




Contents lists available at ScienceDirect

## Computers and Mathematics with Applications

journal homepage: [www.elsevier.com/locate/camwa](http://www.elsevier.com/locate/camwa)

## Fast native-MATLAB stiffness assembly for SIPG linear elasticity

R.E. Bird<sup>\*</sup>, W.M. Coombs, S. Giani

Durham University, School of Engineering and Computing Sciences, South Rd, Durham, DH1 3LE, United Kingdom

## ARTICLE INFO

## Article history:

Received 6 July 2016

Received in revised form 18 July 2017

Accepted 12 August 2017

Available online 9 September 2017

## Keywords:

Efficient

MATLAB

Stiffness matrix

Symmetric interior penalty

Discontinuous Galerkin

Linear elasticity

## ABSTRACT

When written in MATLAB the finite element method (FEM) can be implemented quickly and with significantly fewer lines, when compared to compiled code. MATLAB is also an attractive environment for generating bespoke routines for scientific computation as it contains a library of easily accessible inbuilt functions, effective debugging tools and a simple syntax for generating scripts. However, there is a general view that MATLAB is too inefficient for the analysis of large problems. Here this preconception is challenged by detailing a vectorised and blocked algorithm for the global stiffness matrix computation of the symmetric interior penalty discontinuous Galerkin (SIPG) FEM. The major difference between the computation of the global stiffness matrix for SIPG and conventional continuous Galerkin approximations is the requirement to evaluate inter-element face terms, this significantly increases the computational effort. This paper focuses on the face integrals as they dominate the computation time and have not been addressed in the existing literature. Unlike existing optimised finite element algorithms available in the literature the paper makes use of only native MATLAB functionality and is compatible with GNU Octave. The algorithm is primarily described for 2D analysis for meshes with homogeneous element type and polynomial order. The same structure is also applied to, and results presented for, a 3D analysis. For problem sizes of  $10^6$  degrees of freedom (DOF), 2D computations of the local stiffness matrices were at least  $\approx 24$  times faster, with 13.7 times improvement from vectorisation and a further 1.8 times improvement from blocking. The speed up from blocking and vectorisation is dependent on the computer architecture, with the range of potential improvements shown for two architectures in this paper.

© 2017 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Finite element analysis (FEA) is commonly used as a technique for solving partial differential equations by engineers, mathematicians and scientists. The MATLAB environment, with its library of functions and debugging procedures, allows bespoke FEA routines to be generated quickly with few lines. Examples include Coombs et al. [1], Sigmund [2] and others. However, an unoptimised MATLAB script will often run significantly slower than unoptimised compiled code [3]. This paper demonstrates how the advantage of using only native MATLAB to generate FEA routines is not necessarily penalised with slow run times when written in an optimised form for the symmetric interior penalty Galerkin method (SIPG).

Significant progress on optimising FEA routines in MATLAB was achieved by Dabrowski et al. [3] in 2008. The authors presented MILAMIN, an open source optimised non-native MATLAB implementation of continuous Galerkin (CG) FEA code that is capable of setting up, solving, and post processing 2D unstructured mesh problems with  $10^6$  degrees of freedom (DOF)

<sup>\*</sup> Corresponding author.

E-mail address: [robert.e.bird@durham.ac.uk](mailto:robert.e.bird@durham.ac.uk) (R.E. Bird).

in under a minute. One common method to compute the local element stiffness is to compute each local element matrix by turn through a series of small matrix multiplications. When creating the MILAMIN algorithm in MATLAB the authors recognised that there were two significant bottlenecks with this method.

Firstly, two nested `for` loops are required to generate all the element stiffness matrices in a mesh. The outer loop, to loop through all the elements and the inner loop, to loop through all the Gauss points. As MATLAB loops are inherently slow and the iteration number of the element loop is large when calculating the stiffness matrix for a large mesh (excess of  $10^6$  DOF) this was recognised as the first bottleneck. Secondly, matrix calculations in MATLAB are performed by the Linear Algebra Package (LAPACK) which calls the Basic Linear Algebra Subprogramme (BLAS) package. Each time a BLAS function is called it has an overhead which is significantly larger than the computation time. In conventional FEA codes this overhead is large due to the many small matrix calculations occurring.

Dabrowski et al. [3] removed both bottlenecks by designing an algorithm where an entry in a local stiffness matrix could be computed for all elements simultaneously. This made the iteration size of the MATLAB `for` loops independent of the size of the problem, and also significantly reduced the number of BLAS calls and therefore their respective overhead. The MILAMIN routine was further improved by maximising cache reuse, a technique known as blocking. This work has since been extended by introducing parallel vectorised stiffness matrix calculations in [4].

More recently Rahman and Valdman [5] produced a fast MATLAB script for a volumetric integral of elements with linear nodal shape functions. The focus was to start with a non-vectorised code with a standard finite element structure and then improve its computational speed through vectorisation. One of the key characteristics was to preserve the code's original structure, this ensured that the readability was not lost which is often the case in code optimisation. Lack of readability in optimised codes was also highlighted by Dabrowski et al. [3]. Additionally, Anjam and Valdman [6] produced a vectorised MATLAB script for typical Raviart–Thomas elements used in discretisations of  $H(\text{div})$  spaces and Nédélec elements in discretisations of  $H(\text{curl})$  space. Andreassen et al. [7] provided a comparison and discussion of computational performance between different vector computational languages to assemble a FE global stiffness matrix. Cuvelier et al. [8,9] presented a more general approach to vectorise routines for multiple vector languages.

In a FEA code once all the local element stiffness matrices have been calculated they are assembled together, as a function of element topology, to form a sparse global stiffness matrix. In native MATLAB this is achieved using the command `sparse` which generates a sparse matrix from triplets of data: row position, column position and the associated value. The native performance is slow, Dabrowski et al. [3] use `sparse2` a non-native MATLAB command. Other sparse matrix commands for MATLAB have also been created, investigated, improved and discussed in [10], who also provide their own improvement and graphics processing unit (GPU) implementation.

