2008</sup> [11], [14], and takes a differing approach, coupling equidistribution  
 16 to global mesh regularity and calculating an appropriate scalar *mesh potential*  
 17 from which the mesh can be determined. Optimal transport based methods  
 18 are relatively cheap to implement and have been coupled successfully to com-  
 19 putations of incompressible flows in two-dimensions <sup>Budd2013</sup> [15], and also to large scale  
 20 data assimilation calculations <sup>Piccolo2012, Piccolo2011</sup> [1, 4]. Objections to adaptive mesh redistribution  
 21 methods include the possibilities of mesh tangling and mesh skewness, leading  
 22 to elements with small angles and the loss of balance relationships when rep-  
 23 resenting certain fluid motions. Whilst these objections are often valid, it is  
 24 certainly the case that optimally transported meshes can be computed cheaply,  
 25 even in three dimensions, they have provable regularity <sup>Budd2013</sup> [11],[15], they do not  
 26 suffer from mesh tangling, the reduction in errors due to improved resolution  
 27 can outweigh the extra errors given by mesh skewness, and skewness can also  
 28 be an advantage if it leads to better alignment of the mesh with the underlying  
 29 solution <sup>Cao2005, HR</sup> [16], [7], [7]. Finally the preservation of balance laws can be built into  
 30 the mesh construction through the construction of the monitor function.

31 In this paper we show how the optimal transport method, coupled to a simple  
 32 relaxation approach, can be implemented practically to deal with large three di-  
 33 mensional problems with severe geometric distortion. We then test this method  
 34 on a series of challenging problems including large scale meteorological systems.  
 35 In this implementation the calculation of a three dimensional meteorological grid  
 36 with 21772800 degrees of freedom could be accomplished in five minutes on a  
 37 laptop computer. In principle these meshes can be coupled to data assimilation  
 38 codes using methods of <sup>Piccolo2012, Piccolo2011</sup> [1, 4].

39 The remainder of this paper is structured as follows. In Section 2 we describe  
 40 some of the underlying theory of r-adaptive mesh redistribution and the optimal  
 41 transport method of doing this, leading to a single equation (the Monge-Ampère  
 42 equation) describing the mesh. In Section 3 we describe a relaxation method  
 43 for solving this equation. In Section 4 we describe a simple, practical and effec-  
 44 tive method for discretising this equation and calculating a three dimensional

1 mesh. In Section 5 we consider various static mesh redistribution problems in-  
 2 cluding some which use meteorological data from the Met Office UK4 forecast  
 3 system. Finally in Section 6 we consider an evolving problem with dynamic  
 4 mesh redistribution.


## 5 2 Adaptive mesh redistribution in three dimen- 6 sions

amr

Adaptive mesh redistribution methods work by keeping the number of mesh  
 points and the topology of the mesh fixed but redistribute the mesh in space.  
 For a time evolving problem the mesh can then evolve with the solution of  
 the underlying problem. The simplest three dimensional mesh  $\mathcal{T}_C$  comprises a  
 regular subdivision of the unit cube into identical smaller cubes. We denote the  
 unit cube by  $\Omega_C = [0, 1]^3$ , and it represents a reference or computational space.  
 We can then map the mesh  $\mathcal{T}_C$  into any other logically (or topologically) cuboid  
 mesh  $\mathcal{T}_P$  occupying a *physical* space  $\Omega_P \subset \mathbb{R}^3$ , through the map

$$\mathbf{F}(\cdot, t) : \Omega_C \rightarrow \Omega_P.$$

7 The mesh points in  $\mathcal{T}_P$  are therefore the images of the corners of the cuboids in  
 8  $\mathcal{T}_C$  and these points redistribute as the time  $t$  evolves. For clarity we define a  
 9 point in  $\Omega_C$  by  $\xi \in \Omega_C = (\xi, \eta, \zeta)$ . Similarly we denote a point  $\mathbf{x}$  in the physical  
 10 space  $\Omega_P$  by  $\mathbf{x} \in \Omega_P = (x, y, z)$ . An example of a section of mesh  $\mathcal{T}_C$  in  $\Omega_C$  and  
 11 a section of its image  $\mathcal{T}_P$  in  $\Omega_P$  is given in Figure 1.



(a) A mesh  $\mathcal{T}_C$  in computational space  $\Omega_C$ , denoted  $\xi = (\xi, \eta, \zeta)$       (b) A mesh  $\mathcal{T}_P$  in physical space  $\Omega_P$ , denoted  $x = (x, y, z)$

radaptexample

Figure 1: A mesh  $\mathcal{T}_C \in \Omega_C$  and its image  $\mathcal{T}_P \in \Omega_P$ .

12 For redistribution to be effective we need to concentrate mesh points so that  
 13 they have a high density in certain regions of  $\Omega_P$ . The value of this mesh density  
 14 is taken to be proportional to the size of a monitor function  $m(\mathbf{x}, t) > 0$ , so that

1 if  $|J(\xi, t)|$  is the determinant of the Jacobian of the map from  $\Omega_C$  to  $\Omega_P$  given  
 2 (in 3 dimensions) by

$$|J(\xi, t)| = \begin{vmatrix} x_\xi & x_\eta & x_\zeta \\ y_\xi & y_\eta & y_\zeta \\ z_\xi & z_\eta & z_\zeta \end{vmatrix} \quad (1)$$

3 then

$$m(\mathbf{x}, t) |J(\xi, t)| = \int_{\Omega_p} m(\mathbf{x}, t) \, \mathbf{d}x. \quad (2) \quad \boxed{\text{equi1}}$$

4 We call this the *equidistribution equation*. In one dimension the equidistribution  
 5 equation uniquely defines the map  $\mathbf{F}$  and a number of methods exploit this, most  
 6 particularly the moving mesh PDE methods listed in [17]. In higher dimensions  
 7 additional conditions are required to define the map uniquely. Noting that for  
 8 many computations there are significant advantages to using a uniform mesh, it  
 9 makes initial sense to look for meshes which are close to being uniform in some  
 10 sense. In other words we seek functions  $\mathbf{F}$  which are close to the identity in  
 11 some measure. A convenient such measure is the Wasserstein metric  $I$  given by

$$I = \int_{\Omega_C} |\mathbf{F}(\xi, t) - \xi|^2 \, d\xi \quad (3) \quad \boxed{\text{optim}}$$

12 **Definition 1.** A map  $\mathbf{F}$  which minimises  $I$  is over all invertible mappings  
 13 satisfying (2) called an *optimally transported map*. The resulting mesh  $\mathcal{T}_P$  is an  
 14 *optimally transported mesh*.

15 Finding such a map is an example of a *Monge-Kantorovich problem* (see [18]).  
 16 Although the condition of minimising  $I$  appears to be a coarse global restraint  
 17 on the mesh  $\mathcal{T}_P$ , it not only leads to a system which is easy to calculate, but also  
 18 to meshes with provably excellent regularity, good mesh grading and good mesh  
 19 alignment [11], [15]. We now seek to solve the Monge-Kantorovich problem to  
 20 determine the optimal mesh  $\mathcal{T}_P$ . The key underlying result which allows us to  
 21 compute this mesh is the following

22 **Theorem 1** (Brenier [18]). *There exists a **unique** optimally transported map*  
 23  *$\mathbf{F}(\xi, t)$  which minimises  $I$ , and the Jacobian of which satisfies the equidistribu-*  
 24 *tion equation (2). This map has the same regularity as the monitor function*  
 25  *$m$ . Furthermore,  $\mathbf{F}(\xi, t)$  can be written as the gradient (with respect to  $\xi$ ) of a*  
 26 *convex scalar (mesh) potential  $P(\xi, t)$ , so that*

$$(x, y, z) \equiv \mathbf{x}(\xi, t) = \nabla_\xi P(\xi, t), \quad H_\xi(P(\xi, t)) \succ 0. \quad (4) \quad \boxed{\text{Peqn}}$$

27 Finding the (three dimensional) map  $\mathbf{F}$  and the associated mesh  $\mathcal{T}_P$  is thus  
 28 reduced to the simpler problem of finding the scalar mesh potential  $P$ . As  
 29  $\mathbf{x} = \nabla P$  it follows immediately that  $J(\xi) = H(P)$  where  $H(P)$  is the Hessian  
 30 matrix of  $P$ . Hence the Jacobian  $J(\xi)$  is a *symmetric matrix* which imposes  
 31 certain restrictions on  $\mathbf{F}$ . For example it cannot be a plane rotation. Such

1 maps are called *Legendre Transformations* and play an important role in many  
 2 fields including fluid mechanics and image processing [19]. In 3-dimensions the  
 3 determinant of the Hessian of  $P$  is given by

$$|H(P)| = \begin{vmatrix} P_{\xi\xi} & P_{\xi\eta} & P_{\xi\zeta} \\ P_{\eta\xi} & P_{\eta\eta} & P_{\eta\zeta} \\ P_{\zeta\xi} & P_{\zeta\eta} & P_{\zeta\zeta} \end{vmatrix}. \quad (5)$$

4 The equidistribution equation (2) then becomes the following equation for  $P$ :

$$m(\nabla_{\xi} P, t) |H(P)| = \int_{\Omega_P} m \, \mathbf{d}x. \quad (6) \quad \boxed{\text{mongeampere}}$$

which is a Monge-Ampère equation. To fully specify the mesh we need to impose  
 boundary conditions on  $P$ . Typically we require that the boundary  $\Gamma_C$  of  $\Omega_C$   
 is mapped to the boundary  $\Gamma_P$  of  $\Omega_P$ . If the latter is given implicitly by the  
 condition

$$\Gamma_P = \{(x, y, z) : G(x, y, z) = 0\}$$

5 then we have the nonlinear Neumann boundary condition

$$G(\nabla_{\xi} P) = 0 \quad \text{if } \xi \in \Gamma_P. \quad (7) \quad \boxed{\text{nIBC}}$$

6 Observe that this procedure allocated points to the boundary, but does not  
 7 prescribe their precise location. If  $\Omega_P$  is a cuboid domains so that, for example,  
 8 one face of  $\Omega_P$  is given by the plane  $x = 0$ , then the nonlinear condition (7)  
 9 simplifies to the simpler linear Neumann condition

$$P_{\xi} = 0. \quad (8) \quad \boxed{\text{lbc}}$$

10 For certain problems, for example a number of problems in meteorology, it is  
 11 natural and convenient to use periodic boundary conditions instead. See [15] <sup>Budd2013</sup>  
 12 for an example.

