

Preprint of

R. Sevilla, O. Hassan and K. Morgan

The use of hybrid meshes to improve the efficiency of a discontinuous Galerkin method for the solution of Maxwell's equations

Computers & Structures, 137; 2-13, 2014

The use of hybrid meshes to improve the efficiency of a discontinuous Galerkin method for the solution of Maxwell's equations

Ruben Sevilla, Oubay Hassan, Kenneth Morgan

Civil and Computational Engineering Centre (C²EC), College of Engineering, Swansea University, Singleton Park, Swansea, SA2 8PP, Wales, UK.

Abstract

This paper investigates the efficiency of a high-order nodal discontinuous Galerkin method for the numerical solution of Maxwell's equations using hybrid meshes. An unstructured triangular or tetrahedral mesh is used near curved boundaries and a structured Cartesian mesh is used to fill the remainder of the domain. A quadrature-free implementation is employed for the regular quadrilateral and hexahedral elements which, together with the reduction in the number of internal faces, leads to a reduction in the cpu time requirements. Numerical examples in two and three dimensions are used to illustrate the benefits of using hybrid meshes.

Keywords: Maxwell's equations, discontinuous Galerkin, high-order approximations, hybrid meshes, computational cost

1. Introduction

In the area of computational electromagnetics, the Yee scheme [1] remains the predominant technique for the solution of Maxwell's equations in the time domain. This is due to the simplicity of the algorithm and to its low operation count. However, in its simplest form, this finite difference time domain (FDTD) method requires the use of structured meshes, which creates

Email addresses: R.Sevilla@swansea.ac.uk (Ruben Sevilla), O.Hassan@swansea.ac.uk (Oubay Hassan), K.Morgan@swansea.ac.uk (Kenneth Morgan)

difficulties when attempting to solve problems involving complex geometries. Greater geometric flexibility can be achieved by using unstructured meshes and this has motivated the development of hybrid solution techniques, in which an unstructured finite element (FE) or finite volume (FV) algorithm is combined with the standard FDTD method [2, 3, 4, 5]. However, the linear approximation of curved boundaries, associated with standard low-order FE or FV methods on unstructured meshes, can create non-physical diffraction effects, especially for high-frequency problems. Highly accurate, or exact, geometric representation is known to have a significant impact on the quality of the numerical solution in many electromagnetic applications [6, 7]. In addition, low-order methods suffer from numerical dispersion and/or dissipation, which compromises the accuracy of the computed solution when electromagnetic waves have to be propagated over large distances. Finally, the coupling between FE or FV methods and the FDTD solver requires specific strategies to avoid instabilities when propagating the waves for a very long time [8].

High-order methods have emerged as a promising alternative. They enable the high-fidelity solution of complex problems in electromagnetics, by incorporating an accurate representation of curved geometries [9, 10] and by reducing the levels of numerical dispersion and dissipation [11]. Different approximation approaches have been proposed, including spectral methods, high-order finite difference, finite volume or finite element methods and high-order discontinuous Galerkin (DG) methods [12].

The development of high-order DG methods for the numerical solution of Maxwell's equations in the time domain has been a particularly active area of research [13, 14, 15, 16, 17, 18, 19, 20, 21]. As the automatic generation of unstructured hexahedral meshes is still an open problem [22, 23], most attention has been directed towards the use of unstructured triangular and tetrahedral meshes, in order that fully automatic mesh generators can be employed. In addition, triangular or tetrahedral meshes can be readily generated in which the majority of the elements are affine. The use of affine elements significantly improves computational performance, as element matrices can be pre-computed and scaled using the constant Jacobian of the iso-parametric mapping [24].

The work presented here focuses on the efficiency of a high-order nodal DG method for the numerical solution of Maxwell's equations in the time domain, using different types of elements. Compared to a modal DG method, the nodal approach uses Lagrange polynomials to define the approximation.

Then, the approximation over an element face is only dependent on shape functions associated with nodes belonging to the face, reducing the computational cost incurred by the DG method when integrals over interior faces are computed.

To retain the flexibility and maturity of low-order unstructured mesh generators, a hybrid mesh approach is adopted in which an unstructured triangular or tetrahedral mesh is used near curved boundaries and a structured Cartesian mesh is used to fill the remainder of the domain. With this approach, the benefits of using affine elements can be exploited, not only for triangles or tetrahedra but also, for the quadrilateral, hexahedral and pyramidal elements. Moreover, the use of quadrilateral or hexahedral meshes results in a significant reduction in the number of internal faces, compared to triangular or tetrahedral meshes. It can be expected that this, in turn, will lead to a reduction in the cpu time requirements, as an important portion of the computational cost of a DG scheme is associated with the integration that is required over element faces. In comparison with many hybrid solution methods, that use different numerical schemes on the unstructured and structured portions of the mesh, the proposed approach simplifies the implementation and avoids the requirement for transfer of information between schemes of different nature.

The remainder of the paper is organised as follows. Sections 2 and 3 briefly summarise the DG formulation for the Maxwell's equations and the different finite elements that are employed for the spatial discretisation. In Section 4, the advantages of using affine elements are discussed and the proposed approach of using hybrid meshes is described. Numerical examples in two and three dimensions are used in Section 5, to compare the efficiency of the different finite elements, and, finally, Section 6 summarises the main conclusions of the work that has been presented.

2. Discontinuous Galerkin formulation for Maxwell's equations

Maxwell's curl equations, governing the propagation of electromagnetic waves through free space, can be written in the linear dimensionless conservation form

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}_k}{\partial x_k} = \mathbf{0} \quad k = 1, \dots, \mathbf{n}_{\text{sd}} \quad (1)$$

where n_{sd} ($= 2$ or 3) denotes the number of spatial dimensions and the summation convention is employed. The unknown vector \mathbf{U} is given by

$$\mathbf{U} = \begin{pmatrix} \mathbf{E} \\ \mathbf{H} \end{pmatrix} \quad (2)$$

and the flux vectors, \mathbf{F}_k , are

$$\mathbf{F}_1 = \begin{pmatrix} 0 \\ H_3 \\ -H_2 \\ 0 \\ -E_3 \\ E_2 \end{pmatrix} \quad \mathbf{F}_2 = \begin{pmatrix} -H_3 \\ 0 \\ H_1 \\ E_3 \\ 0 \\ -E_1 \end{pmatrix} \quad \mathbf{F}_3 = \begin{pmatrix} H_2 \\ -H_1 \\ 0 \\ -E_2 \\ E_1 \\ 0 \end{pmatrix} \quad (3)$$

Here, $\mathbf{E} = (E_1, E_2, E_3)^T$ and $\mathbf{H} = (H_1, H_2, H_3)^T$ are the dimensionless electric and magnetic field intensity vectors respectively. In two spatial dimensions, the system of equation (1) decouples into the transverse electric (TE) and transverse magnetic (TM) modes [25].

This system of equations is supplemented with appropriate boundary conditions. Three different boundary conditions are involved in the examples considered in this work: Dirichlet, perfect electrical conductor (PEC) and non-reflecting boundary conditions. As usual in a DG framework, Dirichlet boundary conditions are imposed in a weak form using numerical fluxes. At the surface of a PEC, the tangential component of the electric field vanishes and the condition

$$\mathbf{n} \times \mathbf{E} = \mathbf{0}$$

is applied, where \mathbf{n} denotes the unit normal vector to the PEC surface. For problems posed on unbounded domains, the computational domain is truncated and a non-reflecting boundary condition is imposed at the truncated boundary. This is achieved by the addition of an uniaxial perfectly matched layer (UPML) [26].