In this paper the SIPG method for linear elasticity is implemented. Discontinuous Galerkin (DG) methods were first introduced by Reed et al. [11] for solving the neutron transport equation. Richter [12] prompted an extension of the original DG method to elliptical problems including linear convective-diffusion terms. However, the discontinuous approximation was only applied for the convective terms, with mixed methods for the second-order elliptic operators. Bassi and Rebay in [13] introduced the complete discontinuous approximations for both the convective and second-order elliptical operators.

One arising characteristic of DG methods is that the degrees of freedom are element specific, allowing simple communication at the element interfaces. Specifically,  $hp$ -refinement is simplified due to its capability to incorporate hanging nodes at the element interfaces. These qualities make the DG method very suitable for efficient adaptive refinement to achieve high fidelity simulations [14]. The penalty for allowing this flexibility is that the number of terms to be integrated in the weak form and degrees of freedom is higher for the same number and type of elements when compared to the CG method. The additional integrals are face connectivity stiffness terms which couple the unshared degrees of freedom between elements. This increases the number of calculations required to produce the global stiffness matrix,  $\mathbf{K}$  [15], the need for efficient production of the  $\mathbf{K}$  matrix is therefore necessary even for relatively small problems.

This paper extends the algorithm presented by Dabrowski et al. [3] to include optimised integration of the face terms for SIPG, [16], for linear elastic problems in a vectorised blocked form. Here all the algorithms are designed for native MATLAB functionality only, a clean departure from the majority of the optimised MATLAB algorithms available in the literature [3–5]. The only other known vectorised, non-blocked, MATLAB code on DG methods is by Frank et al. [17]. The authors in [17] consider the time dependent diffusion equation as their model problem, cast within a local DG formulation in 2D. Here we design a block vectorised code in native MATLAB, which exploits the symmetry in SIPG, to model the linearly elastic problem in 2D and 3D.

The paper begins with a brief overview of the bi-linear SIPG formulation for linear elasticity. This is directly followed by a reformulation of the bi-linear SIPG form into a matrix form that can be computed in a vectorised algorithm in Section 2. The vectorised algorithm for computing SIPG face stiffness terms is presented and discussed in Section 3, the volume integral is omitted as it was thoroughly covered in [3]. The vectorised algorithm for computing the SIPG face stiffness terms corresponds to the `Linear2D_DG.m` script, available from [18]. Timing results, validation, and discussions are presented in Section 4 followed by a conclusion in Section 5.

## 2. Optimising the DG method

### 2.1. SIPG weak form for linearly elastic problems

Here we consider the following model problem on a bounded Lipschitz polygonal/polyhedral domain  $\Omega$  in  $\mathbb{R}^d$ ,  $d \in \{2, 3\}$ , with the boundary  $\partial\Omega_N \cup \partial\Omega_D = \partial\Omega$ , where  $\partial\Omega_D$  and  $\partial\Omega_N$  are the portions of the boundary where homogeneous Dirichlet

and Neumann boundary conditions are respectively applied. The strong form of the problem, for small strain hyperelasticity, is defined as

$$\nabla \cdot \boldsymbol{\sigma}(\mathbf{u}) = \mathbf{0} \text{ in } \Omega, \quad \boldsymbol{\sigma}(\mathbf{u}) \cdot \mathbf{n} = \mathbf{g}_N \text{ on } \partial\Omega_N, \text{ and } \mathbf{u} = \mathbf{g}_D \text{ on } \partial\Omega_D. \tag{1}$$

$\mathbf{g}_D$  and  $\mathbf{g}_N$  are respectively the applied Dirichlet and Neumann boundary conditions. The Cauchy stress tensor is defined as  $\boldsymbol{\sigma} = \partial\psi(\boldsymbol{\varepsilon})/\partial\boldsymbol{\varepsilon}(\mathbf{u})$ , where  $\psi$  is the free energy function for hyperelasticity,  $\boldsymbol{\varepsilon}$  is small strain,  $\mathbf{u}$  is displacement and  $\mathbf{n}$  is the normal unit vector to the boundary. The Cauchy stress tensor can also be described  $\boldsymbol{\sigma} = \mathbf{D}\boldsymbol{\varepsilon}(\mathbf{u})$  where  $\mathbf{D}$  is a material stiffness tensor relating stress and strain.

This paper provides only a description of the 2D optimised code, therefore a description of the 3D element spaces and respective mesh is omitted. The polygonal finite element mesh  $\mathcal{T}$  is homogeneous in element type and is in general unstructured. Two element types are defined here, the triangle and quadrilateral, however since only one element type is present in a mesh both types are referred to as  $K$ . The polygonal mesh  $\mathcal{T}$  is comprised of elements  $K$  which are either the image of the reference triangle affine elemental mapping  $F_K : \widehat{K} \rightarrow K$ , or a quadrilateral element under a bi-linear elemental mapping  $B_K : \widehat{K} \rightarrow K$ . The homogeneous discontinuous Galerkin finite element space for triangular elements is defined as  $W_p(\mathcal{T}) = \{\mathbf{w} \in [L^2(\Omega)]^d : \forall K \in \mathcal{T}, \mathbf{w}|_K \in \mathcal{P}_p(K)\}$  and for quadrilateral elements as  $W_p(\mathcal{T}) = \{\mathbf{w} \in [L^2(\Omega)]^d : \forall K \in \mathcal{T}, \mathbf{w}|_K \in Q_p(K)\}$ , where  $P_p(K)$  is the space of polynomials on  $K$  of degree less than or equal to  $p$  and  $Q_p(K)$  is the space of polynomials on  $K$  less or equal to  $p$  in each dimension.