### 13 3 Parabolic Monge-Ampère formulation

14 Equation (6) is a fully non-linear elliptic PDE which is challenging to solve.  
 15 There is a significant literature describing various solution techniques both for  
 16 the equation in its own right [?], <sup>SFU</sup> as part of a meteorological calculation [20, 21] <sup>cullen2006, Chynoweth1991</sup>  
 17 and as part of a mesh generation algorithm [6], [14] <sup>Chapen2008</sup>. Typically these methods use  
 18 a careful finite difference or finite element discretisation of (6) which is then  
 19 solved using a Newton-type algorithm. In [6] <sup>Chacon2011</sup> the resulting linear equations are  
 20 in turn solved using a multi-grid method. These methods are necessarily com-  
 21 putationally expensive and use significant computer time to give an accurate  
 22 answer. However, in the context of mesh generation, we do not want to invest  
 23 much effort in solving (6) as its function is to generate a mesh which is then used  
 24 for other calculations. In this context an accurate solution of (6) is unnecessary,

1 provided that the resulting mesh is sufficiently regular and aligned, and exhibits  
 2 the correct compression properties that we desire. Accordingly, a simple, ex-  
 3 plicit, relaxation method for solving (6) has certain advantages in this context.  
 4 Firstly it can be implemented cheaply using a Forward Euler method, and the  
 5 mesh calculated rapidly. Secondly, the relaxation method can be terminated  
 6 at any time when the mesh generated is sufficient for subsequent computations  
 7 rather than when we have an exact solution of (6). This gives a very signifi-  
 8 cant speed increase in the mesh generation algorithm (with times reduced from  
 9 hours to minutes). In two-dimensions it has been demonstrated [11], [15], that  
 10 a parabolic relaxation of the Monge-Ampère equation, the Parabolic Monge-  
 11 Ampère equation (PMA), is effective for generating meshes. We now extend  
 12 this method to higher dimensions and demonstrate that it continues to be effec-  
 13 tive as a mesh generator. In this formulation we initially consider the true time  
 14  $t$  to be fixed during the computation of the mesh, and introduce a *pseudo-time*  
 15  $\tau \in [0, \infty)$  and a corresponding pseudo-time dependent function  $Q(\xi, \tau)$  so that  
 16  $\nabla_\xi Q \rightarrow \nabla P$  as  $\tau \rightarrow \infty$  where  $P$  solves (6).

17 **Definition 2** (PMA). The Parabolic Monge-Ampère equation in  $d$ -dimensions  
 18 is defined by

$$LQ_\tau \equiv (I - \gamma \Delta_\xi)Q_\tau = (\hat{m}(\nabla_\xi Q)|H(Q)|)^{\frac{1}{d}} \quad (9) \quad \boxed{\text{PMA}}$$

19 where  $\gamma$  is a scalar parameter defining the amount of smoothing applied. The  
 20 function  $\hat{m}$  is a filtered version of the monitor  $m$  obtained by averaging  $m$   
 21 over several mesh points. (The necessity for such filtering for data assimilation  
 22 problems is carefully illustrated in [1].)

23 In this equation the application of  $L^{-1}$  acts as a smoothing operator (described  
 24 first in [22]) which leads to more regular meshes. Furthermore the action of  
 25  $L^{-1}$  on the discrete form of the right hand side of (9) acts to damp out certain  
 26 (mesh dependent) chequer-board instabilities [23]. It can be rapidly calculated  
 27 for cuboid domains by using the FFT or the Fast Cosine Transform (depending  
 28 upon whether we have periodic or Neumann boundary conditions). The oper-  
 29 ator  $(H(Q))^{1/d}$  is used on the RHS (instead of  $H(Q)$ ) as it has the property  
 30 that  $(H(\lambda Q))^{1/d} = \lambda(H(Q))^{1/d}$ . Thus both sides of (9) scale linearly. This  
 31 is useful both to ensure global existence of the solutions of (9) and to give it  
 32 certain desirable scaling properties [11]. It is further shown in [11] that the  
 33 equation (9) is *locally stable* so that, if  $\nabla Q$  is sufficiently close to  $\nabla P$  then  
 34  $\nabla_\xi Q \rightarrow \nabla P_\xi$  as  $\tau \rightarrow \infty$ . with standard linear convergence. Furthermore,  
 35 during the evolution of (9) both  $H(P)$  and  $\nabla^2 Q$  are bounded away from zero.  
 36 This prevents mesh tangling provided that the equation (9) has a sufficiently  
 37 fine discretisation [11].

38 The evolutionary system (9) is subject to the same boundary conditions as (6).  
 39 It is convenient when solving the PMA equation, especially when using periodic  
 40 boundary conditions, to consider instead of  $Q$  the difference between it and the  
 41 function  $|\xi|^2/2$ . Consider the displacement of the periodic potential,  $\tilde{Q}$ , such

1 that

$$\tilde{Q} = Q - \frac{|\xi|^2}{2}. \quad (10)$$

2 This gives

$$\nabla_\xi \tilde{Q} = \nabla_\xi Q - \xi \quad (11)$$

3 and hence

$$\mathbf{x} = \nabla_\xi \tilde{Q} + \xi \quad (12)$$

4 as  $\mathbf{x} = \nabla_\xi Q$ . The PMA equation can then be rewritten as

$$(I - \gamma \Delta_\xi) \tilde{Q}_\tau = (\hat{m}(\nabla_\xi \tilde{Q} + \xi) |I + H(\tilde{Q})|)^{\frac{1}{d}} \quad (13) \quad \boxed{\tilde{Q}}$$

5 In the absence of a better initial guess, we use the initial conditions for (13)  
6  $\tilde{Q}(0) = 0$ . In the case of a dynamically evolving monitor function, it is sub-  
7 stantially more efficient to evolve  $\tilde{Q}$  starting from the most recently computed  
8 value of  $\tilde{Q}$ . If the monitor function  $\hat{m}$  is known then a corresponding mesh can  
9 be found by evolving (13) in time, either until a steady state is reached or until  
10 the resulting mesh is sufficient, in compression and regularity, for solving any  
11 coupled PDE or data assimilation problem. This latter option results in very  
12 significant time savings.

13 If the mesh is used to solve a time dependent PDE then the monitor function  
14  $m$  will evolve in the true time  $t$ . In this case the mesh is evolved in the pseudo-  
15 time until it is adapted to the solution of the PDE. The solution of the PDE  
16 is then interpolated onto the new mesh. The true time is then advanced by  
17 an appropriate amount and the new solution to the PDE, and hence the new  
18 value of  $m$  is calculated. The process of finding the new mesh by evolution in  
19 pseudo-time is then repeated. We now consider the practical issues with solving  
20 (13) forwards in pseudo-time on the assumption that the monitor function is  
21 known a-priori. In our examples we will consider cases where  $m$  is fixed and  
22 also where  $m$  evolves in time.

## 23 4 Implementation

24 When implementing (13) to find  $\tilde{Q}$  and hence the mesh, it is essential that the  
25 algorithm used is fast and robust as it will typically be part of a much larger  
26 solution process. For example, the UK4 model, a model with 4km resolution  
27 over the UK used by the Met Office for both numerical weather prediction and  
28 for data assimilation, has dimension  $288 \times 360 \times 70 = 7257600$  grid points. Each  
29 of these has 3 degrees of freedom (latitudinal, longitudinal and vertical) and  
30 each degree of freedom is stored in double precision and thus requires 8 bytes of  
31 storage. Hence to store one grid requires  $288 \times 360 \times 70 \times 3 \times 8 = 174182400$  bytes  
32  $= 166.11$  MB. This shows the scale of the problem we are considering and why an  
33 efficient implementation of the algorithm to redistribute the mesh is essential.  
34 However, for mesh generation it need not be especially accurate.