The domain Ω is represented by an unstructured assembly of elements and a DG weak formulation for equation (1) may be expressed over element Ω_e as

$$\int_{\Omega_e} \mathbf{W} \cdot \frac{\partial \mathbf{U}_e}{\partial t} d\Omega - \int_{\Omega_e} \frac{\partial \mathbf{W}}{\partial x_k} \cdot \mathbf{F}_k(\mathbf{U}_e) d\Omega + \int_{\Gamma_e} \mathbf{W} \cdot \mathbf{F}_n(\mathbf{U}_e) d\Gamma = 0 \quad (4)$$

Here, \mathbf{U}_e denotes the restriction of \mathbf{U} to the element Ω_e , \mathbf{n} is the outward unit normal vector to the boundary Γ_e of Ω_e , \mathbf{F}_n is the normal flux on Γ_e and \mathbf{W} is a test function. In DG methods, the discontinuous nature of the approximation is accounted for by replacing the physical normal flux at the boundary by a consistent numerical flux, $\tilde{\mathbf{F}}_n(\mathbf{U}_e, \mathbf{U}^{\text{out}})$. This numerical flux is evaluated in terms of the trace of the solution on element Ω_e and the trace of the solution, \mathbf{U}^{out} , on the other element adjacent to Γ_e . A natural choice, for the linear hyperbolic system of interest here, is to employ a flux splitting technique [27], which corresponds to an upwind approximation [17]. The normal flux \mathbf{F}_n is decomposed into incoming flux (superscript $-$) and outgoing flux (superscript $+$)

$$\mathbf{F}_n(\mathbf{U}) = \mathbf{F}_n^-(\mathbf{U}) + \mathbf{F}_n^+(\mathbf{U})$$

where the incoming and outgoing fluxes are associated with the negative and positive eigenvalues of the jacobian matrix $\mathbf{A}_n = \frac{\partial \mathbf{F}_n}{\partial \mathbf{U}}$ respectively. The numerical flux is then computed as

$$\tilde{\mathbf{F}}_n(\mathbf{U}, \mathbf{U}^{\text{out}}) = \mathbf{F}_n^+(\mathbf{U}) + \mathbf{F}_n^-(\mathbf{U}^{\text{out}}) \quad (5)$$

Introducing the resulting numerical normal flux into equation (4), the DG weak formulation for element Ω_e can be re-written as

$$\int_{\Omega_e} \mathbf{W} \cdot \frac{\partial \mathbf{U}_e}{\partial t} d\Omega + \int_{\Omega_e} \mathbf{W} \cdot \frac{\partial \mathbf{F}_k(\mathbf{U}_e)}{\partial x_k} d\Omega + \int_{\Gamma_e} \mathbf{W} \cdot \mathbf{A}_n^- [[\mathbf{U}_e]] d\Gamma = 0 \quad (6)$$

where $[[\mathbf{U}_e]] = \mathbf{U}_e - \mathbf{U}^{\text{out}}$ denotes the jump in the solution across Γ_e . In the evaluation of the boundary term, the expression

$$\mathbf{A}_n^- [[\mathbf{U}]] = \frac{1}{2} \begin{pmatrix} -\mathbf{n} \times [[\mathbf{H}]] + \mathbf{n} \times (\mathbf{n} \times [[\mathbf{E}]]) \\ \mathbf{n} \times [[\mathbf{E}]] + \mathbf{n} \times (\mathbf{n} \times [[\mathbf{H}]]) \end{pmatrix}$$

is used.

3. Discretisation in space

Triangles and quadrilaterals are employed to provide a consistent discretisation of the spatial solution domain, Ω , for two dimensional problems. In three dimensions, consistent meshes consisting of tetrahedra, hexahedra,

prisms and pyramids are used. Apart from the pyramid, which requires special attention, optimal nodal finite elements of arbitrary order are readily defined for all these shapes. For the pyramid, a recently proposed approximation space [28] is adopted. This space is well suited for both continuous and discontinuous approximations and is optimal, i.e. the *a priori* error estimate is $\mathcal{O}(h^{p+1})$ in the $\mathcal{L}^2(\Omega)$ norm, where p denotes the order of the approximation. The approximation spaces that are employed are summarised in Table 1.

<i>Element</i>	<i>Approximation Space</i>	<i>Dimension</i>
triangle	$\mathbb{P}_p(\xi_1, \xi_2) = \{\xi_1^i \xi_2^j, i + j \leq p\}$	$\frac{1}{2}(p+1)(p+2)$
quadrilateral	$\mathbb{Q}_p(\xi_1, \xi_2) = \{\xi_1^i \xi_2^j, i, j \leq p\}$	$(p+1)^2$
tetrahedron	$\mathbb{P}_p(\xi_1, \xi_2, \xi_3) = \{\xi_1^i \xi_2^j \xi_3^k, i + j + k \leq p\}$	$\frac{1}{6}(p+1)(p+2)(p+3)$
hexahedron	$\mathbb{Q}_p(\xi_1, \xi_2, \xi_3) = \{\xi_1^i \xi_2^j \xi_3^k, i, j, k \leq p\}$	$(p+1)^3$
prism	$\mathbb{P}_p(\xi_1, \xi_2) \otimes \mathbb{P}_p(\xi_3)$	$\frac{1}{2}(p+1)^2(p+2)$
pyramid	$\mathbb{P}_p(\xi_1, \xi_2, \xi_3) \oplus \sum_{k=0}^{r-1} \left(\frac{\xi_1 \xi_2}{1 - \xi_3} \right)^{r-k} \mathbb{P}_p(\xi_1, \xi_2)$	$\frac{1}{6}(p+1)(p+2)(2p+3)$

Table 1: The reference elements of order p , their approximation spaces and the dimension of these spaces.

To obtain a semi-discrete formulation, the solution, \mathbf{U} , is approximated as

$$\mathbf{U}(\boldsymbol{\xi}) \simeq \sum_{j=1}^{\mathbf{n}_{\text{en}}} \mathbf{U}_j(t) N_j(\boldsymbol{\xi})$$

on a reference element with local coordinates $\boldsymbol{\xi} = (\xi_1, \xi_2, \xi_3)$. Here $\mathbf{U}_j(t)$ denotes the value of the solution at node j , N_j is the shape function associated with node j and \mathbf{n}_{en} is the total number of nodes on the element. The nodal shape functions span the approximation spaces described in Table 1 and the number of nodes per element is equal to the dimension of the corresponding approximation space.

In two dimensions, a Fekete nodal distribution is adopted for the triangle [29] and a tensor product of one dimensional Fekete nodal distributions

for the quadrilateral. Figure 1 shows the reference triangle and the reference quadrilateral, with a nodal distribution corresponding to $p = 5$. In three dimensions, the nodal distributions proposed in [30] for the tetrahedron and in [28] for the pyramid are used. A tensor product of one dimensional Fekete nodal distributions is used for the hexahedron and a tensor product of triangular and one dimensional Fekete nodal distributions is used for the prism. Figure 2 shows the reference tetrahedron, hexahedron, prism and pyramid with a nodal distribution corresponding to $p = 3$.




Figure 1: Reference elements with $p = 5$ showing the location of the nodes: (a) triangle; (b) quadrilateral.

Introducing the approximate solution into the weak form of equation (6), the system

$$\sum_{j=1}^{\mathbf{n}_{\text{en}}} \mathbf{M}_{ij} \mathbf{I} \frac{d\mathbf{U}_j}{dt} + \sum_{j=1}^{\mathbf{n}_{\text{en}}} \left(\mathbf{C}_{ij}^k \mathbf{A}_k \right) \mathbf{U}_j - \sum_{j=1}^{\mathbf{n}_{\text{fn}}} \mathbf{M}_{ij}^f \mathbf{A}_n^- [\mathbf{U}_j] = \mathbf{0} \quad (7)$$

of ordinary differential equations is obtained for every node i of element Ω_e . This system is advanced in time using a standard explicit fourth-order Runge–Kutta method with a time step small enough to guarantee that the numerical error is dominated by the spatial discretisation error. Here, \mathbf{M} denotes the element mass matrix, \mathbf{C}^k the convection matrix in the x_k direction, \mathbf{M}^f is the face mass matrix, \mathbf{I} is the identity matrix and \mathbf{n}_{fn} is the number of nodes per face. These matrices are defined by

$$\mathbf{M}_{ij} = \int_{\Omega_e} N_i N_j d\Omega \quad \mathbf{C}_{ij}^k = \int_{\Omega_e} N_i \frac{\partial N_j}{\partial x_k} d\Omega \quad \mathbf{M}_{ij}^f = \int_{\Gamma_e} N_i N_j d\Gamma$$




Figure 2: Reference elements with $p = 3$ showing the location of the nodes: (a) tetrahedron; (b) hexahedron; (c) prism; (d) pyramid.