We denote by  $\mathcal{F}(K)$  the set of the three elemental faces for the triangle, or as the set of the four elemental faces for the quadrilateral, of an element  $K$ . If the intersection  $F = \partial K^+ \cap \partial K^-$  of two elements  $K^+, K^- \in \mathcal{T}$  is a segment, we call  $F$  an interior face of  $\mathcal{T}$ . The set of all interior faces is denoted by  $\mathcal{F}_I(\mathcal{T})$ . Analogously, if the intersection  $F = \partial K \cap \partial\Omega$  of an element  $K \in \mathcal{T}$  and  $\partial\Omega$  is a segment, we call  $F$  a boundary face of  $\mathcal{T}$ .

The SIPG method for the approximation of the model problem (1) is now introduced in the bi-linear form where the homogeneous Dirichlet boundary conditions on  $\partial\Omega_D$  are applied strongly and no body force exists. Find the displacement  $\mathbf{u}_h \in W_p(\mathcal{T})$  such that  $\mathbf{a}(\mathbf{u}_h, \mathbf{w}) = \mathbf{l}(\mathbf{w})$  for all  $\mathbf{w} \in W_p(\mathcal{T})$ , where

$$\begin{aligned} \mathbf{a}_K(\mathbf{u}_h, \mathbf{w}) &= \sum_{K \in \mathcal{T}} (\boldsymbol{\sigma}(\mathbf{u}_h), \boldsymbol{\varepsilon}(\mathbf{w}))_K - \sum_{F \in \mathcal{F}_I(\mathcal{T})} \{ \boldsymbol{\sigma}(\mathbf{u}_h) \}, \llbracket \mathbf{w} \rrbracket \}_F \\ &\quad - \sum_{F \in \cup \mathcal{F}_I(\mathcal{T})} \llbracket \mathbf{u}_h \rrbracket, \{ \boldsymbol{\sigma}(\mathbf{w}) \} \}_F + \sum_{F \in \mathcal{F}_I(\mathcal{T})} \beta \langle p_F^2 h_F^{-1} \llbracket \mathbf{u}_h \rrbracket, \llbracket \mathbf{w} \rrbracket \rangle_F, \end{aligned} \tag{2}$$

and

$$\mathbf{l}(\mathbf{w}) = \sum_{F \in \mathcal{F}_N(\mathcal{T})} \langle \mathbf{g}_N, \mathbf{w} \rangle_F. \tag{3}$$

$\beta$  is a penalty term for linear elastic SIPG defined in [19],  $h_F$  is this size of an element face, and

$$\llbracket \mathbf{v} \rrbracket = \mathbf{v} \Big|_{F^+} \cdot \mathbf{n}^+ - \mathbf{v} \Big|_{F^-} \cdot \mathbf{n}^+, \tag{4}$$

$$\{ \mathbf{v} \} = \frac{1}{2} \left( \mathbf{v} \Big|_{F^+} + \mathbf{v} \Big|_{F^-} \right), \tag{5}$$

where the element faces of  $K^+$  and  $K^-$  on an intersection  $F \in \mathcal{F}_I(\mathcal{T})$  are respectively referred to as  $F^+$  and  $F^-$ , with  $\mathbf{n}^+$  and  $\mathbf{n}^-$  as their respective outward normals. Additionally for convenience  $(\cdot, \cdot)$  and  $\langle \cdot, \cdot \rangle$  are used, where  $(a, b)_\Omega = \int_\Omega ab$  and  $\langle a, b \rangle_{\partial\Omega} = \int_{\partial\Omega} ab$ .

### 2.2. Matrix form of the SIPG method

Now that the weak form of the problem has been described it is possible to express the stress, strain and displacements in (2) as function of, nodal displacements, shape functions and their derivatives, and material stiffness. Once expressed, each term in the bi-linear form can be reformulated as a set of matrix multiplications which can be used to compute the stiffness matrix for SIPG. The first step to achieving the matrix formulation is decomposing the element displacements  $\mathbf{u}_n$  into a matrix of element shape functions  $\mathbf{N}_n$  and their corresponding coefficients  $\mathbf{u}_n$ , such that  $\mathbf{u}_n = \mathbf{N}_n \mathbf{u}_n$  where

$$\mathbf{N}_n = \begin{bmatrix} N_1 & 0 & N_2 & 0 & \dots & N_{\text{nen}} & 0 \\ 0 & N_1 & 0 & N_2 & \dots & 0 & N_{\text{nen}} \end{bmatrix}. \tag{6}$$

$\mathbf{u}_n = [u_1, v_1, \dots, u_{\text{nen}}, v_{\text{nen}}]^T$ , nen is the number of element nodes, and,  $u$  and  $v$  are respectively the displacements in the  $x$  and  $y$  directions of the Cartesian coordinate system. Similarly the test function can be represented as  $\mathbf{w} = \mathbf{N}_n \mathbf{w}_n$ .

As the small strain tensor is a function of  $\mathbf{u}_n$ , the strain can also be expressed as a set of matrix multiplications  $\boldsymbol{\varepsilon} = \mathbf{L}\mathbf{N}_n\mathbf{u}_n$  with the additional term

$$\mathbf{L} = \begin{bmatrix} \frac{\partial}{\partial x} & 0 \\ 0 & \frac{\partial}{\partial y} \\ \frac{\partial}{\partial y} & \frac{\partial}{\partial x} \end{bmatrix}, \tag{7}$$

as the small strain partial differential matrix operator [15]. From hyperelasticity the Cauchy stress is simply expressed as  $\boldsymbol{\sigma} = \mathbf{D}\boldsymbol{\varepsilon} = \mathbf{D}\mathbf{L}\mathbf{N}_n\mathbf{u}_n$ , where  $\mathbf{D}$  is the plane stress or strain stiffness matrix. Substituting the matrix forms of the stress, strain and displacement into (2), and setting