1 Accordingly when calculating  $\tilde{Q}$ , we seek an explicit method where possible, for  
 2 both time and memory considerations. One such method uses a forward Euler  
 3 discretisation of (9) to evolve  $\tilde{Q}$  so that

$$\tilde{Q}(\tau + \delta\tau) = \tilde{Q}(\tau) + \delta\tau\tilde{Q}_\tau(\tau) \quad (14) \quad \boxed{\text{model\_evo}}$$

4 where  $\tilde{Q}_\tau(\tau)$  is given by

$$\tilde{Q}_\tau = L^{-1}(\hat{m}(\nabla_\xi \tilde{Q} + \xi)|I + H(\tilde{Q})|)^{\frac{1}{d}}. \quad (15) \quad \boxed{\text{mesh\_evo}}$$

5 To compute the RHS of (15) we discretise the Hessian operator in (15). This can  
 6 be done most simply by using a finite difference scheme in the computational  
 7 space  $\Omega_C$ . We assume that  $\Omega_C$  is divided into regular cuboids with the values of  
 8  $\tilde{Q}$  given at the vertices of the cuboid. The location  $(x, y, z)$  of the mesh in the  
 9 physical space  $\Omega_P$  at these vertices can then be recovered from  $\tilde{Q}$  by taking a  
 10 discrete gradient (most simply by using central differences). The  $d$ -dimensional  
 11 mesh can then be stored as  $d$   $d$ -dimensional arrays, each containing one of the  
 12 degrees of freedom of the mesh. So in a 2-dimensional case, with  $n_x$  grid points  
 13 in the  $x$ -direction and  $n_y$  grid points in the  $y$ -direction, the mesh is stored as 2  
 14  $n_x \times n_y$  arrays. The first of which contains the  $x$  coordinates of the grid and  
 15 the second containing the  $y$  coordinates. Similarly in the three dimensional case  
 16 there are 3 arrays,  $x$ ,  $y$  and  $z$ , each of size  $n_x \times n_y \times n_z$  where  $n_z$  is the number  
 17 of grid points in the  $z$ -direction. The connectivity of the grid is then implicitly  
 18 defined by the relationship within the  $d$ -dimensional array. Algorithms 1 and  
 19 2 outline the steps taken to find a solution of the Monge-Ampère equation (6)  
 20 and determine the corresponding mesh in the static and dynamic situations  
 21 respectively. Due to memory constraints for the meteorological test problem,  
 22 these algorithms to solve the PMA equation were implemented in Fortran95.

---

**Algorithm 1** The PMA algorithm in 3D for a static monitor function
 

---

- 1: Read initial mesh  $\xi = (\xi, \eta, \zeta)$
- 2:  $\tau \leftarrow 0$
- 3: Initialise  $\tilde{Q}(\tau) = \tilde{Q}_0$
- 4: Store the grid  $\mathbf{x}(\tau) = (x(\tau), y(\tau), z(\tau))$  as

$$x(\tau) \leftarrow \xi + \frac{\partial \tilde{Q}(\tau)}{\partial \xi}, \quad y(\tau) \leftarrow \eta + \frac{\partial \tilde{Q}(\tau)}{\partial \eta}, \quad z(\tau) \leftarrow \zeta + \frac{\partial \tilde{Q}(\tau)}{\partial \zeta}.$$

- 5: **while**  $r > tol$  &  $\tau < \tau_{\max}$  **do**
- 6:   Compute  $\tilde{Q}_\tau(\tau)$  via:
  - Compute the monitor function at the current grid points  $m(\mathbf{x}(\tau))$ . This may be analytically defined or interpolated from a given data set
  - Filter the monitor function

$$\hat{m}(\mathbf{x}(\tau)) \leftarrow m(\mathbf{x}(\tau))$$

- Compute the second derivatives of  $\tilde{Q}(\tau)$  in the computational space by using via finite differences to give discrete approximations to:

$$\tilde{Q}_{\xi\xi}(\tau), \tilde{Q}_{\eta\eta}(\tau), \tilde{Q}_{\zeta\zeta}(\tau), \tilde{Q}_{\xi\eta}(\tau), \tilde{Q}_{\xi\zeta}(\tau), \tilde{Q}_{\eta\zeta}(\tau)$$

- Calculate the determinant,  $\rho(\tau)$ , of the Hessian of the mesh potential  $\tilde{Q}(\tau)$  at every current grid point:

$$\rho(\tau) \leftarrow |I + H(\tilde{Q}(\tau))|$$

- Calculate the smoothing operator  $L^{-1}$  by applying the Fast Cosine Transform to the 3-dimensional array  $(\hat{m}(\mathbf{x}(\tau))\rho(\tau))^{\frac{1}{3}}$ , so

$$\tilde{Q}_\tau \leftarrow L^{-1}(\hat{m}(\mathbf{x}(\tau))\rho(\tau))^{\frac{1}{3}}$$

- 7:   Take a Forward Euler step

$$\tilde{Q}(\tau + \delta\tau) = \tilde{Q}(\tau) + \delta\tau \tilde{Q}_\tau(\tau)$$

- 8:   Compute the finite difference approximations to  $\frac{\partial \tilde{Q}(\tau)}{\partial \xi}$ ,  $\frac{\partial \tilde{Q}(\tau)}{\partial \eta}$  and  $\frac{\partial \tilde{Q}(\tau)}{\partial \zeta}$
- 9:   Store the new grid as

$$x(\tau) \leftarrow \xi + \frac{\partial \tilde{Q}(\tau)}{\partial \xi}, \quad y(\tau) \leftarrow \eta + \frac{\partial \tilde{Q}(\tau)}{\partial \eta}, \quad z(\tau) \leftarrow \zeta + \frac{\partial \tilde{Q}(\tau)}{\partial \zeta}$$

- 10:   Compute the change in the mesh

$$r \leftarrow \|\nabla_\xi \tilde{Q}(\tau + \delta\tau) - \nabla_\xi \tilde{Q}(\tau)\|_2$$

- 11:    $\tau \leftarrow \tau + \delta\tau$
  - 12: **end while**
-

- 1 For completeness we now augment Algorithm 1 with an outer loop that is applied  
 2 in the case of a time-dependent monitor function, giving Algorithm 2.

PMAalgo\_dynamic

---

**Algorithm 2** The PMA algorithm in 3D for a dynamic monitor function

---

- 1:  $t \leftarrow 0$   
 2: Apply Algorithm 1 with  $m(\mathbf{x}) = m(\mathbf{x}, 0)$ ,  $\tilde{Q}_0 \equiv 0$  and  $\tau_{\max} = \infty$   
 3: **while**  $t < t_{\max}$  **do**  
 4:   Apply Algorithm 1 with  $m(\mathbf{x}) = m(\mathbf{x}, t)$  and the initial potential  $\tilde{Q}_0$   
    given by the final value of  $\tilde{Q}(\tau)$  from the previous iteration of Algorithm 1  
 5:    $t \leftarrow t + \delta t$   
 6: **end while**
- 

Now we elaborate on the details of the algorithms to show how the PMA method can be implemented in practice in 3 dimensions for a problem in which a cuboid region  $\Omega_C$  of dimensions  $[0, 1]^3$  is mapped to a corresponding cuboid region  $\Omega_P$  of dimensions  $[0, 1]^3$ . As described in Section 2 this leads to a problem with Neumann boundary conditions of the form

$$\tilde{Q}_\xi(0, \dots) = \tilde{Q}_\xi(1, \dots) = \tilde{Q}_\eta(\dots, 0, \dots) = \tilde{Q}_\eta(\dots, 1, \dots) = \tilde{Q}_\zeta(\dots, \dots, 0) = \tilde{Q}_\zeta(\dots, \dots, 1) = 0.$$

- 3 For this implementation we assume that  $\Omega_C$  has a regular cubic mesh with,  
 4 respectively,  $n_x, n_y$  and  $n_z$  cubes in the the three coordinate directions, of cor-  
 5 responding side lengths  $h_x, h_y$  and  $h_z$ .

#### 6 4.1 First order differentiation

- 7 With the mesh potential  $Q$  stored in an  $d$ -dimensional ordered array, comput-  
 8 ing the first order derivatives is straight forward to implement using a central  
 9 differencing scheme. So for instance in the 3 dimensional case, the derivative  
 10 with respect to  $\xi$  is given by

$$\tilde{Q}_\xi(j, :, :) = \frac{\tilde{Q}(j+1, :, :) - \tilde{Q}(j-1, :, :)}{2h_x}, \quad j = 2 : n_x - 1$$

- 11 At the boundaries we invoke the Neumann boundary conditions so that

$$\tilde{Q}_\xi(1, :, :) = \tilde{Q}_\xi(n_x, :, :) = 0.$$

- 12 Derivatives with respect to other variables follow similarly.

#### 13 4.2 Second order differentiation

- 14 Firstly, we note that the enforcing the Neumann boundary conditions implies  
 15 that all mixed derivatives on the boundary are zero, i.e. on the boundary

$$\tilde{Q}_{\xi\eta} = \tilde{Q}_{\xi\zeta} = \tilde{Q}_{\eta\zeta} = 0.$$

1 For the non-mixed derivatives on the boundary, a simple forward (or backward)  
 2 scheme is used, so that for example

$$\tilde{Q}_{\eta\eta}(:, 1, :) = \frac{\tilde{Q}(:, 1, :) - 2\tilde{Q}(:, 2, :) + \tilde{Q}(:, 3, :)}{4h_y^2}.$$

3 In the interior of the domain, central differencing is employed such that

$$\tilde{Q}_{\eta\eta}(:, j, :) = \frac{\tilde{Q}(:, j+1, :) - 2\tilde{Q}(:, j, :) + \tilde{Q}(:, j-1, :)}{h_y^2}, \quad j = 2 : n_y - 1$$

and similarly for mixed second derivatives away from the boundary, so that for  
 example

$$\tilde{Q}_{\xi\xi}(i, :, k) = \frac{1}{4h_x h_z} (\tilde{Q}(i+1, :, k) - \tilde{Q}(i-1, :, k) - \tilde{Q}(i+1, :, k-1) + \tilde{Q}(i-1, :, k-1))$$

4 for all  $i \in \{2, \dots, n_x - 1\}$  and  $k \in \{2, \dots, n_z - 1\}$ .