Using an isoparametric mapping, the integrals over element Ω_e are evaluated on the reference element, $\hat{\Omega}$ as

$$\begin{aligned}
 \mathbf{M}_{ij} &= \int_{\hat{\Omega}} N_i N_j |\mathbf{J}| d\Omega & \mathbf{J} &= \frac{\partial \mathbf{x}}{\partial \xi} \\
 \mathbf{C}_{ij}^k &= \int_{\hat{\Omega}} N_i \left(\sum_{l=1}^{n_{sd}} \mathbf{J}_{kl}^{-1} \frac{\partial N_j}{\partial \xi_l} \right) |\mathbf{J}| d\Omega & \mathbf{M}_{ij}^f &= \int_{\hat{\Gamma}} N_i N_j \|\mathbf{J}^f\| d\Gamma
 \end{aligned} \tag{8}$$

where \mathbf{J} denotes the Jacobian of the isoparametric mapping and \mathbf{J}^f is the Jacobian of the restriction of the isoparametric mapping to the element face f .

The computational effort required to evaluate the element and face matri-

ces is highly dependent upon the number of integration points used to approximate the integrals. For quadrilateral and hexahedral elements, quadrature based on the tensor product of well known one-dimensional Gauss–Legendre rules is readily implemented for any order of approximation. Note, however, that other quadrature formulae, with fewer integration points, exist [31, 32]. For triangles, specific quadrature rules, such as the symmetric quadrature proposed in [33, 34], are used. Analogously, efficient specific quadrature rules are used for tetrahedra, prisms and pyramids [34, 35].

4. Reducing the computational cost using affine elements

4.1. Comparison of affine and non-affine elements

The isoparametric transformation that relates the reference and physical coordinates is, in general, a non-linear mapping. This means that the computation of the element and face matrices requires the evaluation of the Jacobian of the transformation at each quadrature point. Assuming that the cost of evaluating the Jacobian is negligible compared to the cost of computing the element matrices, the number of operations required to compute an element mass matrix is $\mathbf{n}_{\text{en}}^2(4\mathbf{n}_{\text{eq}} - 1)$, where \mathbf{n}_{eq} is the number of quadrature points per element. The computation of the convection matrices requires $\mathbf{n}_{\text{en}}^2(\mathbf{n}_{\text{eq}}\mathbf{n}_{\text{sd}}(2\mathbf{n}_{\text{sd}} + 3) - 1)$ operations and the computation of a face mass matrices requires $\mathbf{n}_{\text{fn}}^2(4\mathbf{n}_{\text{fq}} - 1)$ operations, where \mathbf{n}_{fq} is the number of quadrature points per face.

However, when the physical element is affine with respect to the reference element, the isoparametric transformation is linear and the element matrices simplify to

$$\mathbf{M}_{ij} = |\mathbf{J}|\widehat{\mathbf{M}}_{ij} \quad \mathbf{C}_{ij}^k = |\mathbf{J}| \sum_{l=1}^{\mathbf{n}_{\text{sd}}} \mathbf{J}_{kl}^{-1} \widehat{\mathbf{C}}_{ij}^k \quad \mathbf{M}_{ij}^f = \sum_{f=1}^{\mathbf{n}_{\text{ef}}} \|\mathbf{J}^f\| \widehat{\mathbf{M}}_{ij}^f \quad (9)$$

where

$$\widehat{\mathbf{M}}_{ij} = \int_{\widehat{\Omega}} N_i N_j d\Omega$$

is the reference element mass matrix,

$$\widehat{\mathbf{C}}_{ij}^k = \int_{\widehat{\Omega}} N_i \frac{\partial N_j}{\partial \xi_k} d\Omega$$




Figure 3: Variation, with element order, of the ratio between the number of operations required by the standard and the quadrature-free implementations in computing, for different elements: (a) the element matrices; (b) the face mass matrix.

is the reference element convection matrix, in the local ξ_k direction, and

$$\widehat{\mathbf{M}}_{ij}^f = \int_{\widehat{\Gamma}} N_i N_j d\Gamma$$

is the reference face mass matrix. In this case, a quadrature-free implementation can be employed [24] in which the reference element and face matrices are pre-computed and scaled, using the constant Jacobian. This avoids the requirement for numerical integration during the time marching process. If all Jacobians are computed and stored, the quadrature-free implementation requires \mathbf{n}_{en}^2 operations for the computation of an element mass matrix, $\mathbf{n}_{\text{sd}} \mathbf{n}_{\text{en}}^2 (2\mathbf{n}_{\text{sd}} - 1)$ operations for the computation of the convection matrices and \mathbf{n}_{fn}^2 operations for the computation of a face mass matrix. Figure 3(a) illustrates the variation in the ratio of the total computational cost of the standard and the quadrature-free implementations, when computing the element matrices for different elements, with the order of approximation ranging from $p = 1$ up to $p = 7$. A similar comparison is shown in Figure 3(b) for the computation of the face mass matrix. In this figures, triangular, quadrilateral, tetrahedral and hexahedral elements are denoted by TRI, QUA, TET and HEX respectively. The comparison confirms that the standard implementation is much more expensive and that the quadrature-free implementation becomes more efficient as the order of the approximation

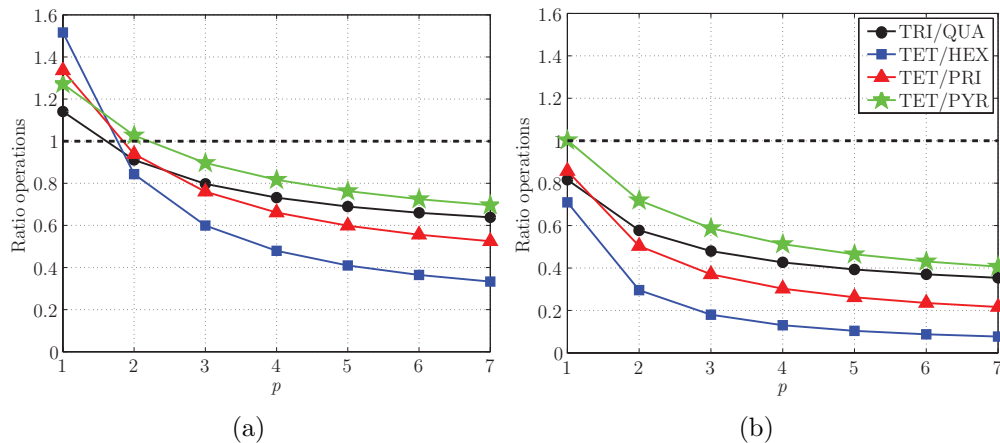


Figure 4: Variation, with element order, of the ratio of operations needed by different finite elements in order to (a) compute the element and face matrices; and (b) evaluate the residual of the ODE system of Equation (7).

increases, e.g. for linear hexahedra, the quadrature-free implementation reduces the number of operations by a factor of ten compared to the standard implementation, while this factor increases to 100, when the order of approximation is increased to $p = 4$, and to 450, for the case $p = 7$. In two dimensions, for both triangular and quadrilateral elements, the quadrature-free version is ten times faster for the case $p = 1$ and up to 100 times faster for an order of approximation $p = 7$.

For the quadrature-free implementation, the computational cost of the time marching process can be further reduced by storing the inverse of the Jacobian of the isoparametric mapping, the determinant of the Jacobian and the norm of the face Jacobians for each element, i.e. by storing $n_{\text{sd}}^2 + 2$ scalars per element. The required memory is, approximately, 50MB per million elements in two dimensions and between 90MB and 100MB per million elements in three dimensions. Of course, memory access can adversely affect the performance of the algorithm, if the information stored at the start of the computation does not fit into the cache memory. In addition, it is apparent that the computational cost, understood as the number of floating point operations, will not be directly related to cpu time, as the actual cpu time required is highly dependent on the computer architecture.

4.2. Hybrid meshes

The benefit of using affine elements can be clearly seen in Figure 3, by comparing the computational cost incurred when evaluating the element and face matrices. Triangular and tetrahedral meshes are preferred in this context, as the only requirement for triangles or tetrahedra to be affine is that their edges or faces, respectively, are planar. However, in the DG framework, for a given number of degrees of freedom, triangular and tetrahedral meshes possess more elements and more internal faces than quadrilateral or hexahedral meshes. Since a DG method defines the approximation element by element, the computation of numerical fluxes at internal faces represents an important part of the overall computational cost of the algorithm.