$$\mathbf{B}_n = \mathbf{L}\mathbf{N}_n = \begin{bmatrix} \frac{\partial N_1}{\partial x} & 0 & \dots & \frac{\partial N_{\text{nen}}}{\partial x} & 0 \\ 0 & \frac{\partial N_1}{\partial y} & \dots & 0 & \frac{\partial N_{\text{nen}}}{\partial y} \\ \frac{\partial N_1}{\partial y} & \frac{\partial N_1}{\partial x} & \dots & \frac{\partial N_{\text{nen}}}{\partial y} & \frac{\partial N_{\text{nen}}}{\partial x} \end{bmatrix}, \tag{8}$$

gives,

$$\begin{aligned} \mathbf{a}_K(\mathbf{u}_n, \mathbf{w}) &= \sum_{K \in \mathcal{T}} (\mathbf{D}\mathbf{B}_n\mathbf{u}_n, \mathbf{B}_n\mathbf{w}_n)_K - \sum_{F \in \mathcal{F}_I(\mathcal{T})} \{ \mathbf{D}\mathbf{B}_n\mathbf{u}_n \}, \llbracket \mathbf{N}_n\mathbf{w}_n \rrbracket \}_F \\ &- \sum_{F \in \cup \mathcal{F}_I(\mathcal{T})} \{ \llbracket \mathbf{N}_n\mathbf{u}_n \rrbracket, \mathbf{D}\mathbf{B}_n\mathbf{w}_n \}_F + \sum_{F \in \mathcal{F}_I(\mathcal{T})} \beta (p_F^2 h_F^{-1} \llbracket \mathbf{N}_n\mathbf{u}_n \rrbracket, \llbracket \mathbf{N}_n\mathbf{w}_n \rrbracket \}_F. \end{aligned} \tag{9}$$

The test function term in the Neumann boundary condition (3) is also expressed as a matrix multiplication

$$\mathbf{l}(\mathbf{w}) = \sum_{F \in \mathcal{F}_N(\mathcal{T})} \langle \mathbf{g}_N, \mathbf{N}_n\mathbf{w}_n \rangle_F. \tag{10}$$

Each term in the bi-linear form can now be reformulated into a set of matrix multiplications by setting (10) equal to (9), multiplying out the brackets and dividing by  $\mathbf{w}_n$  to give

$$\begin{aligned} \sum_{F \in \mathcal{F}_N(\mathcal{T})} \int_F \mathbf{N}^T \mathbf{g}_N &= \sum_{K \in \mathcal{T}} \int_K \mathbf{B}_n^T \mathbf{D}\mathbf{B}_n\mathbf{u}_n - \sum_{F \in \mathcal{F}_I(\mathcal{T})} \int_F (\mathbf{C}_1 + \mathbf{C}_2 + \mathbf{C}_3 + \mathbf{C}_4 \\ &+ \mathbf{D}_1 + \mathbf{D}_2 + \mathbf{D}_3 + \mathbf{D}_4 + \mathbf{E}_1 + \mathbf{E}_2 + \mathbf{E}_3 + \mathbf{E}_4) \\ &= \sum_{K \in \mathcal{T}} \mathbf{K}_{\text{CC}}\mathbf{u}_n + \sum_{F \in \mathcal{F}_I(\mathcal{T})} \mathbf{K}_{\text{LF}}\mathbf{u}_n \\ &= (\mathbf{K}_K + \mathbf{K}_F)\mathbf{U}_n = \mathbf{K}\mathbf{U}_n, \end{aligned} \tag{11}$$

where  $(\mathbf{K}_K + \mathbf{K}_F)$  is the global stiffness matrix,  $\mathbf{K}$ , comprised of the global element stiffness matrix and the face stiffness matrix respectively.  $\mathbf{U}_n$  is a vector containing all the nodal displacements for all  $K \in \mathcal{T}$  such that with respect to the mesh topology  $\mathbf{U}_n = \sum_{K \in \mathcal{T}} \mathbf{u}_n(K)$ . The terms in (11) in their full form are

$$\mathbf{C}_1 = \mathbf{B}_n^{+T} \mathbf{D}\mathbf{n}^{+T} \mathbf{N}_n^+ \mathbf{u}_n^+ / 2 = \mathbf{M}_{\mathbf{C}1} \mathbf{u}_n^+, \tag{12}$$

$$\mathbf{C}_2 = -\mathbf{B}_n^{+T} \mathbf{D}\mathbf{n}^{+T} \mathbf{N}_n^- \mathbf{u}_n^- / 2 = \mathbf{M}_{\mathbf{C}2} \mathbf{u}_n^-, \tag{13}$$

$$\mathbf{C}_3 = \mathbf{B}_n^{-T} \mathbf{D}\mathbf{n}^{+T} \mathbf{N}_n^+ \mathbf{u}_n^+ / 2 = \mathbf{M}_{\mathbf{C}3} \mathbf{u}_n^+, \tag{14}$$

$$\mathbf{C}_4 = -\mathbf{B}_n^{-T} \mathbf{D}\mathbf{n}^{+T} \mathbf{N}_n^- \mathbf{u}_n^- / 2 = \mathbf{M}_{\mathbf{C}4} \mathbf{u}_n^-, \tag{15}$$

$$\mathbf{D}_1 = \mathbf{N}_n^{+T} \mathbf{n}^+ \mathbf{D}\mathbf{B}_n^+ \mathbf{u}_n^+ / 2 = \mathbf{M}_{\mathbf{D}1} \mathbf{u}_n^+, \tag{16}$$

$$\mathbf{D}_2 = \mathbf{N}_n^{+T} \mathbf{n}^+ \mathbf{D}\mathbf{B}_n^- \mathbf{u}_n^- / 2 = \mathbf{M}_{\mathbf{D}2} \mathbf{u}_n^-, \tag{17}$$