5 Similar approximations can be used for the other second order derivatives of  $\tilde{Q}$ .

### 6 4.3 Filtering of the monitor function

7 As described above, some form of filtering of the monitor function is required  
 8 in practice Piccolo2009a [1], [11] to produce sufficiently smooth meshes in a reasonable time.  
 9 This is typically achieved in numerical weather prediction and other similar ap-  
 10 plications by applying an appropriate low pass filter Budd2009a [11] to the monitor function  
 11  $m$ . For a three dimensional isotropic problem this most conveniently can take  
 12 the form:

$$\hat{m}(i, j, k) = \frac{\sum_{\ell_1=-1}^1 \sum_{\ell_2=-1}^1 \sum_{\ell_3=-1}^1 m(i + \ell_1, j + \ell_2, k + \ell_3) \beta^{|\ell_1|+|\ell_2|+|\ell_3|}}{\sum_{\ell_1=-1}^1 \sum_{\ell_2=-1}^1 \sum_{\ell_3=-1}^1 \beta^{|\ell_1|+|\ell_2|+|\ell_3|}}. \quad (16)$$

13 Here  $\beta$  is a smoothing parameter such that  $\beta \in [0, 1]$ . However, this type of  
 14 filtering of the monitor function is not suitable for highly anisotropic cases, for  
 15 example the highly stratified flows treated in the data assimilation application  
 16 of Piccolo2012 [1]. However, filtering only within horizontal atmospheric layers retains this  
 17 stratified structure Piccolo2012 [1]. Thus a filtering operator that is more suitable for data  
 18 assimilation contexts is as follows:

$$\hat{m}(i, j, k) = \frac{\sum_{\ell_1=-1}^1 \sum_{\ell_2=-1}^1 m(i + \ell_1, j + \ell_2, k) \beta^{|\ell_1|+|\ell_2|}}{\sum_{\ell_1=-1}^1 \sum_{\ell_2=-1}^1 \beta^{|\ell_1|+|\ell_2|}} \quad (17)$$

19 This produces much sharper monitor functions and hence gives better refinement  
 20 of the grid around the structures of interest. With real data this filtering has to  
 21 be applied several times in order to get a monitor function which will produce  
 22 a grid with sufficient regularity.

## 1 4.4 Applying the smoothing operator $L^{-1}$

2 For the solution of PMA on a domain with purely Neumann boundary condi-  
 3 tions, the Fast Cosine Transform can be employed to calculate  $L^{-1}$  and hence  
 4 to apply the smoothing operator of the left hand side of the PMA equation (9)  
 5 in  $\mathcal{O}(N \log(N))$  operations. In an  $d$ -dimensional problem this transform has to  
 6 be applied  $d$  times; once along each dimension of the mesh. The freely avail-  
 7 able software FFTW [24] was used to apply the Fast Cosine transform as it  
 8 has the ability to work on multidimensional arrays *in-place*. That is to say the  
 9 data structures do not need to be manually altered to perform a Fast Cosine  
 10 Transform along different dimensions. In the 3-dimensional case, the routine  
 11 `dfftw_plan_r2r_3d` is used with the option `FFTW_REDFT10` along each dimen-  
 12 sion to signify the forward fast cosine transform. When the forward transform  
 13 has been applied, the transformed variable is multiplied by the factor

$$1/(1 + \gamma(k_x^2 + k_y^2 + k_z^2)). \quad (18) \quad \boxed{\text{fftfactor}}$$

14 where the frequency-space coefficients  $k_x$ ,  $k_y$  and  $k_z$  are 3D vector fields given  
 15 by

$$k_x(i, j, k) = \frac{i-1}{n_x-1} \pi n_x, \quad k_y(i, j, k) = \frac{j-1}{n_y-1} \pi n_y, \quad \& \quad k_z(i, j, k) = \frac{k-1}{n_z-1} \pi n_z$$

16 for all  $i \in \{1, \dots, n_x\}$ ,  $j \in \{1, \dots, n_y\}$  and  $k \in \{1, \dots, n_z\}$ . Then the inverse  
 17 Fast Cosine Transform is applied via `dfftw_plan_r2r_3d` used with the option  
 18 `FFTW_REDFT01` along each dimension. This whole operation is equivalent to  
 19 applying the operator  $(I - \gamma \Delta)^{-1}$  and can be seen to explicitly damp the higher  
 20 order frequency components in the mesh, such as the potential chequer-board  
 21 modes which can arise in the discretisation of the Hessian operator.

## 22 4.5 Choice of the tuning parameters

When applying the PMA algorithm we must make decisions on how rapidly  
 the mesh must be updated, the degree of convergence at each iteration, and  
 the degree of smoothing which must be applied. This requires us to initially  
 determine appropriate values for the three parameters used in the static case  
 (Algorithm 1), namely  $\delta\tau$  and  $\gamma$ . In this static case, if the time-step  $\delta\tau$   
 is too large, the Hessian matrix  $H$  will typically become indefinite, leading to  
 mesh crossing and other undesirable features. If it is too small then the system  
 becomes overly stiff. This parameter can be controlled adaptively, however it is  
 generally robust to being set at a small constant value. Noting that the intrinsic  
 time-scale of this system is given by  $m^{-1/d}$  a robust choice is to take

$$\delta\tau = \epsilon m^{-1/d}$$

23 where  $\epsilon$  is a small constant value typically in the range  $0.1 \leq \epsilon \leq 1$ . (This choice  
 24 also has certain useful features when scaling symmetries act on the system [12].) Budd2009

1 The parameter  $\gamma$  appears in the smoothing operator  $L \equiv (I - \gamma \Delta_\xi)^{-1}$  as part  
 2 of equation (15) and is applied in (18). Larger values of  $\gamma$  correspond to higher  
 3 smoothing of the calculated mesh. Typically we have found that the smaller the  
 4 value of  $\gamma$ , the faster that PMA converges to an equidistributed mesh. However  
 5 with  $\gamma$  too small mesh tangling can occur. Hence once the step length for the  
 6 Euler method ( $\delta\tau$ ) has been chosen above then  $\gamma$  is chosen to balance the speed  
 7 of convergence with the robustness of the method. Values of  $\gamma$  in the range  
 8  $\gamma \in [0.1, 0.6]$  are typical and, as above, these could be set adaptively for best  
 9 performance.

10 In the case of a dynamically evolving monitor function,  $\delta t$  corresponds to the  
 11 natural time-scale of the model (ie. the underlying solution of the PDE). If  
 12 the PDE is calculated numerically then it is sensible (and usual) to take  $\delta t$  to  
 13 be the same as the time-step used to evolve the solution of the PDE, although  
 14 occasionally we might interpolate the value of  $m$  between time steps allowing us  
 15 to use values of  $\delta t$  which are smaller than the time-step in the method. When the  
 16 initial redistributed mesh has been found in step 2 of Algorithm 2, it is desirable  
 17 that the mesh is updated more rapidly than the solution of the underlying PDE,  
 18 so that it can track it effectively, but not much more rapidly, so that we are not  
 19 working too hard to calculate the mesh. For the inner loop of Algorithm 2 (step  
 20 4), a value of  $\delta\tau = 0.1 \delta t$  is appropriate for many applications. In the inner loop  
 21 of Algorithm 2 it is not always necessary to run the pseudotime iterations for a  
 22 long time, as the mesh remains close to equidistribution provided  $\delta t$  is not too  
 23 large. Instead we set  $\tau_{\max} = \delta t$  and take  $K$  iterations of the inner inner loop  
 24 with time-step of  $\delta\tau = \delta t/K$ . In correspondence with the above, a typical value  
 25 of  $K$  may be in the range  $[1, 10]$ , with larger values necessary if the difference  
 26  $||m(\mathbf{x}, t + \delta t) - m(\mathbf{x}, t)||$  is large.

## 27 5 Static mesh results

28 We now present a series of examples chosen to demonstrate the performance of  
 29 the PMA algorithm on various challenging problems. In particular the examples  
 30 are chosen to investigate the correspondence of the symmetry and regularity  
 31 of the mesh to that of the underlying monitor function, to demonstrate the  
 32 avoidance of mesh tangling when calculating the meshes in three dimensions  
 33 and also to show that the PMA algorithm can cope with very large problems for  
 34 which the monitor function is defined only at data points. In this section results  
 35 are presented for a series of time invariant test problems in which  $m(\mathbf{x}, t) \equiv m(\mathbf{x})$   
 36 is taken to be a constant (in time) function, and only Algorithm 1 is used,  
 37 starting from an initial potential  $\tilde{Q}_0 = 0$ .

38 The **first example** is a simple symmetrical case in which we present meshes  
 39 generated by considering a monitor function which is large near the boundary of  
 40 a sphere. This serves to show the symmetry preserving properties of the PMA  
 41 equation and the regularity and alignment of the resulting meshes.

1 The **second example** is a more complicated, but still analytically determined,  
 2 monitor function describing a helical feature. This will show more clearly the  
 3 meshes which it is possible to construct which can represent a complex three  
 4 dimensional geometry.

5 **Finally** in this section we will consider the very large and practical problem of  
 6 generating adapted three-dimensional meshes for the purposes of meteorological  
 7 data assimilation calculations. In this example we use forecast data from the  
 8 Met Office UK4 model to define a monitor function based on an estimate of  
 9 the potential vorticity, looking at a sequence of different meteorological events.  
 10 This example illustrates the effectiveness of the PMA algorithm to generate a  
 11 mesh when used on a large scale practical three dimensional problem, with a  
 12 monitor function defined by data.

13 For all of the examples, the codes for the PMA algorithm were executed on a  
 14 laptop with an Intel® Core™2 Duo CPU P9400 @ 2.4Ghz with 4GB RAM run-  
 15 ning a 32-bit Linux OS and were compiled with the gfortran compiler in double  
 16 precision. All reported times are wall-clock times measured using `system_clock`,  
 17 averaged over 3 runs.