For simple geometries, the potential computational advantages of quadrilateral and hexahedral elements, with the quadrature-free implementation, can often be exploited. However, this is totally unfeasible for complex geometries. For this reason, we adopt the use of hybrid meshes, employing triangular or tetrahedral elements near objects of complex geometrical shape, with affine quadrilateral or hexahedral elements used to fill the remainder of the computational domain.

With the outer boundary of the computational domain taken in the form of a rectangle or a regular hexahedron, hybrid meshes of this form may be generated using a three stage procedure [36, 37, 38]. In the first stage, discretisation of the interior boundary curves or surfaces is performed [39] and the region inside the rectangle or hexahedron is discretised using a uniform structured mesh of square or cube elements. Squares or cubes within a prescribed distance of the scatterer, or lying internal to the scatterer, are removed in a second stage, to create a staircase-shaped surface or curve that completely encloses the scatterer. In three dimensions, a pyramid is constructed from each quadrilateral face of the staircase-shaped surface. In the third stage a Delaunay algorithm with automatic point insertion [40] is used to produce an unstructured mesh with the appropriate mesh spacing in the region between the staircase-shaped surface and the scatterer. It is worth noting that in a DG framework it is also possible to avoid the use of pyramids by using non-conforming meshes, but this approach is not considered here.

4.3. Operation count for different affine elements

A comparison of the number of operations required to compute the element and face matrices for the different finite elements with the quadrature-free implementation is displayed in Figure 4(a), where PRI and PYR refer to

prismatic and pyramidal elements respectively. This figure shows the variation in the ratio between the number of operations required by triangular and tetrahedral elements and the other elements, as the order of approximation changes from $p = 1$ to $p = 7$. In two dimensions, it is assumed that a triangular mesh is obtained from a quadrilateral mesh by splitting each quadrilateral into two triangles. In three dimensions, when comparing hexahedral and tetrahedral meshes, it is assumed that each hexahedron is split into six tetrahedra, that each prism is split into three tetrahedra and that each pyramid is split into two tetrahedra. The results reveal that, although triangular and tetrahedral meshes involve more elements and internal faces than meshes consisting of the other elements, they require fewer operations to compute the element and face matrices in the quadrature-free implementation. In two dimensions, only quadrilateral meshes with $p = 1$ require fewer operations than triangular meshes. A similar conclusion is reached when comparing hexahedra and prisms to tetrahedra. For pyramids, fewer operations are required not only with $p = 1$ and $p = 2$, but tetrahedra require fewer operations for higher order approximations.

A similar comparison is performed in Figure 4(b) for the number of operations required to evaluate the residual of the differential equation system of equation (7). For each component of the solution vector, the divergence term requires the computation of a matrix vector product of dimension \mathbf{n}_{en} for each of the \mathbf{n}_{sd} Cartesian directions, i.e. $\mathbf{n}_{\text{en}}^2(2\mathbf{n}_{\text{en}} - 1)$ operations. The contribution from the boundary integral requires the computation of the jump of the solution and a matrix vector product of dimension \mathbf{n}_{fn} for each solution component, i.e. $\mathbf{n}_{\text{fn}}(\mathbf{n}_{\text{fn}}(2\mathbf{n}_{\text{fn}} - 1) + 1)$ operations. The vectors obtained from the contribution of divergence and face terms are added to the residual vector, which is multiplied finally by the pre-computed inverse of the mass matrix. This requires an extra matrix vector product of dimension \mathbf{n}_{en} . The total number of operations per element is therefore $\mathbf{n}_{\text{en}}(2\mathbf{n}_{\text{en}}(2\mathbf{n}_{\text{en}} - 1) + 1)$ and the number of operations per face is $\mathbf{n}_{\text{fn}}(\mathbf{n}_{\text{fn}}(2\mathbf{n}_{\text{fn}} - 1) + 2)$. In this case, the results show that, for any order of approximation, the computation of the residual vector involves fewer operations when triangular or tetrahedral meshes are used.

It is important to remark that different approximation spaces are employed for each element, so different accuracy can be expected when using different elements with the same number of degrees of freedom. The numerical examples presented in Section 5 compare the cpu time and number of degrees of freedom required to achieve a desired accuracy. This compari-

son aims to identify if the expected extra accuracy introduced by the richer spaces of the quadrilateral or hexahedral elements, compared to triangular or tetrahedral elements, compensates for the extra computational cost observed in Figures 4(a) and (b).

5. Performance comparison

This section evaluates the efficiency of different elements employed for the numerical solution of the Maxwell's equations in two and three dimensions using a DG formulation. Simple test cases are used to show that the optimal rate of convergence is achieved and more complex scattering applications are used to compare the performance of the different elements. Although the hybrid meshes employed in the scattering examples do not involve prismatic elements, an analysis of their performance is presented for completeness. It is worth emphasising that all the examples considered here use the same order of approximation in all elements, but different orders could be easily handled within the adopted DG formulation.

5.1. PEC resonator

The first example involves the propagation of an electromagnetic wave inside a resonator consisting of two concentric PEC cylinders of diameter $d_1 = 1/3$ and $d_2 = 1$. The exact solution of this problem is known [41] and is used to enable the evaluation of the relative error in the $\mathcal{L}^2(\Omega)$ norm in the solution computed for the TM mode. An h -convergence study is performed initially, to verify the optimal convergence of the DG code using high-order triangular and quadrilateral elements. Five meshes are considered for each element type, with an order of approximation ranging from $p = 1$ to $p = 5$. The five quadrilateral meshes are displayed in Figure 5 for $p = 1$. The triangular meshes are obtained by splitting each quadrilateral in these meshes into two triangles. The corresponding high-order meshes are generated by placing the boundary nodes on the true geometry and adapting the internal nodes of curved elements to obtain an optimal isoparametric nodal distribution [42].

Figure 6 show the $\mathcal{L}^2(\Omega)$ norm of the relative error in the third component of the electric field, E_3 , as a function of the characteristic mesh size h . Each symbol represents an order of approximation and each point represents a level of mesh refinement. A solid line is used for triangular meshes and a dashed line for quadrilateral meshes. The optimal convergence rate of




Figure 5: The five quadrilateral meshes, with $p = 1$, employed for the PEC resonator with (a) 16 elements; (b) 64 elements; (c) 256 elements; (d) 1024 elements; (e) 4096 elements.

$\mathcal{O}(h^{p+1})$ is obtained in all cases. It is observed that, for the same characteristic mesh size, the additional terms that appear in the approximation space, \mathbb{Q}_p , for quadrilateral elements, with respect to the approximation space, \mathbb{P}_p , used for triangular elements, introduce significant advantages when the coarsest mesh is considered but, as the mesh is refined, this advantage disappears. On the coarsest mesh, the accuracy obtained with quadrilateral elements with $p = 4$ is almost the same as that obtained with triangular elements with $p = 5$.

Triangular and quadrilateral meshes can also be compared in terms of the number of degrees of freedom (n_{dof}) and the cpu time required to reach a desired accuracy. It is important to note here that the quadrilateral meshes contain only non-affine elements, whereas the triangular meshes contain both non-affine elements and affine elements. This means that this comparison is aimed at illustrating the superiority of the quadrature-free implementation.

The comparison in terms of n_{dof} is displayed in Figure 7(a), which shows that quadrilateral elements are more efficient only for coarse levels of mesh refinement. Obviously, for the same characteristic mesh size, h , triangular meshes have more degrees of freedom than quadrilateral meshes because, in a DG framework, the interior edge nodes are duplicated when a quadrilateral is split into two triangles. A comparison of the variation in the cpu time required to reach a desired accuracy is represented in Figure 7(b). For coarse meshes, quadrilateral elements not only reduce the number of degrees of freedom but also reduce the cpu time required to reach a desired accuracy when compared to triangular elements. For approximations of order $p = 1$ to $p = 3$, triangular and quadrilateral meshes offer the same performance as the mesh is refined but, for higher order approximations, triangular meshes turn out to be more efficient. The percentage of affine elements in triangular




Figure 6: For the PEC resonator, h -convergence, on triangular (TRI) and quadrilateral (QUA) meshes, of the $\mathcal{L}^2(\Omega)$ norm of the relative error in H_1 .

meshes clearly dominates as the mesh is refined and the quadrature-free implementation shows better performance, e.g. at the fourth level of mesh refinement, the accuracy obtained by quadrilateral and triangular elements with $p = 5$ is almost identical, while the triangular mesh requires four times less cpu time.