$$\mathbf{D}_3 = -\mathbf{N}_n^{-T} \mathbf{n}^+ \mathbf{D}\mathbf{B}_n^+ \mathbf{u}_n^+ / 2 = \mathbf{M}_{\mathbf{D}3} \mathbf{u}_n^+, \tag{18}$$

$$\mathbf{D}_4 = -\mathbf{N}_n^{-T} \mathbf{n}^+ \mathbf{D}\mathbf{B}_n^- \mathbf{u}_n^- / 2 = \mathbf{M}_{\mathbf{D}4} \mathbf{u}_n^-, \tag{19}$$

$$\mathbf{E}_1 = \beta \frac{p^2}{h_F} \mathbf{N}_n^{+T} \mathbf{N}_n^+ \mathbf{u}_n^+ = \mathbf{M}_{\mathbf{E}1} \mathbf{u}_n^+, \tag{20}$$

$$\mathbf{E}_2 = -\beta \frac{p^2}{h_F} \mathbf{N}_n^{+T} \mathbf{N}_n^- \mathbf{u}_n^- = \mathbf{M}_{\mathbf{E}2} \mathbf{u}_n^-, \tag{21}$$

$$\mathbf{E}_3 = -\beta \frac{p^2}{h_F} \mathbf{N}_n^{-T} \mathbf{N}_n^+ \mathbf{u}_n^+ = \mathbf{M}_{E3} \mathbf{u}_n^+, \tag{22}$$

$$\mathbf{E}_4 = \beta \frac{p^2}{h_F} \mathbf{N}_n^{-T} \mathbf{N}_n^- \mathbf{u}_n^- = \mathbf{M}_{E4} \mathbf{u}_n^-. \tag{23}$$

The superscripts + and – in Eqs. (12) to (23) correspond to variables existing in  $K^+$  and  $K^-$  respectively. The variable  $\mathbf{n}^+$  is a matrix of normal components to  $F^+$ , its form for the SIPG linear elastic 2D problem is

$$\mathbf{n}^{+T} = \begin{bmatrix} n_x & 0 & n_y \\ 0 & n_y & n_x \end{bmatrix}. \tag{24}$$

Last the set  $M$  is defined as  $M = \{\mathbf{M}_{C1}, \mathbf{M}_{C2}, \mathbf{M}_{C3}, \mathbf{M}_{C4}, \mathbf{M}_{D1}, \mathbf{M}_{D2}, \mathbf{M}_{D3}, \mathbf{M}_{D4}, \mathbf{M}_{E1}, \mathbf{M}_{E2}, \mathbf{M}_{E3}, \mathbf{M}_{E4}\}$ .  $M$  is the set of partial stiffness matrices which when integrated over a single face  $F$ , and assembled together with respect to the element topology of  $K^+$  and  $K^-$ , produce the local SIPG face stiffness matrix  $\mathbf{K}_{LF}$  for the face  $F$ .

### 2.3. Vectorising the SIPG face integral

Traditionally in matrix formulated FE code two loops exist, an outer loop over the elements, and faces for SIPG, and an inner loop over the integration points. Within the inner loop small matrix operations occur to produce each partial stiffness matrix in  $M$ . The outer loop over the elements, and faces, is proportional to the size of the problem; loops in MATLAB are significantly slower than compiled code therefore an algorithm with this structure is unacceptable to use for large problems. Further, the overhead associated with each BLAS call is large due to the many small matrix multiplications [3]. In order to avoid loops which are proportional to the size of the problem, each partial stiffness matrix in  $M$  is reformulated into a corresponding matrix of scalar equations. Each scalar equation can be computed for all elements simultaneously in an entry-wise vector calculation, meaning the outer element loop can be removed. The improved algorithm now has the form: two inner loops over the entries in the reformulated matrices of  $M$  and an outer loop over the integration points. This speeds up computation for large problems since the size of loops are no longer dependent on the size of the problem being modelled. The overhead associated with the BLAS calls is also significantly reduced since the number of BLAS calls is minimised.

The method for reformulating each partial stiffness matrix in Eqs. (12) to (23) is the same, here the integral matrix term  $\mathbf{M}_{C2}$  from (13) is used as an example,

$$\sum_{F \in \mathcal{F}_1} \left( \int_F \mathbf{M}_{C2} \right) \mathbf{u}^- = -\frac{1}{2} \sum_{F \in \mathcal{F}_1} \left( \int_F \mathbf{B}^{+T} \mathbf{D} \mathbf{n}^{+T} \mathbf{N}^- \right) \mathbf{u}^-. \tag{25}$$

The vector  $\mathbf{u}^-$  is omitted in the integral as it is the solution to the linear elastic problem and so is unknown.

To reformulate  $\mathbf{M}_{C2}$  the shape functions and the derivatives,  $\mathbf{N}^-$  and  $\mathbf{B}^+$ , are expanded into their full form so  $\mathbf{M}_{C2}$  becomes,

$$\mathbf{M}_{C2} = \begin{bmatrix} \frac{\partial N_1^+}{\partial x} & 0 & \frac{\partial N_1^+}{\partial y} \\ 0 & \frac{\partial N_1^+}{\partial y} & \frac{\partial N_1^+}{\partial x} \\ \vdots & \vdots & \vdots \\ \frac{\partial N_{nen}^+}{\partial x} & 0 & \frac{\partial N_{nen}^+}{\partial y} \\ 0 & \frac{\partial N_{nen}^+}{\partial y} & \frac{\partial N_{nen}^+}{\partial x} \end{bmatrix} \mathbf{D} \mathbf{n}^+ \begin{bmatrix} N_1^- & 0 & \dots & N_{nen}^- & 0 \\ 0 & N_1^- & \dots & 0 & N_{nen}^- \end{bmatrix}. \tag{26}$$