## 18 5.1 Simple test cases

### 19 5.1.1 Example 1: A three dimensional shell

`sec:shell`

20 We define the density  $f(\mathbf{x})$  of a smooth three dimensional ball with a (graded)  
 21 boundary of width  $r_2$  and centred on the point  $(x_0, y_0, z_0)$  as follows. Let  $s$  be  
 22 the distance of a point in our domain to the centre of the ball given by

$$s(\mathbf{x}) = s(x, y, z) = \sqrt{(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2}. \quad (19)$$

23 We then define the density of the ball via the function

$$f(\mathbf{x}) = f(x, y, z) = \begin{cases} 1 & \text{for } s(x, y, z) \leq r_1 \\ \frac{1}{2} \cos\left(\frac{(s(x, y, z) - r_1)\pi}{r_2}\right) + \frac{1}{2} & \text{for } s(x, y, z) \leq r_1 + r_2 \\ 0 & \text{for } s(x, y, z) > r_1 + r_2 \end{cases} \quad (20) \quad \text{cartball}$$

24 where  $r_1$  and  $r_2$  are scalars defining the width of the ball. For this problem  
 25 we will consider generating a mesh which concentrates points close to the shell  
 26 forming the boundary of the ball. This can be achieved by using a monitor  
 27 function which is large when the derivatives of the density function  $f(\mathbf{x})$  are  
 28 also large. Accordingly, we define the monitor function  $m(x, y, z)$  by

$$m(x, y, z) = \sqrt{(1 + c^2(f_x(x, y, z)^2 + f_y(x, y, z)^2 + f_z(x, y, z)^2))}. \quad (21) \quad \text{ballmonitor3d}$$

29 Here  $c$  is a regularisation constant, which we set in our examples to be  $c = 0.75$ .  
 We now consider a three dimensional mesh, constructed within the unit cube,  
 and adapted to this monitor function in which we set the parameters defining




Figure 2: Plot of distance  $r$  that the mesh moved in each iteration (The Euclidean change in the mesh) for the shell problem.

conv\_ball

the width of the ball to be  $r_1 = r_2 = \frac{1}{6}$ , and centred in the domain so that

$$(x_0, y_0, z_0)^T = \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)^T.$$



1 In the examples shown the computational domain  $\Omega_C = [0, 1]^3$  is split into a  
 2 grid of  $n_x \times n_y \times n_z$  points, with  $n_x = n_y = n_z = 100$  and is mapped into  
 3 the same physical domain (so that the solution of the PMA equation satisfies  
 4 Neumann boundary conditions).

5 The PMA algorithm was applied to this problem with  $\delta\tau = 0.2$  and  $\gamma = 0.2$ . The  
 6 convergence of the mesh to an equidistributed state to a tolerance of  $5E - 11$   
 7 is shown in Figure 2. The calculation terminated after 41 iterations, taking  
 8 34 seconds on the laptop computer described earlier. From this figure we can  
 9 clearly see the rapid, linear convergence of the algorithm.

10 The resulting mesh is presented in Figure 3. From this simple test problem  
 11 it is possible to see how the solution of the PMA equation is equidistributing  
 12 the monitor function. There are many more grid points in the region where  
 13 the monitor function is high than outside of that region, and the mesh shows  
 14 excellent alignment with the boundary of the ball. In Figure 3a we plot the  
 15 values of the monitor function in three dimension, with part of the ball cut  
 16 away to show the variation in value across the shell. In Figure 3b we show a  
 17 plane in the mesh that precisely follows the contours of the monitor function.  
 18 Figures 3c and 3d show the grid from the centre of the computational domain



1 projected onto the  $x$ - $y$  plane in physical space. Figure 3d shows the regularity  
2 of the grid that is generated and that the PMA equation aligns the mesh with  
3 the contours of the monitor function. This elegant behaviour arises because  
4 symmetries in the monitor function lead to symmetries in the PMA equation  
5 and hence in the function  $Q$ .





(a) 3D plot of the monitor function for  $m > 1.05$

(b) 3D view of grid in physical space of the grid from  $z = 1/3$  in computational space.

ball1

ball2



(c) Projected view of a plane of the mesh that was at  $z = 49/99$  in computational space

(d) Zoomed view of projected mesh from  $z = 49/99$  around the high monitor function.

ball3

ball4

Figure 3: Monitor function and resulting sections from the mesh for the shell test problem.

ballmesh


1 **5.1.2 Example 2: A three dimensional helix**

2 We next consider an analytically defined monitor function that describes a com-  
 3 plex three dimensional helical surface without the symmetries of the shell. Tak-  
 4 ing  $\mathbf{x} = (x, y, z)^T$  then a monitor function  $m(\mathbf{x})$  which is large in a neighbour-  
 5 hood of such a helix is given by

$$m(x, y, z) = \exp(-w_1[(x - (w_2 \cos(4z\pi) + 0.5))^2 + (y - (w_2 \sin(4z\pi) + 0.5))^2]) \quad (22)$$

helixm

6 Here the parameter  $w_1$  describes the width of this boundary neighbourhood, and  
 7 the parameter  $w_2$  gives the width of the helix. These are set to be  $w_1 = 100$   
 8 and  $w_2 = \frac{1}{4}$ . The domain is split into  $100 \times 100 \times 100$  grid points and the three  
 9 dimensional values of the monitor function are shown in Figure 4.



(a) 3D plot of the helical monitor function (b) Cut away plot of 3d monitor function

Figure 4: 3D plots of the helical monitor function showing only those points with  $m > 0.05$

helix\_mon

10 The PMA algorithm was applied to the helical problem with  $\delta\tau = 0.1$  and  
 11  $\gamma = 0.2$  and was successful in generating a highly non-uniform mesh without  
 12 any evidence of mesh tangling at any stage of the application of the algorithm.  
 13 The non-monotone convergence of the mesh to an equidistributed state to a  
 14 tolerance of  $5E - 11$  is shown in Figure 5. The calculation terminated after 473  
 15 iterations, taking 7.9 minutes on the laptop computer described earlier. We note  
 16 that this convergence is significantly slower than for the shell in the example in  
 17 Section 5.1.1. In Figure 6 we show the mesh generated by the PMA algorithm  
 18 when applied to the problem taking  $m$  as defined in (22). In Figure 6b we show  
 19 where the two horizontal planes in Figure 6a are mapped to in physical space.





Figure 5: Plot of distance the mesh moved in each iteration (Euclidean change in mesh) for the helical problem.

conv\_helix


1 Similarly Figures 6c and 6d show where the vertical planes in Figure 6a are  
 2 mapped to in physical space. These show that the redistributed grid is closely  
 3 following the monitor function and very clearly show the fully 3D nature of the  
 4 problem.

5 For this helical problem, compared with the other examples considered in this  
 6 paper, a substantially larger number of iterations are required in order to have a  
 7 suitable mesh and for the algorithm to converge. This is due to the complex and  
 8 highly non-uniform structure of the monitor function which we are considering  
 9 and the fact that the original uniform mesh has to encounter significant deviation  
 10 to wrap around the helical structure. Due to the twisted nature of the monitor  
 11 function, if we were to be aggressive with our strategy to find a mesh for this,  
 12 mesh tangling would easily occur on the path to the final mesh. To avoid this we  
 13 decreased our step size ( $\delta\tau$ ) in the forward Euler method and hence we increase  
 14 the number of iterations required to reach the same tolerance as considered in  
 15 the other examples. It is a significant achievement of the algorithm that it was  
 16 able to evolve the mesh in this manner without ever encountering a state where  
 17 the mesh was tangled.




(a) Planes in the computational mesh showing where the meshes in Figures 6b–6d originate in computational space

helix\_planes




(b) Location of the two horizontal planes in the physical space corresponding to the horizontal planes shown in computational space

helix\_hor



(c) Location of the plane in physical space corresponding to  $y = 1/3$  in computational space

helix\_ver1



(d) Location of the plane in physical space corresponding to  $y = 7/9$  in computational space

helix\_ver2

Figure 6: 3D plots of the mesh generated by the helical monitor function at various slices.

helix\_res

## 1 5.2 Meteorological test problems

2 We now consider a large scale meteorological problem for which the monitor  
3 function is not given as an analytic function, but is instead defined at a set  
4 of discrete data points. This is a commonly encountered situation both in the  
5 numerical solution of PDES or (as in this case) of function approximation where  
6 the function is only known at discrete points.

7 Data assimilation is the technique of matching noisy data to models of a process  
8 which also may have error. It is widely, and successfully, used in meteorology to  
9 determine an atmospheric state consistent both with observations and with the  
10 underlying physics of the atmosphere. In order to implement data assimilation  
11 methods effectively, it is important that the underlying covariance matrix of  
12 the errors is well represented. This matrix is too large to store explicitly. In  
13 this context adaptive mesh redistribution can be applied to create a simplified  
14 and thus manageable representation of the background error covariance matrix,  
15 and in particular include a reasonable representation of the spatially varying  
16 structure of the covariances [4, 1]. <sup>Piccolo2011, Piccolo2012</sup> The Met Office data assimilation system  
17 already implements a 1D adaptive meshing procedure for the vertical component  
18 of their grid used for their data assimilation algorithms. The improvement in  
19 data correlations represented by doing this has resulted in a measurable increase  
20 in forecasting accuracy [4, 1]. <sup>Piccolo2011, Piccolo2012</sup> In this paper we consider the first step of extending  
21 this work by considering how to use the PMA algorithm to generate a suitable  
22 3D mesh for data assimilation in a variety of meteorological conditions. A  
23 discussion of the implementation and testing of the adapted meshes within the  
24 data assimilation system will follow in a later paper.