As expected for a problem with a smooth solution, the linear approximation is clearly outperformed by higher order approximations, not only in terms of a reduction in the number of degrees of freedom but, more importantly, in terms of the cpu time required to reach a desired accuracy, e.g. the cpu time required by quadrilateral elements with $p = 1$ to reach an accuracy of 5×10^{-3} is more than 140 times higher than the cpu time required with $p = 2$ and more than 860 times higher than the cpu time required with $p = 3$.

5.2. Scattering by a 15λ PEC cylinder

The next example considers the scattering, by a circular PEC cylinder, of an incident plane TM wave, generated by a source located in the far field. The diameter of the cylinder is 15λ , where λ is the wavelength of the incident wave. This problem has an analytical solution [43] and the relative error of

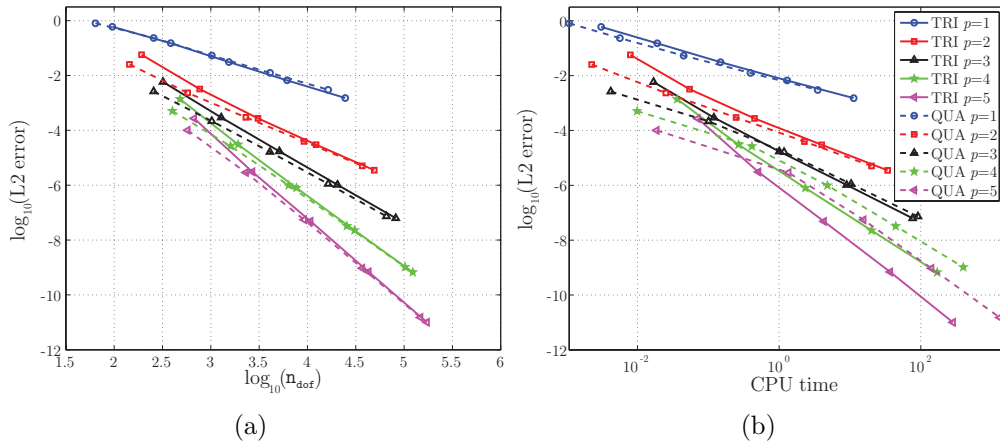


Figure 7: PEC resonator: $\mathcal{L}^2(\Omega)$ norm of the relative error in H_1 as a function of (a) the number of degrees of freedom; (b) the cpu time.

the scattering width (RCS) [26] is used to measure the accuracy of numerical simulation. The truncated boundary, forming the interface between the free-space and the absorbing PML, is taken to be a square. The outer boundary of the PML is also taken to be a square and the thickness of the PML is set equal to λ . A set of hybrid meshes is generated, composed of triangular elements near the scatterer and affine quadrilaterals elsewhere. Figure 8 shows the first five, of a series of 14, meshes employed, with the order of approximation ranging from $p = 1$ to $p = 7$.




Figure 8: Scattering of a plane TM wave by a PEC circular cylinder: the first five hybrid meshes used to discretise the truncated domain and the PML.

The main properties of a representative selection of five of the 14 meshes, together with the properties of the triangular meshes formed by subdivision of each quadrilateral into two triangles, are summarised in Table 2. The number

of elements, n_{e1} , and the percentage of affine elements, n_{e1}^{Aff} , is given for both the triangular and the hybrid meshes and the percentage of quadrilateral elements, n_{e1}^Q , is also detailed for the hybrid meshes.

<i>Mesh</i>	n_{e1}	% n_{e1}^{Aff}	% n_{e1}^Q	<i>Mesh</i>	n_{e1}	% n_{e1}^{Aff}
HYB 1	152	82	50	TRI 1	228	88
HYB 2	516	88	71	TRI 2	884	93
HYB 5	3076	95	88	TRI 5	5788	97
HYB 9	9744	97	93	TRI 9	18840	99
HYB 14	23352	98	96	TRI 14	45700	99

Table 2: Scattering of a plane TM wave by a PEC circular cylinder: properties of five representative hybrid and triangular meshes.

To compare the performance of triangular and hybrid meshes, Figure 9 (a) shows how the ratio between the n_{dof} required with triangular meshes and the n_{dof} required with hybrid meshes, to reach relative errors of 10^{-1} , 10^{-2} , 10^{-3} and 10^{-4} in RCS, varies as the element order is increased. Each point in this figure, corresponding to a given order of approximation, is obtained by dividing the cpu time required by the first triangular mesh and by the first hybrid mesh which achieve the desired accuracy. It is important to remark that triangular and hybrid meshes may achieve the required accuracy with different levels of mesh refinement. This means that the ratios shown in this figure are not necessarily obtained with the same level of mesh refinement.

For lower order approximations, hybrid meshes are more competitive, e.g. to achieve an error of 10^{-1} with $p = 1$, hybrid meshes use approximately 70% of the degrees of freedom required by triangular meshes. As the order of the approximation is increased, this difference becomes less significant, e.g. with $p = 4$ and $p = 5$, hybrid meshes require 90% of the number of degrees of freedom employed by triangular meshes. For $p > 6$, the number of degrees of freedom required by triangular meshes to give an error of 10^{-1} is slightly lower than the number required using hybrid meshes. This shows that the advantage of using hybrid meshes is highly dependent upon the order of the approximation employed. A similar comparison, in terms of the cpu time required, may be gained from Figure 9 (b).




Figure 9: Scattering of a plane TM wave by a PEC circular cylinder showing how the accuracy level achieved varies with the element order and: (a) the ratio between the number of degrees of freedom with triangular meshes and the number required with quadrilateral meshes; (b) the ratio between cpu time required by triangular and quadrilateral meshes.

The results show that higher order approximations reduce substantially both the n_{dof} and the cpu time required to reach a desired accuracy when compared to lower order elements. It is worth emphasising that, for this example, this is not only true for high fidelity computations but also for lower desired accuracy levels, e.g. to achieve an error of just 10^{-1} , using hybrid meshes and $p = 1$, the computation requires 203 904 degrees of freedom, whereas only 49 068 degrees of freedom are required with $p = 5$. This reduction in the number of unknowns is also accompanied by a significant reduction in cpu time, as the computation with $p = 5$ is 90 times faster than the computation with $p = 1$. The results also reveal that the cpu time required to obtain a relative error in the RCS of between 10^{-2} and 10^{-3} does not change significantly when the order of approximation is increased from $p = 5$ to $p = 7$. This suggests that there is an optimum value for the order of the approximation, if the minimum cpu time required to compute the solution with a given accuracy is of interest.

5.3. Scattering by a 100λ PEC NACA0012 aerofoil

The next example involves the scattering of a plane incident electromagnetic TE wave by a PEC NACA0012 aerofoil. The chord length of the aerofoil is equal to 100λ . The singularity introduced into the solution at the sharp

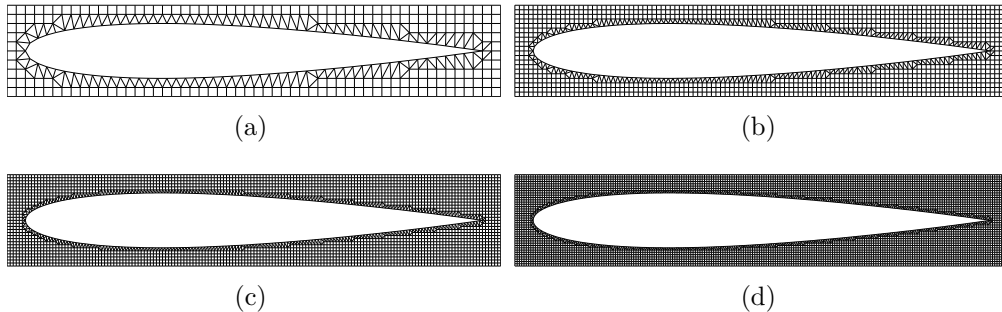


Figure 10: Scattering of a plane TE wave by a PEC NACA0012 aerofoil: the first four hybrid meshes used to discretise the truncated domain and the PML.

trailing edge, the relatively high frequency of the incident wave and the variation in the curvature of the aerofoil near the leading edge make the accurate computation of the scattered field a challenging task for this problem. As an analytical solution is not available, a reference numerical solution is used to measure the accuracy of the computed results. The reference solution, using a spatial discretisation with approximately 35 nodes per wave length, is computed on a mesh with more than 40 000 elements and an order of a approximation $p = 7$. Solutions computed on an additionally refined mesh produced variations of the order of 10^{-6} in the RCS, which is almost two orders of magnitude lower than the errors in the computations that are compared in this example. A set of 11 hybrid meshes are used, with the corresponding triangular meshes obtained by splitting each quadrilateral into two triangles. Figure 10 shows the first four hybrid meshes and the main properties of some of the 11 meshes, and the corresponding triangular meshes, are summarised in Table 3. Figure 11 shows the contours of the transverse component of the magnetic field that has been computed on the second mesh using an order of a approximation $p = 7$. The estimated error in the RCS is 1×10^{-2} using less than one element per wavelength.