The forms of  $\mathbf{B}^+$  and  $\mathbf{N}^-$  are repeated down the rows and along the columns respectively so  $\mathbf{M}_{C2}$  can therefore be rewritten in the condensed form

$$\mathbf{M}_{C2} = \sum_{i=1}^{nen} \sum_{j=1}^{nen} \begin{bmatrix} \frac{\partial N_i^+}{\partial x} & 0 & \frac{\partial N_i^+}{\partial y} \\ 0 & \frac{\partial N_i^+}{\partial y} & \frac{\partial N_i^+}{\partial x} \end{bmatrix} \mathbf{D} \mathbf{n}^+ \begin{bmatrix} N_j^- & 0 \\ 0 & N_j^- \end{bmatrix}, \tag{27}$$

where  $i$  and  $j$  are respectively the local finite element node numbers for elements  $K^+$  and  $K^-$  with corresponding shape functions that pre-and-post multiply  $\mathbf{M}_{C2}$ . The material stiffness matrix  $\mathbf{D}$  is either acting in plane strain or stress and so is represented as

$$\mathbf{D} = \begin{bmatrix} A & B & 0 \\ B & A & 0 \\ 0 & 0 & C \end{bmatrix}. \tag{28}$$

When multiplied out, (27) becomes

$$\mathbf{M}_{C_2}^r = \sum_{i=1}^{\text{nen}} \sum_{j=1}^{\text{nen}} \begin{bmatrix} N_j^- \left( A \frac{\partial N_i^+}{\partial x} n_x^+ + C \frac{\partial N_i^+}{\partial y} n_y^+ \right) & N_j^- \left( B \frac{\partial N_i^+}{\partial x} n_y^+ + C \frac{\partial N_i^+}{\partial y} n_x^+ \right) \\ N_j^- \left( B \frac{\partial N_i^+}{\partial y} n_x^+ + C \frac{\partial N_i^+}{\partial x} n_y^+ \right) & N_j^- \left( A \frac{\partial N_i^+}{\partial y} n_y^+ + C \frac{\partial N_i^+}{\partial x} n_x^+ \right) \end{bmatrix}. \quad (29)$$

$\mathbf{M}_{C_2}^r \equiv \mathbf{M}_{C_2}$ , however for the sake of clarity the reduced 2-by-2 matrix form of  $\mathbf{M}_{C_2}$  is redefined. An equivalent matrix exists for all the partial stiffness matrices in  $M$ . The new set  $M^r$  is now defined and contains the equivalent 2-by-2 matrix forms of matrices in the set  $M$ , denoted with the superscript  $r$ , such that  $M^r = \{\mathbf{M}_{C_1}^r, \mathbf{M}_{C_2}^r, \mathbf{M}_{C_3}^r, \mathbf{M}_{C_4}^r, \mathbf{M}_{D_1}^r, \mathbf{M}_{D_2}^r, \mathbf{M}_{D_3}^r, \mathbf{M}_{D_4}^r, \mathbf{M}_{E_1}^r, \mathbf{M}_{E_2}^r, \mathbf{M}_{E_3}^r, \mathbf{M}_{E_4}^r\}$ . All the entries in  $\mathbf{M}_{C_2}$  are now represented by 4 scalar equations which are looped over the indices  $(i, j)$ .

### 3. Code assembly

The complete code layout is summarised by Algorithm 1, and correlates to lines of Linear2D\_DG.m, the optimised SIPG.m script provided by [18]. Algorithm 1 contains three stages:

1. Area integral: lines 1–5. MATLAB code: lines 22–64.
2. SIPG face integral: lines 6–11. MATLAB code: lines 93–310.
3. Sparse storage: lines 12. MATLAB code: lines 311–350.

In stage 2 the SIPG face integral computes the local face stiffness matrix  $\mathbf{K}_{LF}$  for all faces in the mesh. Stage 2 dominates the number of lines in the code due to having 12 terms to evaluate rather than just one like in stage 1 (11). In stage 3 the local face stiffness matrices are assembled into a sparse global stiffness matrix completing the algorithm.

---

#### Algorithm 1 Complete Code layout.

---

```

1: for Area block do
2:   for Area Gauss blocks do
3:     Area integral
4:   end for
5: end for
6: for DG face block do
7:   for Gauss point loop do
8:     DG face integral set-up
9:     DG face integral
10:  end for
11: end for
12: Sparse storage

```

---

The optimised SIPG area integral is not discussed as it is identical to the optimised CG area integral in [3]. This paper focuses on the novel implementation of the blocked vectorised integration of the partial stiffness matrices in  $M$  to produce the global face stiffness matrix  $\mathbf{K}_F$ . When considering the fast vectorised computation of the SIPG face terms in (11) there are six main aspects to address: optimising the CPU cache reuse (blocking) to further increase the speed obtained by vectorisation, the structure of the vectorised algorithm, memory allocation, reducing the number of BLAS operations, generating variables for the blocked algorithm, and the reference Gauss point locations. In the following sections each of these points will be addressed in turn.

Section 2.3 demonstrated that the matrices in  $M$  could be represented by a repeated 2-by-2 matrix of scalar equations looped over the nodal indices  $(i, j)$ . Arranging the partial stiffness matrices in  $M$  into their equivalent form in  $M^r$  means that each partial stiffness matrix is constructed from four entries, each of which are scalar equations that are applicable to all finite elements in the mesh. Reformulating the matrices in  $M$  into the form in  $M^r$  allows the integral of each entry in the partial stiffness entry to be computed for all faces simultaneously in a vectorised algorithm; the schematic for such an algorithm in MATLAB is represented in Fig. 1. The following subsections use the matrix  $\mathbf{M}_{C_2}^r$  as an example to describe stage 2 of Algorithm 1.