25 To be effective within the context of a data assimilation calculation, the mesh  
26 generation code must be both fast and robust to use, and must also be easily  
27 linked to the existing data assimilation software. For the Met Office application,  
28 the goal is to produce a weather forecast after using data assimilation to get a  
29 best guess for the current state of the atmosphere. This imposes an immediate  
30 operational time restriction on the time-frame in which the computations can  
31 be made, as a forecast delivered after the event is useless. As a rule of thumb, a  
32 mesh which takes more than five minutes to generate is not useful operationally.  
33 This paper considers adapting the UK4 grid (4km horizontal spacing local area  
34 model over the British Isles) with efficiency a key consideration for any future  
35 operational implementation. As a code for an operational centre, the meshes  
36 produced will have to run automatically and hence be robust to all weather  
37 conditions. Thus it is essential to have a monitor function which is well scaled  
38 to maintain good global resolution while still refining sufficiently around features  
39 of interest.

40 This specific application of adaptive meshing is as an aide to help calculate the  
41 background error covariance matrix within the data assimilation algorithm, and  
42 thus to ease the finding of the minimum in the variational problem as in [1, 4]. <sup>Piccolo2012, Piccolo2011</sup>

### 5.3 Defining a monitor function

In this example, the physical coordinates  $\mathbf{x} = (x, y, z)$  correspond to longitude, latitude and vertical levels respectively. The vertical levels are defined using a terrain-following coordinate  $\eta$  which is a monotone function of height. It is plausible to assume that the correlation structure is isotropic in geostrophic and isentropic coordinates, which implies the use of the semi-geostrophic potential vorticity as a monitor function [20]. The PV is the Jacobian of the transformation from physical to geostrophic and isentropic coordinates. This is given in terms of the primitive variables  $u, v$  and  $\theta$  by

$$PV = \begin{vmatrix} f + v_x & v_y & v_z \\ -u_x & f - u_y & -u_z \\ g\theta_x/\theta_0 & g\theta_y/\theta_0 & g\theta_z/\theta_0 \end{vmatrix}$$

where  $f$  is the Coriolis parameter (assumed constant),  $u$  and  $v$  are the wind velocities in the longitudinal and latitudinal directions respectively,  $g$  is the force due to gravity,  $\theta$  is potential temperature and  $\theta_0$  a reference potential temperature [20]. Since the PV calculated from real data may not be positive, we use only the dominant diagonal terms of semigeostrophic potential vorticity to form the basis for the monitor function which we use to control the adapted mesh. Each of the diagonal terms is regularised to take account of the typical scale of the individual terms and ensure positivity. This resulting monitor function then has the following form

$$m = \begin{vmatrix} \sqrt{1 + c_1(1 + \frac{v_x}{10f})^2} & 0 & 0 \\ 0 & \sqrt{1 + c_2(1 - \frac{u_y}{10f})^2} & 0 \\ 0 & 0 & \sqrt{1 + c_3(\frac{\theta_z}{\theta_0})^2} \end{vmatrix}.$$

Note that the wind gradients  $u_y$  and  $v_x$  have been rescaled by a factor of 10 to remove some of the greater variability in the wind speeds than in the potential temperature. The constants  $c_1$ ,  $c_2$  and  $c_3$  are regularisation parameters which allow for different weightings to be given to the different components. With a great deal of testing, it was found that all the normalisation parameters equal 0.75 gave good results. Note that  $c_1 = c_2 = 0$  reduces this three dimensional monitor function to the one dimensional static stability based monitor function, which is currently used operationally [4, 1].

In the application to atmospheric data assimilation it is important to respect the stratified structure of the atmosphere. Though the monitor function should be smoothed to avoid computational difficulties caused by rapid grid variations, the smoothing should be applied only in the horizontal and not the vertical. Thus the filtering operator that is applied is

$$\tilde{m}_{i,j,k} = \frac{\sum_{\ell_1=-1}^1 \sum_{\ell_2=-1}^1 m_{i+\ell_1, j+\ell_2, k} \beta^{|\ell_1|+|\ell_2|}}{\sum_{\ell_1=-1}^1 \sum_{\ell_2=-1}^1 \beta^{|\ell_1|+|\ell_2|}} \quad (23)$$

1 This produces much sharper monitor functions and hence gives better refinement  
 2 of the grid around the structures of interest.

### 3 5.4 Test cases

4 In our calculations we considered three different meteorological data sets to test  
 5 the grid generation capabilities of the 3D PMA algorithm. These data sets were  
 6 actual forecast data provided by the UK Met Office for periods of very different  
 7 weather conditions, in particular: (a) a stable boundary layer, (b) scattered  
 8 showers, and (c) a frontal system. In keeping with the possible operational  
 9 restrictions on adapted grid generation, all parameters used in the subsequent  
 10 results will be fixed across all cases to show the robustness of the method.  
 11 In all of these calculations, the parameters used were  $\delta\tau = 0.5, \gamma = 0.5$  and  
 12 the convergence tolerance was set to  $5E - 11$ . The PMA algorithm performed  
 13 very well in each case and the meshes obtained captured all the features of the  
 14 underlying localised systems (identified by the monitor function). Consequently  
 15 we are confident that the resulting meshes should perform very well when used  
 16 for data assimilation calculations. The table below shows the convergence results  
 17 from the three test cases.

Test case	Iterations	CPU time (minutes)
Stable boundary layer	22	4.0
Scattered showers	22	4.2
Frontal system	21	5.4

tab:met

Table 1: Results for the three meteorological test cases

18 Observe, that even in these large data sets, the PMA algorithm converges  
 19 rapidly. We now show the resulting meshes in each case. For each figure we give  
 20 the monitor function and the mesh at appropriate sections through the domain.

#### 21 5.4.1 Stable boundary layer

22 This test case uses the same UK4 model data described in [Piccolo2012](#)  
 23 scenario when UK was mainly covered by low-level clouds. The synoptic situa-  
 24 tion over the UK at the time (3rd January 2011 at 00UTC) was characterised  
 25 by a weak flow within a large anticyclone of 1030 hPa surface pressure. Ob-  
 26 served vertical profiles show saturated boundary layer below an inversion of  
 27 850 hPa. There is a warm front in the south-west with some likely enhancements  
 28 from a vorticity anomaly aloft. This is associated with extensive low clouds  
 29 particularly in the south-west. Figure 7 shows a cross section (longitude versus  
 30 levels) of the monitor function described in Section 5.3 for 3 January 2011 at  
 31 00 UTC and the corresponding mesh. The three dimensional monitor function  
 32 clearly captures the vertical structures in the troposphere which indicates the  
 33 presence of clouds at different levels in agreement with the results showed when  
 34 using the one dimensional static stability monitor function described in [Piccolo2012](#). The



1 mesh follows the monitor function by moving the vertical height levels further  
2 together when the monitor function is greater than one and further apart when  
3 it is smaller than one. This is in agreement with the one dimensional results. In  
4 addition the three dimensional monitor function moves the mesh horizontally  
5 capturing more realistically local variations of the cloud layering.

6 Another cross section is shown in Figure 8. Again the mesh (latitudes versus  
7 height levels) follows the structure of the corresponding monitor function and  
8 captures local variability both vertically and horizontally.




Figure 7: The monitor function and the resulting mesh for the stable boundary layer system at a  $94^{\text{th}}$  latitude increment ad with increasing longitude. The function is shown in the vertical plane from  $(50.68N, 11.51W)$  to  $(50.80N, 4.84E)$

fig7

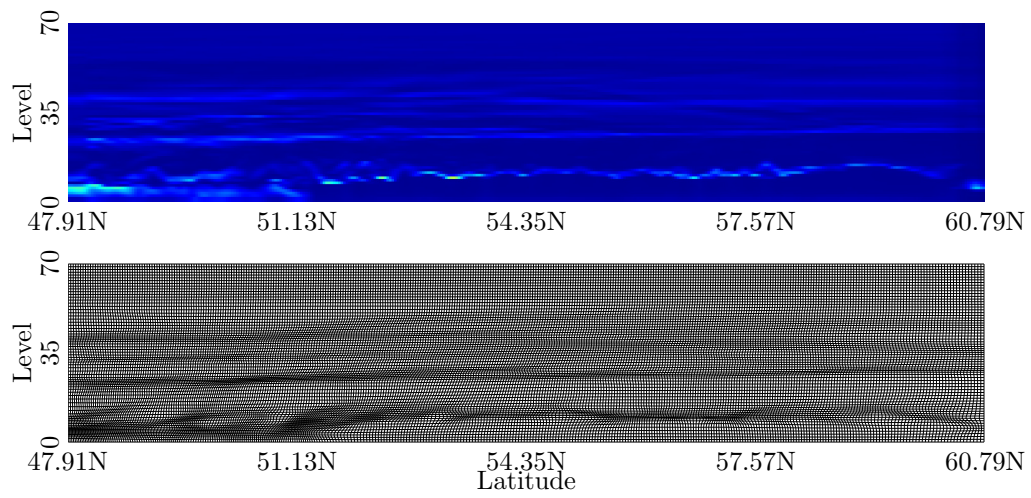


Figure 8: The monitor function and the mesh for the stable boundary layer system at a 260<sup>th</sup> longitude increment. In the vertical plane from (47.91N, 2.89E) to (60.79N, 4.86E).

fig8

### 1 **5.4.2 Scattered showers**

2 The next two cases have been selected to test the capability of the scheme  
3 to capture two different extremes, i.e. localised convective activity as in the  
4 scenario of scattered showers and a large scale weather system as in the case of  
5 a front. The synoptic situation over the UK on the 24 April 2012 at 12UTC was  
6 characterised by a weak flow within a large scale upper trough with an upper  
7 filament of vorticity in the south-west of England giving focus to the convective  
8 activity. The latter gives large values of the monitor function. The convective  
9 activity is shown by the radar image in Figure 9. The adaptive mesh scheme  
10 here needs to pick up very small and localised showers scattered over the UK  
11 as well as the response to the large scale forcing over SW England.