Figure 11: Scattering of a plane TE wave by a PEC NACA0012 aerofoil: numerical solution, H_3 , computed on the second hybrid mesh with $p = 7$.

<i>Mesh</i>	n_{e1}	$\% n_{e1}^{Aff}$	$\% n_{e1}^Q$	<i>Mesh</i>	n_{e1}	$\% n_{e1}^{Aff}$
HYB 1	412	76	47	TRI 1	604	83
HYB 2	1498	87	71	TRI 2	2562	92
HYB 5	8774	94	88	TRI 5	16472	97
HYB 8	22130	96	92	TRI 8	42542	98
HYB 11	41476	97	94	TRI 11	80592	99

Table 3: Scattering of a plane TE wave by a PEC NACA0012 aerofoil: properties of five representative hybrid and triangular meshes.

To compare the performance of triangular and hybrid meshes, Figure 12 (a) shows the variation in the ratio between the n_{dof} required by a triangular mesh and the n_{dof} required by a hybrid mesh in order to reach relative accuracy levels of 10^{-1} , 10^{-2} , 10^{-3} and 10^{-4} in the RCS, as the element order is increased. The comparison in terms of the cpu time required is displayed in Figure 12 (b). The use of hybrid meshes is preferable for low order approxi-




Figure 12: Scattering of a plane TE wave by a PEC NACA0012 aerofoil showing how the accuracy level varies with the element order and the ratio of (a) number of degrees of freedom and (b) cpu time required by triangular and hybrid meshes.

mations whereas higher order approximations require less degrees of freedom

and cpu time when using triangular meshes. It is worth remarking that triangular meshes are computationally more efficient for $p > 5$ for both examples, but the differences in cpu time are less important for this example.

Even though the solution to this example includes a singularity at the trailing edge, higher order approximations are able to reduce the n_{dof} and the cpu time required to reach a desired accuracy compared to lower order approximations. With linear elements, it is very difficult to converge the RCS to an acceptable level of accuracy, e.g. the error in the RCS computed on the finest mesh with $p = 1$ and 490 632 degrees of freedom is more than 10^{-1} , whereas equivalent accuracy can be obtained with $p = 7$ and just 60 624 degrees of freedom. This implies a factor of seven reduction in the cpu time. The efficiency of high-order approximations in this example is mainly attributed to two reasons. Firstly, dispersion and dissipation errors in wave propagation problems are known to be best overcome by using higher order approximations [11]. Secondly, accurate geometry representations are crucial in the computation of accurate RCS patterns [44] and approximating the boundary with linear elements introduces singularities into the numerical solution that can dramatically affect the accuracy of the computed solution [45].

5.4. Electromagnetic wave propagation in a cube

This three dimensional example involves the propagation of an electromagnetic wave inside a cube. The exact solution of this problem is taken as in [46] and is used to determine the relative error in the $\mathcal{L}^2(\Omega)$ norm of the computed solution. Initially, an h -convergence study is performed to verify the optimal convergence of the three dimensional DG code, using high-order tetrahedral, hexahedral, prismatic and pyramidal elements. For each element type, five structured meshes, containing only affine elements, are considered and the computation is performed for an order of approximation ranging from $p = 1$ up to $p = 5$. Figure 13 shows the variation in the error in H_3 as the characteristic mesh size h changes. Each symbol represents an order of approximation, each point represents a level of mesh refinement and different line styles are used for tetrahedra, hexahedra, prisms and pyramids, denoted by TET, HEX, PRI and PYR respectively. The optimal convergence rate of $\mathcal{O}(h^{p+1})$ is obtained in each case. For the same characteristic mesh size, the extra terms appearing in the approximation spaces, for hexahedral, prismatic and pyramidal elements, compared to the approximation space used for tetrahedral elements, introduce significant advantages for both coarse and




Figure 13: Wave propagation in a cube: h -convergence of the $\mathcal{L}^2(\Omega)$ norm of the relative error in H_3 .

fine meshes. For a given characteristic mesh size, the least accurate approximation is provided by tetrahedral elements. Similar accuracy is obtained by prismatic and pyramidal elements, and the most accurate results are provided by hexahedral elements. For higher order approximations, the difference in accuracy between the elements used here is significant. Hexahedra provide more accurate results than tetrahedra, with the improvement being less than one order of magnitude for $p = 1$ and about two orders of magnitude for $p = 5$.

To compare the performance of the different elements, Figure 14 shows the ratio between the n_{dof} required by tetrahedra and hexahedra, prisms and pyramids, to reach a relative error in the $\mathcal{L}^2(\Omega)$ norm of 10^{-1} , 10^{-2} , 10^{-3} and 10^{-4} . The results clearly indicate that, in all cases, hexahedral, prismatic and pyramidal elements are able to reduce the n_{dof} to reach a required accuracy compared to tetrahedral elements and the best performance is always obtained by hexahedral elements. In addition, the results reveal that the benefit of using hexahedral, prismatic or pyramidal elements instead of tetrahedral elements is more pronounced as the order of the approximation is increased. For instance, to achieve an accuracy of 10^{-2} , linear hexahedral




Figure 14: Wave propagation in a cube showing how the accuracy level varies with the element order and the ratio of number of degrees of freedom required by tetrahedral and (a) hexahedral, (b) prismatic and (c) pyramidal meshes.

meshes need four times less degrees of freedom than linear tetrahedral meshes and, for the same accuracy, hexahedral meshes with $p = 5$ reduce the number of degrees of freedom by a factor of 25 compared to tetrahedral meshes.

The comparison in terms of cpu time is displayed in Figure 15. It is ob-




Figure 15: Wave propagation in a cube showing how the accuracy level varies with the element order and the ratio of cpu time required by tetrahedral and (a) hexahedral, (b) prismatic and (c) pyramidal meshes.

served that the reduction in n_{dof} is translated into a corresponding reduction in computational cost. For an order of approximation up to $p = 4$, hexahedral elements provide the same accuracy as tetrahedral elements by reducing the cpu time by a factor of at least ten. For an order of approximation $p = 5$, hexahedral elements remain more efficient, reducing the cpu time by a factor of a seven. Similar performance is observed for prismatic and pyramidal finite elements, although the former are slightly more efficient. It is worth

remarking on the different behavior obtained in two and three dimensions. In two dimensional examples, meshes of quadrilateral elements have been shown to be marginally better in terms of cpu time than meshes of triangular elements. However, in this three dimensional example the improvement in cpu time attained by using hexahedral elements is more than a factor of five.