#### 3.1. Blocking

When the CPU performs a BLAS operation the best performance is achieved when all the data required for the operation resides in the cache. However when the data size is too large to reside entirely in the cache, sections are stored in the RAM, which is slower to access. The technique for maximising the vector size, with the condition that the data for a BLAS operation resides in the cache memory, is called blocking.

If the number of faces included in a matrix entry calculation is too large such that data for the calculation resides outside the CPU cache, the set of faces  $\mathcal{F}(\mathcal{T})$  is split into blocks of faces, defined as SIPG face blocks. The vector calculation for a matrix entry is now performed for each block in turn so the cache reuse is maximised, reducing the overall run time. This process is dictated by the for loop on line 2, Fig. 1, which runs through all the SIPG face blocks in the mesh.

```

1  glob_2=zeros(tot_num_faces,ndof,ndof);
2  for Block_n = 1:num_blocks % MATLAB code: lines 93-310.
3      for gp = 1:ngp % MATLAB code: lines 125-299.
4          glob_pn=zeros(nel_block,ndof,ndof);
5          % Variable generation for current block Figure 3.
6          for i = 1:nen % MATLAB code: lines 175-297.
7              j2=(nen+1)-i;
8              Ai=(i-1)*2; Bj=(j2-1)*2;

9              C2t_11=((A.*dNx_p(:,i).*nx)+(C.*dNy_p(:,i).*ny)).*int_W;
              % Components of the remaining entries which vary in 'i' of:
               $M_{C1}$ ,  $M_{C2}$ ,  $M_{C3}$ , and  $M_{C4}$ 
              % are also computed here.

10             D2t_12=((B.*dNy_p(:,j2).*nx)+(C.*dNx_p(:,j2).*ny)).*int_w;
              % Components of the remaining entries which vary in 'j2' of:
               $M_{D1}$ ,  $M_{D2}$ ,  $M_{D3}$ , and  $M_{D4}$ 
              % are also computed here.

11             for j = i:nen
12                 i2=(nen+1)-j;
13                 Aj=(j-1)*2; Bi=(i2-1)*2;

14                 glob_pn(:,Ai,Aj)=Nn(:,j).*C2t_11.*half(i,j)+glob_pn(:,Ai,Aj);
              % Computations of all components of:
              % glob_pp, glob_pn, glob_np and glob_nn occur here.

15                 if i<j
16                     glob_pn(:,Bi,Bj+1)=Np(i2).*D2t_12+glob_pn(:,Bi,Bj+1);
17                 end
18             end
19         end
20     end
21     glob_2(bl_index, :, :) = glob_pn;
22 end

```

Fig. 1. Vectorised calculation schematic of entry (1,1) in  $M_{C2}^T$  and (1,2) in  $M_{D2}^T$ .

### 3.2. Structure

The structure of the algorithm which generates all the SIPG local face stiffness matrices, for all faces in the mesh, is described in Fig. 1. It is characterised by four `for` loops appearing on lines 2, 3, 6, and 11. The first loop, loops through all the SIPG face blocks. The second loop is the Gauss point loop which numerically integrates the partial stiffness matrices in  $M^T$  to

generate the local face stiffness matrix (11) for all faces in SIPG face block. The final two loops are over nodal combinations (i, j) of the matrices in the set  $M^r$ , which when integrated form the local face stiffness matrices for (11).

### 3.3. Reducing the number of BLAS operations

It is possible to take advantage of the structure of matrices in  $M^r$  to reduce the number of BLAS operations. As an example, the entry (1,1) of  $M_{C2}^r$  can be split into two components; a component which varies with row number i, represented in Fig. 1 as C2t\_11 on line 9, and one with column number j, line 14. All entries of  $M_{C1}^r$ ,  $M_{C2}^r$ ,  $M_{C3}^r$ , and  $M_{C4}^r$  can be split into two components in the same way. The component which is a function of i requires more BLAS operations and is calculated outside the inner node loop, line 9 of Fig. 1. The i component is then multiplied with the j component and added to glob\_pn within the inner loop on line 14. This reduces the number of BLAS calls, the computation time, and the time associated with calling the routine.

An equivalent  $M_{D2}^r$  matrix can be constructed for  $M_{D2}$ , (11). Unlike  $M_{C2}^r$ , the components of entries in  $M_{D2}^r$  which are a function of column number require more BLAS operation than those which are a function of row number. As an example, the entry (1,2) of  $M_{D2}^r$  can be split into two components; a component which varies with column number j2 represented in Fig. 1 as D2t\_12 on line 10, and one with row number i2, line 16. The column index j2 is defined on line 7, and the row index i2 is defined on line 12 of Fig. 1. The multiplication of the column and row dependent variables, as with  $M_{C2}^r$ , still occurs in the inner loop, line 16 of Fig. 1. Equivalently all entries of  $M_{D1}^r$ ,  $M_{D2}^r$ ,  $M_{D3}^r$ , and  $M_{D4}^r$  can be split into two components.

The number of BLAS calls can be reduced further by utilising the symmetry of the global stiffness matrix so that only the upper triangular components of the local stiffness matrices need to be calculated. The for loop matrix indices, lines 6 and 11 of Fig. 1, are therefore restricted to this part of the  $M_{C2}$  matrix. However, in order to keep the size of the node index loops the same between  $M_{C2}^r$  and  $M_{D2}^r$ , j2 and i2 of  $M_{D2}^r$  are looped through in reverse order, lines 7 and 12 of Fig. 1. Lastly, the MATLAB indices j2 and i2 refer to variables in  $M_{D1}^r$ ,  $M_{D2}^r$ ,  $M_{D3}^r$  and  $M_{D4}^r$ , and the, i and j, indices refer to variables in  $M_{C1}^r$ ,  $M_{C2}^r$ ,  $M_{C3}^r$ ,  $M_{C4}^r$ ,  $M_{E1}^r$ ,  $M_{E2}^r$ ,  $M_{E3}^r$  and  $M_{E4}^r$ .