12 Figure 10 shows an horizontal cross section of the monitor function on the left  
13 and the corresponding mesh on the right for a low height level of the model.  
14 The monitor function tends to capture local and small scale phenomena. These  
15 do not coincide with the radar image in Figure 9, this is because the monitor  
16 function is calculated from a T+3h forecast and not from observations. The  
17 monitor function does not respond to the random showers over Ireland, but  
18 does pick up the area with no showers over central England. The mesh follows  
19 the monitor function behaviour and clustered mesh points near the high values of  
20 the monitor function. When the showers are better organised and less random,  
21 like the filament over North Scotland, the mesh nicely aligns with this feature.  
22 Figure 11 shows instead a vertical cross section (latitudes versus height levels)  
23 for the same case. As well as capturing the small scale variations due to the  
24 showers the monitor function picks up the upper filament of vorticity (around  
25 level 35) and the lower filament over north Scotland (around level 8). The  
26 mesh nicely follows the behaviour of the monitor function both horizontally and  
27 vertically.

### 28 **5.4.3 A Frontal system**

29 The last case described in this section follows from the scattered showers weather  
30 system. The large upper trough described in the previous section extends south  
31 and by 00UTC on the 25 April 2012 it drives the surface cyclonic system east-  
32 ward bringing a warm front system into the south-west of UK. The activity  
33 on the front is strongly enhanced by vorticity forcing at 250 hPa. Figure 12  
34 shows the radar image for the front system on the 25 April 2012 at 03UTC. The  
35 horizontal cross section of the monitor function and the corresponding mesh for  
36 this case are shown in Figure 13. The front is clearly depicted in both pictures  
37 and the refinement of the mesh is high in correspondence with the front. Figure  
38 14 shows the vertical cross section (latitude versus levels) of the monitor func-  
39 tion and the resulting mesh. It clearly picks up the three dimensional structure  
40 of the front (around latitude 50N) as a function of height and latitude. The  
41 monitor function also displays extra vertical structures over the UK. Again the  
42 mesh nicely follows the behaviour of the monitor function both horizontally and  
43 vertically.




fig9

Figure 9: Radar image of the scattered showers system.




Figure 10: The monitor function and the mesh for the scattered shower system at a 8<sup>th</sup> vertical level. At 261.7m

fig10

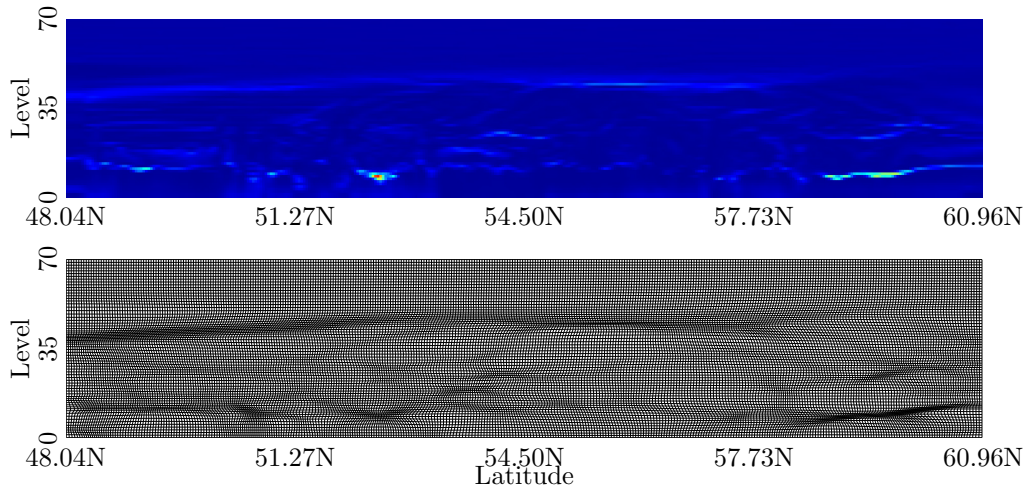


Figure 11: The monitor function and the mesh for the scattered shower system at a 135<sup>th</sup> longitude increment. Vertical plane from (48.04N, 3.81W) to (60.96N, 4.29W).

fig11




fig12

Figure 12: The radar image of the frontal system crossing the South West coast of the British Isles




Figure 13: The monitor function and the mesh of the frontal system at a 23<sup>rd</sup> vertical level at 1911.7m

fig13




Figure 14: Monitor function and mesh of frontal system at 16<sup>th</sup> longitude increment. Vertical plane from (47.77N, 10.17W) to (60.60N, 12.94W).

fig14

## 1 6 A moving mesh test problem

2 We now consider the performance of the PMA algorithm when used to compute a  
 3 time varying three-dimensional mesh when the monitor function  $m(\mathbf{x}, t)$  is itself  
 4 a function of time. This situation of course is closer to a typical implementation  
 5 of a mesh redistribution method when it would be used to as part of the solution  
 6 of a time varying PDE. In this section the example considered is the same as  
 7 that studied by Chacón et al. [6] which also considers calculating a mesh by  
 8 solving the Monge-Ampère equation, but which uses a Newton method coupled  
 9 with a multi-grid solver to do this. To find the mesh in this case we implement  
 10 Algorithm 2 as described earlier.

11 The time-varying, analytically defined, monitor function considered is given by:

$$m(x, y, z, t) = 1 + 4 \exp \left( -r(x, y, z)^2 \left( \frac{\cos^2(\kappa(x, y, z, t))}{\sigma_x^2} + \frac{\sin^2(\kappa(x, y, z, t))}{\sigma_y^2} \right) \right) \quad (24)$$

12 where  $r(x, y, z)$  is the distance to the centre of the domain at  $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ ,  $\sigma_x =$   
 13  $\sqrt{0.05}$ ,  $\sigma_y = \sqrt{0.001}$  are scaling factors and

$$\kappa(x, y, z, t) = \arctan \left( \frac{y - \frac{1}{2}}{x - \frac{1}{2}} \right) + 1.6 \sin(\pi z) \max[(\frac{1}{2} - r)t, 0]. \quad (25)$$

14 The goal of this test problem is to find meshes at times  $t \in \{0, 1, \dots, 100\}$ . The  
 15 problem of finding the mesh for this time dependent system is then solved in  
 16 two stages in a manner analogous to the MMPDE method described in [13].

17 **Firstly** at time  $t = 0$  Algorithm 2 sets the monitor function  $m(\mathbf{x}, 0)$  and,  
 18 starting from a uniform mesh, the system (14) is evolved forward in pseudo-time  
 19 using Algorithm 1 with  $m(\mathbf{x}, 0)$  fixed until the mesh satisfies the equidistribution  
 20 condition to a high tolerance. For this calculation we take  $\delta\tau = 0.1$ ,  $\gamma = 0.2$   
 21 and  $tol = 5E - 11$ .

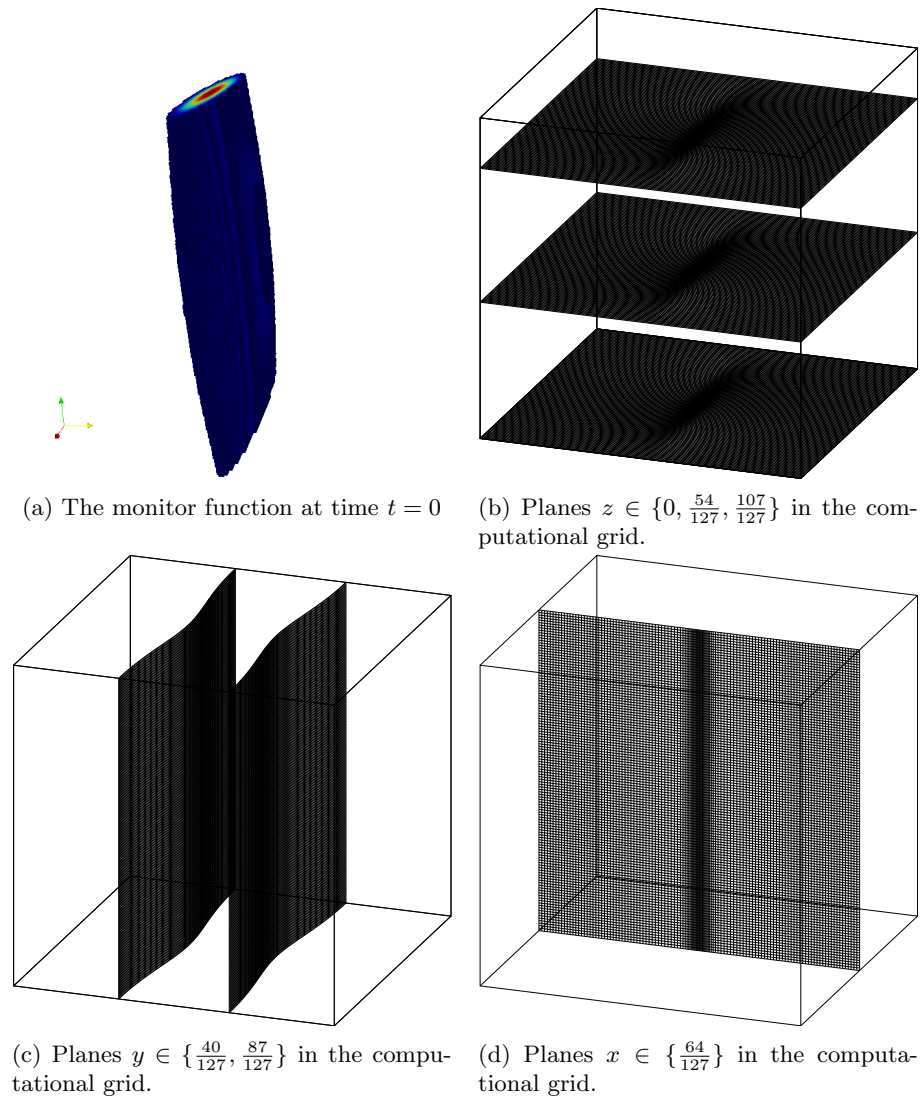
22 **Secondly** Algorithm 2 evolves the monitor function in real time, with the value  
 23 of  $t$  increased in intervals of  $\delta t = 1.0$ . For each of these outer timesteps, we set  
 24  $\tau_{\max} = \delta t$  and  $\delta\tau = \delta t/5$ , ensuring at least 5 pseudo-timesteps per inner loop.