Finally, the performance of low and higher order approximations is briefly discussed. Again, as expected for a problem with smooth solution, the use of higher order approximations implies that a desired accuracy can be obtained with fewer degrees of freedom compared to low order approximations. For instance, the computation with tetrahedral linear elements in the finest mesh, with almost 5 million degrees of freedom, provides an error comparable to the computation with tetrahedral elements and $p = 4$ in the first mesh, with less than 25 000 degrees of freedom. Similar conclusions are obtained for other finite elements, e.g. to achieve an error of 10^{-3} , hexahedral elements with $p = 2$ employ almost 100 times more degrees of freedom than hexahedral elements with $p = 5$. The results also show that higher order approximations are able to reduce the cpu time required to reach a desired accuracy when compared to lower order approximations. For instance, the cpu time required to achieve an accuracy of 10^{-2} can be reduced by a factor of 30 when using tetrahedral elements with $p = 3$ compared to linear tetrahedral elements. Similar performance is obtained when comparing other linear elements with their corresponding higher order equivalents, but it is apparent that there is an optimal order of approximation depending on the desired accuracy. For instance, to achieve an error of 5×10^{-2} , quadratic hexahedral elements are slightly more efficient than cubic hexahedral elements, but for higher fidelity results, say 5×10^{-3} , cubic elements are more efficient. However, it is important to remark that for any degree of accuracy, even for very low accuracy levels, the use of higher order elements (i.e. $p > 1$) is advantageous in terms of reducing the n_{dof} and, more importantly, the cpu time for a required accuracy. This example also gives an indication of the maximum speed up that can be achieved when using hybrid meshes instead of pure tetrahedral meshes.

5.5. Scattering by a 2λ PEC sphere

The final example involves the scattering of a plane incident electromagnetic wave by a PEC sphere of diameter 2λ . This problem has an analytical solution [43] and the relative error of the radar cross section (RCS) is used

to measure the accuracy of numerical simulation. The truncated boundary, forming the interface between the free-space and the absorbing PML, is taken to be a cube. The outer boundary of the PML is also taken to be a cube and the thickness of the PML is set equal to λ . Using the procedure described in Section 4.2, two hybrid meshes are generated, composed of tetrahedral elements near the scatterer, a transition region of affine pyramids and affine hexahedrons elsewhere. The unstructured curved portion of the mesh is obtained following the procedure described in [47].

Figure 16 shows a view of the two hybrid meshes, where the transition between tetrahedral and hexahedral elements using pyramids is displayed. The first mesh has 3 834 tetrahedra, 5 232 hexahedra and 452 pyramids and the second mesh has 25 339 tetrahedra, 42 769 hexahedra and 1 596 pyramids. It is anticipated that the savings obtained when using hybrid meshes for electromagnetic scattering problems will be related to the percentage of tetrahedral and hexahedral elements in the mesh. In the first mesh 40% of the elements are tetrahedra while 36% of the elements are tetrahedral in the second mesh. To compare the performance of tetrahedral and hybrid meshes, two tetrahedral meshes are generated by considering the hybrid meshes of Figure 16 and splitting each pyramid into two tetrahedrons and each hexahedra into six tetrahedra. The total number of elements in the corresponding tetrahedral meshes is 36 130 and 285 145 respectively.

As expected, the results reveal that the accuracy of the numerical solution computed on the hybrid mesh is almost identical to the accuracy on the corresponding tetrahedral mesh. As the error is measured in the RCS, the accuracy is controlled by the solution in the unstructured region, which is common to both tetrahedral and hybrid meshes. For instance, on the coarsest mesh and with $p = 1$, the \mathcal{L}^2 norm of the RCS error in the horizontal plane is 8.0×10^{-2} for the hybrid mesh and 8.6×10^{-2} for the tetrahedral mesh. With hybrid meshes, the cpu time is reduced by a factor of 2.8 and the degrees of freedom by a factor of 2.4. Similar conclusions are reached when the order of the approximation is increased. Using $p = 3$ in the coarsest mesh, the \mathcal{L}^2 norm of the RCS error in the horizontal plane is 2.6×10^{-4} for both the hybrid and the tetrahedral meshes. With hybrid meshes, the cpu time is reduced by a factor of 1.8 and the degrees of freedom by a factor of 1.7. Figure 17 shows the contours of the first component of the electric field, E_1 , computed on the coarse hybrid mesh of Figure 16 using an order of approximation $p = 4$.

Similar conclusions are obtained for the second mesh. With $p = 1$, the




Figure 16: Scattering of a plane wave by a PEC sphere: a view of the two hybrid meshes used to discretise the truncated domain and the PML.

\mathcal{L}^2 norm of the RCS error in the horizontal plane is 2.3×10^{-2} for the hybrid mesh and 2.4×10^{-2} for the tetrahedral mesh. With hybrid meshes, the cpu time is reduced by a factor of 2.8 and the degrees of freedom by a factor of 2.5. Similar conclusions are derived for $p = 2$. The \mathcal{L}^2 norm of the RCS error in the horizontal plane is 1.5×10^{-4} for both the hybrid and the tetrahedral meshes. With hybrid meshes, the cpu time is reduced by a factor of 2.2 and the degrees of freedom by a factor of 2.0.

It is worth noting that a large difference in cpu time saving is observed when comparing this example to the previous example. This is due to the use of hybrid meshes instead of meshes with only one element type. With 40% of the elements in the coarsest hybrid mesh being tetrahedra, the cpu time is mainly controlled by the computations required on the unstructured portion of the mesh. It is apparent that the saving in cpu time associated with the use of hybrid meshes is highly dependent on the geometry under consideration. For geometries such as cylinders or spheres, where the maximum and minimum distance from a point in the scatterer to the PML region




Figure 17: Scattering of a plane wave by a PEC sphere: numerical solution, E_1 , computed on the coarse hybrid mesh with $p = 4$.

is similar, the percentage of elements in the unstructured part of the mesh is expected to be very high. For geometries in which a significant variation on the distance to the PML region is expected, such as a complete aircraft configuration, this percentage is expected to be much lower. This means that, for complex configurations the saving in cpu time will be significantly higher as many more hexahedral elements will be used to fill the truncated domain.

6. Concluding remarks

The efficiency of a high-order DG solver for the solution of Maxwell's equations in the time domain using hybrid meshes has been studied. Geometric flexibility is achieved by using an unstructured mesh around complex geometric objects and computational efficiency is then improved by using a

Cartesian mesh of affine quadrilaterals or hexahedrons to fill the remainder of the domain. The proposed approach benefits from the efficient quadrature-free implementation in the majority of the domain as all elements, apart from a number curved elements, are affine.

The optimal rate of convergence has been verified, for triangular and quadrilateral elements in two dimensions and for tetrahedral, prismatic, pyramidal and hexahedral elements in three dimensions, using numerical examples with smooth analytical solutions. The efficiency of using hybrid meshes has been analysed using more complex electromagnetic scattering applications in two and three dimensions. The results show that the saving in cpu time is highly dependent on the order of the approximation and the percentage of Cartesian elements in the hybrid mesh. In two dimensions, it has been found that for low to moderate degrees of approximation, i.e. $p < 5$, hybrid meshes are able to produce the same accuracy with a small reduction in cpu time and in the number of degrees of freedom. For very high-order approximations, the results indicate that it is best to consider triangular meshes. In three dimensions, the performance improvement offered by hybrid meshes is more significant, as the cpu time and the number of degrees of freedom can be reduced by a factor of two. The saving in cpu time is directly related to the percentage of tetrahedral elements in the hybrid mesh and, therefore, to the geometry under consideration. When a hexahedral mesh is compared to a tetrahedral mesh, it has been found that, to obtain the same accuracy, hexahedral elements are between seven to 15 times more efficient than tetrahedral elements.

References

References

- [1] K. S. Yee, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media, IEEE Trans. Antennas Propagat. 14 (3) (1966) 302–307.
- [2] D. J. Riley, C. D. Turner, VOLMAX: a solid-model-based, transient volumetric Maxwell solver using hybrid grids, IEEE Trans. Antennas Propagat. 39 (1) (1997) 20–33.
- [3] K. S. Yee, J. S. Chen, The finite-difference time-domain (FDTD) and the finite-volume time-domain (FVTD) methods in solving Maxwell's equations, IEEE Trans. Antennas Propagat. 45 (3) (1997) 354–363.