25 Some of the resulting meshes for the case of a  $128 \times 128 \times 128$  mesh are presented  
 26 as follows. In Figure 15 we show the monitor function and the resulting mesh  
 27 at the initial time  $t = 0$ . In Figures 16 and 17 we then show the evolved meshes  
 28 at the later times  $t = 50$  and  $t = 100$ .

29 We can see at time  $t = 100$  that the mesh closely follows the contours of the  
 30 monitor function and is very regular with no hint of mesh tangling.


31 We next consider the computational cost of calculating these meshes. To do this  
 32 the unit cube is discretised into a grid of  $N \times N \times N$ , where  $N = 32, 64, 128$ ,






cfig1

Figure 15: The monitor function and the resulting meshes at the time  $t = 0$



cfi2

Figure 16: The monitor function and the resulting meshes at the time  $t = 50$



cfig3 Figure 17: The monitor function and the resulting meshes at the time  $t = 100$

1 and we list the number of iterations to converge to the given tolerance in the  
 2 pseudo-time calculation at  $t = 0$  and the total CPU time required to compute  
 3 the 101 meshes until  $t = 100$ . These results are presented in Table 2.

Grid resolution $N$	DOFs	Initial iterations of static PMA	CPU(s)
$32 \times 32 \times 32$	98304	44	15.91
$64 \times 64 \times 64$	786432	45	180.28
$128 \times 128 \times 128$	6291456	44	2356.29

Table 2: Timings for the evolution of the mesh to an equidistributed state for varying spatial discretisations

table1

4 A direct comparison with the results in Chacón et. al. <sup>Chacon2011</sup> [6] can be made with  
 5 these results. In their paper they describe and implement a Newton-Krylov  
 6 iteration using modern multigrid methods to solve exactly the fully non-linear  
 7 Monge-Ampère equation. For the high resolution  $128 \times 128 \times 128$  grid this  
 8 computation required 32000s of CPU time. The methods produced by the PMA  
 9 algorithm appear similar in structure, despite the significantly reduced cost of  
 10 their calculation.

## 1 7 Conclusion

2 In this paper we have demonstrated that the Parabolic Monge-Ampère algo-  
3 rithm can be extended from two dimensions to three, and that it is effective  
4 in generating meshes with good regularity in a short time. In particular it can  
5 deliver effective meshes for three dimensional meteorological data assimilation  
6 calculations using large data sets with 21 million degrees of freedom, in times  
7 commensurate with those required for actual weather forecasting. When applied  
8 to test problems it shows rapid convergence, with meshes rapidly (and without  
9 any hint of tangling) converging to an equidistributed state. In particular it is  
10 an order of magnitude faster in converging than other similar mesh generation  
11 methods. We therefore think that this method should be considered seriously  
12 as a fast and effective method for redistributing a large three dimensional mesh.

## 13 References

- [Piccolo2012](#)<sub>14</sub> [1] Chiara Piccolo and Mike Cullen. A new implementation of the adaptive  
15 mesh transform in the Met Office 3D-Var System. *Quarterly Journal of the*  
16 *Royal Meteorological Society*, 138(667):1560–1570, July 2012.
- [Gullbrand2003](#)<sub>17</sub> [2] Jessica Gullbrand and Fotini Katopodes Chow. The effect of numerical  
18 errors and turbulence models in large-eddy simulations of channel flow,  
19 with and without explicit filtering. *Journal of Fluid Mechanics*, 495:323–  
20 341, November 2003.
- [Huang1996](#)<sub>21</sub> [3] Weizhang Huang and Robert D Russell. A moving collocation method for  
22 solving time dependent partial differential equations. *Applied Numerical*  
23 *Mathematics*, 20:101–116, 1996.
- [Piccolo2011](#)<sub>24</sub> [4] Chiara Piccolo and Mike Cullen. Adaptive mesh method in the Met Of-  
25 fice variational data assimilation system. *Quarterly Journal of the Royal*  
26 *Meteorological Society*, 137(656):631–640, April 2011.
- [Pain2005](#)<sub>27</sub> [5] C.C. Pain, M.D. Piggott, A.J.H. Goddard, F. Fang, G.J. Gorman, D.P.  
28 Marshall, M.D. Eaton, P.W. Power, and C.R.E. de Oliveira. Three-  
29 dimensional unstructured mesh ocean modelling. *Ocean Modelling*, 10(1-  
30 2):5–33, January 2005.
- [Chacon2011](#)<sub>31</sub> [6] L. Chacón, G.L. Delzanno, and J.M. Finn. Robust, multidimensional mesh-  
32 motion based on Monge-Kantorovich equidistribution. *Journal of Computa-*  
33 *tional Physics*, 230(1):87–103, January 2011.
- [Behrens1998](#)<sub>34</sub> [7] Jörn Behrens. Atmospheric and ocean modeling with an adaptive finite  
35 element solver for the shallow-water equations. *Applied Numerical Mathe-*  
36 *matics*, 26(1-2):217–226, January 1998.
- [Weller2009](#)<sub>37</sub> [8] Hilary Weller. Predicting mesh density for adaptive modelling of the global  
38 atmosphere. *Philosophical transactions. Series A, Mathematical, physical,*  
39 *and engineering sciences*, 367(1907):4523–4542, November 2009.

- Ainsworth1997 [9] Mark Ainsworth and Bill Senior. Aspects of an adaptive hp-finite element method: Adaptive strategy, conforming approximation and efficient solvers. *Computer Methods in Applied Mechanics and Engineering*, 150(1-4):65–87, December 1997.
- behrens2006 [10] Jörn Behrens. *Adaptive atmospheric modeling: key techniques in grid generation, data structures, and numerical operations with applications*. Springer, 2006.
- Budd2009a [11] C J Budd and J F Williams. Moving mesh generation using the parabolic Monge-Ampère Equation. *SIAM Journal on Scientific Computing*, 31(5):3438–3465, 2009.
- Budd2009b [12] Chris J. Budd, Weizhang Huang, and Robert D. Russell. Adaptivity with moving grids. *Acta Numerica*, 18:1–131, May 2009.
- huang2011 [13] Weizhang Huang and R Robert D Russell. *Adaptive moving mesh methods*, volume 174. Springer, 2011.
- Delzanno2008 [14] G.L. Delzanno, L. Chacón, J.M. Finn, Y. Chung, and G. Lapenta. An optimal robust equidistribution method for two-dimensional grid adaptation based on Monge-Kantorovich optimization. *Journal of Computational Physics*, 227(23):9841–9864, December 2008.
- Budd2013 [15] C.J. Budd, M.J.P. Cullen, and E.J. Walsh. Monge-Ampère based moving mesh methods for numerical weather prediction, with applications to the Eady problem. *Journal of Computational Physics*, 236:247–270, March 2013.
- Cao2005 [16] Weiming Cao. On the Error of Linear Interpolation and the Orientation, Aspect Ratio, and Internal Angles of a Triangle. *SIAM Journal on Numerical Analysis*, 43(1):19–40, January 2005.
- Huang1994a [17] Weizhang Huang, Yuhe Ren, and Robert D. Russell. Moving Mesh Partial Differential Equations (MMPDES) Based on the Equidistribution Principle. *SIAM Journal on Numerical Analysis*, 31(3):709–730, 1994.
- Brenier1991 [18] Yann Brenier. Polar Factorization and Monotone Rearrangement of Vector-Valued Functions. *Communications on Pure and Applied Mathematics*, XLIV:375–417, 1991.
- sewell12002 [19] M J Sewell. Some applications of transformation theory in mechanics. *Large Scale Atmosphere–Ocean Dynamics*, 2:143–223, 2002.
- cullen2006 [20] M.J.P. Cullen. *A Mathematical Theory of Large-Scale Atmosphere/Ocean Flow*. Imperial College Press, 2006.
- Chynoweth1991 [21] S. Chynoweth and J Sewell. A concise derivation of the semi-geostrophic equations. *Quarterly Journal of the Royal Meteorological Society*, 117(502):1109–1128, 1991.

- `Ceniceros2001` [22] Hector D. Ceniceros and Thomas Y. Hou. An Efficient Dynamically Adaptive Mesh for Potentially Singular Solutions. *Journal of Computational Physics*, 172(2):609–639, September 2001.
- `Budd2011` [23] C J Budd and V A Galaktionov. On Self-similar blow-up in evolution equations of Monge-Ampère Type: A view from Reaction-Diffusion Theory. *IMA Journal of Applied Mathematics*, 2011.
- `fftw` [24] Matteo Frigo and Steven G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. Special issue on “Program Generation, Optimization, and Platform Adaptation”.