- [4] M. El hachemi, O. Hassan, K. Morgan, D. P. Rowse, N. P. Weatherill, Hybrid methods for electromagnetic scattering simulations on overlapping grids, *Comm. Numer. Methods Engrg.* 19 (9) (2003) 749–760.
- [5] K. Morgan, M. El hachemi, O. Hassan, N. P. Weatherill, Explicit time domain finite element methods for electromagnetics, in: D. A. Caughey, M. M. Hafez (Eds.), *Frontiers of Computational Fluid Dynamics*, 2006, World Sci. Publishing, 2006, pp. 161–181.
- [6] R. Sevilla, S. Fernández–Méndez, A. Huerta, NURBS–enhanced finite element method (NEFEM), *Internat. J. Numer. Methods Engrg.* 76 (1) (2008) 56–83.
- [7] D. Xue, L. Demkowicz, Control of geometry induced error in *hp* finite element (FE) simulations. I. Evaluation of FE error for curvilinear geometries, *Internat. J. Numer. Anal. Model.* 2 (3) (2005) 283–300.
- [8] F. Edelvik, G. Ledfelt, Explicit hybrid time domain solver for the Maxwell equations in 3D, *J. Sci. Comp.* 15 (1) (2000) 61–78.
- [9] R. Sevilla, S. Fernández–Méndez, A. Huerta, Comparison of high–order curved finite elements, *Internat. J. Numer. Methods Engrg.* 87 (8) (2011) 719–734.
- [10] R. Sevilla, S. Fernández–Méndez, A. Huerta, 3D NURBS–enhanced finite element method (NEFEM), *Internat. J. Numer. Methods Engrg.* 88 (2) (2011) 103–125.
- [11] M. Ainsworth, Dispersive and dissipative behaviour of high order discontinuous Galerkin finite element methods, *J. Comput. Phys.* 198 (1) (2004) 106–130.
- [12] J. S. Hesthaven, High–order accurate methods in time–domain computational electromagnetics: A review, *Adv. Imag. Electron Phys.* 127 (2003) 59–123.
- [13] J. S. Hesthaven, T. Warburton, Nodal high–order methods on unstructured grids I. Time–domain solution of Maxwell's equations, *J. Comput. Phys.* 181 (1) (2002) 186–221.

- [14] B. Cockburn, F. Li, C.-W. Shu, Locally divergence-free discontinuous Galerkin methods for the Maxwell equations, *J. Comput. Phys.* 194 (2004) 588–610.
- [15] T. Lu, P. W. Zhang, W. Cai, Discontinuous Galerkin methods for dispersive and lossy Maxwell's equations and PML boundary conditions, *J. Comp. Phys* 200 (2) (2004) 549–580.
- [16] A. V. Kabakian, V. Shankar, W. F. Hall, Unstructured grid-based discontinuous Galerkin method for broadband electromagnetic simulations, *J. Sci. Comput.* 20 (3) (2004) 405–431.
- [17] M.-H. Chen, B. Cockburn, F. Reitich, High-order RKDG methods for computational electromagnetics, *J. Sci. Comp.* 22 (1) (2005) 205–226.
- [18] X. Ji, W. Cai, P. Zhang, High-order DGTD methods for dispersive Maxwell's equations and modelling of silver nanowire coupling, *Internat. J. Numer. Methods Engrg.* 69 (2007) 308–325.
- [19] H. Fahs, Development of a *hp*-like discontinuous Galerkin time-domain method on non-conforming simplicial meshes for electromagnetic wave propagation, *Internat. J. Numer. Anal. Model.* 6 (2) (2009) 193–216.
- [20] H. Fahs, L. Fezoui, S. Lanteri, V. Dolean, F. Rapetti, Recent achievements on a DGTD method for time-domain electromagnetics, *IEEE Trans. Magn.* 46 (8) (2010) 3061–3064.
- [21] M. König, K. Busch, J. Niegemann, The discontinuous Galerkin time-domain method for Maxwell's equations with anisotropic materials, *Photonics Nanostruct.* 8 (2010) 303–309.
- [22] X. Roca, Paving the path towards automatic hexahedral mesh generation, Ph.D. thesis, Universitat Politècnica de Catalunya (2009).
- [23] E. Ruiz-Gironés, Automatic hexahedral meshing algorithms: from structured to unstructured meshes, Ph.D. thesis, Universitat Politècnica de Catalunya (2011).
- [24] H. L. Atkins, C. W. Shu, Quadrature-free implementation of discontinuous Galerkin method for hyperbolic equations, *AIAA Journal* 36 (5) (1998) 775–782.

- [25] C. A. Balanis, *Advanced Engineering Electromagnetics*, John Wiley and Sons, New York, 1989.
- [26] A. Taflov, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, Artech House, Inc., 1995.
- [27] J. Donea, A. Huerta, *Finite Element Methods for Flow Problems*, Wiley, 2005.
- [28] M. Bergot, G. Cohen, M. Duruflé, Higher-order finite elements for hybrid meshes using new nodal pyramidal elements, *J. Sci. Comput.* 42 (3) (2010) 345–381.
- [29] M. A. Taylor, B. A. Wingate, R. E. Vincent, An algorithm for computing Fekete points in the triangle, *SIAM J. Numer. Anal.* 38 (5) (2000) 1707–1720.
- [30] H. Luo, C. Pozrikidis, A Lobatto interpolation grid in the tetrahedron, *IMA J. Appl. Math.* 71 (2) (2006) 298–313.
- [31] D. A. Dunavant, Economical symmetrical quadrature rules for complete polynomials over a square domain, *Internat. J. Numer. Methods Engrg.* 21 (10) (1985) 1777–1784.
- [32] D. A. Dunavant, Efficient symmetrical cubature rules for complete polynomials of high degree over the unit cube, *Internat. J. Numer. Methods Engrg.* 23 (10) (1986) 397–407.
- [33] S. Wandzura, H. Xiao, Symmetric quadrature rules on a triangle, *Comput. Math. Appl.* 45 (12) (2003) 1829–1840.
- [34] L. Zhang, T. Cui, H. Liu, A set of symmetric quadrature rules on triangles and tetrahedra, *J. Comput. Math.* 27 (1) (2009) 89–96.
- [35] P. Solin, K. Segeth, *Higher-Order Finite Element Methods*, Chapman & Hall, 2003.
- [36] R. Davies, K. Morgan, O. Hassan, A high order hybrid finite element method applied to the solution of electromagnetic wave scattering problems in the time domain, *Comp. Mech.* 44 (3) (2009) 321–331.

- [37] M. El hachemi, O. Hassan, K. Morgan, D. Rowse, N. P. Weatherill, A low-order unstructured-mesh approach for computational electromagnetics in the time domain, *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* 362 (1816) (2004) 445–469.
- [38] O. H. Z. Q. Xie, K. Morgan, Tailoring unstructured meshes for use with a 3D time domain co-volume algorithm for computational electromagnetics, *Internat. J. Numer. Methods Engrg.* 87 (1–5) (2011) 48–65.
- [39] J. Peiró, *Handbook of grid generation*, CRC Press, 1999, Ch. Surface grid generation, pp. 19.1–19.20.
- [40] N. P. Weatherill, O. Hassan, Efficient three-dimensional Delaunay triangulation with automatic point creation and imposed boundary constraints, *Internat. J. Numer. Methods Engrg.* 37 (12) (1994) 2005–2039.
- [41] K. Dridi, J. Hesthaven, A. Ditkowski, Staircase-free finite-difference time-domain formulation for general materials in complex geometries, *IEEE Trans. Antennas Propagat.* 49 (5) (2001) 749–756.
- [42] C. Bernardi, Optimal finite-element interpolation on curved domains, *SIAM J. Numer. Anal.* 26 (5) (1989) 1212–1240.
- [43] R. F. Harrington, *Time-Harmonic Electromagnetic Fields*, McGraw-Hill, New York, 1961.
- [44] R. Sevilla, S. Fernández-Méndez, A. Huerta, NURBS-Enhanced Finite Element Method (NEFEM): a seamless bridge between CAD and FEM, *Arch. Comput. Methods Engrg.* 18 (4) (2011) 441–484.
- [45] R. Sevilla, S. Fernández-Méndez, A. Huerta, NURBS-enhanced finite element method (NEFEM) for Euler equations, *Internat. J. Numer. Methods Fluids* 57 (9) (2008) 1051–1069.
- [46] J. Lee, B. Fornberg, A split step approach for the 3-D Maxwell's equations, *J. Comput. Appl. Math.* 158 (2) (2003) 485–505.
- [47] Z. Q. Xie, R. Sevilla, O. Hassan, K. Morgan, The generation of arbitrary order curved meshes for 3D finite element analysis, *Comp. Mech.* DOI: 10.1007/s00466-012-0736-4.