The loop indices, i and j, correspond to the node number of elements in the local matrix. When considering nodes on the leading diagonal, i.e. when i==j, only the upper triangle components of the four entries in the matrices in  $M^r$  are required. An if statement is present, line 15 of Fig. 1, so all the entries in a matrix of  $M^r$  are computed only if i < j, and the lower triangular components are omitted if i == j.

To complete the global stiffness matrix formulation, the transpose of the global matrix is added to itself. To avoid doubling values on the leading diagonals of the local matrices, diagonal entries of the matrices in  $M^r$  are divided by 2 when i==j by half(i, j). half(i, j) is a simple script which returns a value 0.5 if i==j and 1 otherwise.

### 3.4. Memory allocation

Memory for variables which increase in size during the nodal loops is preallocated, this prevents reallocation of the variables on the RAM which reduces the run time. The partial stiffness matrices in  $M$  can be split into four sets. Each set can be summed together to form

$$G_s = \sum_{F \in \mathcal{F}_1} \int_F M_{C_s} + M_{D_s} + M_{E_s} \quad \text{where } s = 1 \dots 4. \tag{30}$$

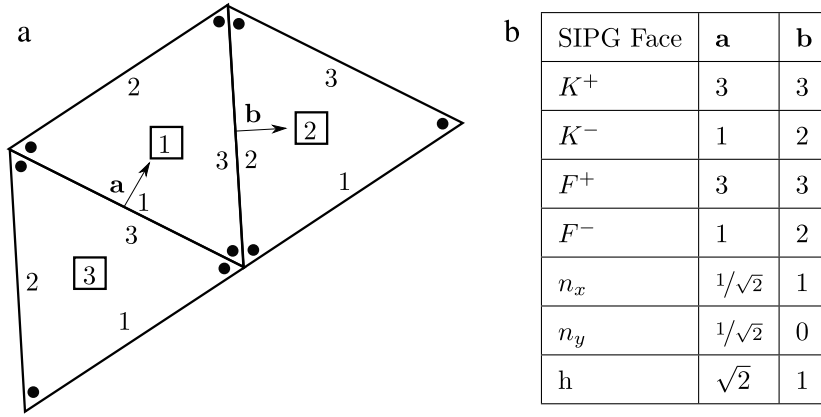
For a face  $F$ , when  $s = 1$ ,  $M_{C1}$ ,  $M_{D1}$  and  $M_{E1}$  reside in the same location in the global stiffness matrix,

$$G_1 = \sum_{F \in \mathcal{F}_1} \int_F M_{C1} + M_{D1} + M_{E1}. \tag{31}$$

Therefore only one storage variable needs to be preallocated for  $M_{C1}$ ,  $M_{D1}$  and  $M_{E1}$ . In Fig. 1 the storage variable glob\_2 is defined for  $G_2$ , line 1, for all  $F \in \mathcal{F}(\mathcal{T})$ . Performance improvements were found when a second temporary storage variable was used during the local matrix calculation, glob\_pn defined on line 4, which corresponded to all faces in the current SIPG face block for the set  $G_2$ . Once integration is completed for the current SIPG face block, glob\_pn is stored into glob\_2 on line 21 of Fig. 1.

The variable glob\_pn is a three dimensional array; the first dimension corresponds to the face numbers in the SIPG block, the second and third dimensions respectively correspond to the degrees of freedom of the finite element that pre-and-post-multiplied the partial stiffness matrix. As an example the local degrees of freedom of the entry (1,1) of  $M_{C2}^r$  are a function of node numbers i and j. The degrees of freedom are provided by the variables Ai and Aj, they are used to steer the entry (1,1) into the appropriate second and third dimensions of glob\_pn, line 14 of Fig. 1. Equivalently entry (1,2) of  $M_{D2}^r$  is a function of the node numbers i2 and j2. Here the degree of freedom numbers Bi and Bj+1 stores entry (1,2) into the appropriate position in glob\_pn, line 16 of Fig. 1. The +1 term of Bj+1 derives from the entry being in the second column of  $M_{D2}^r$ .





**Fig. 2.** (a) Example of 3 elements in a 2D DG mesh. Arrows indicate the outward normal direction, values in a box are the element number and values on the element edge are the local face number. (b) The transpose of the matrix `etp1_face` for DG faces **a** and **b** in Fig. 2(a).  $n_x$  and  $n_y$  are the outward normal components, and  $h$  is the length of the face,  $K^+$  is the positive element number with face  $F^+$ , and  $K^-$  is the negative element number with face  $F^-$ .

### 3.5. Generating variables for blocked algorithm

When integrating an entry of a partial stiffness matrix simultaneously for multiple SIPG faces, the shape function and their derivatives for the scalar equation for that entry must be in vector form. To compute a local SIPG face stiffness matrix, information is required from both the  $K^+$  and  $K^-$  elements sharing a face. During mesh generation the face connectivity for all faces in the mesh  $\mathcal{F}(\mathcal{T})$  is stored in the face connectivity matrix `etp1_face`, where a column correlates to face number in  $\mathcal{F}(\mathcal{T})$ , Fig. 2(b).

The script to generate the variables for face integration is represented in Fig. 3 and runs on line 5 of Fig. 1. The outer two loops, the same block and Gauss point loop as of Fig. 1, are represented on lines 1 and 4 of Fig. 3. The third loop is over the local element face number `fn`. This is required as local shape function values, `Nr`, and their local derivatives, `dNr`, are unique to a local element face. To compute global shape function derivative terms, `dNx_p` and `dNy_p`, for multiple elements simultaneously, only one local face can be considered at time. Therefore manipulation of `etp1_face` is required to only consider SIPG faces in the current block and current local face number. The information for the faces in the current SIPG face block is selected from `etp1_face` and is stored in `etp1_face_block` by the index `block_index`, line 3 Fig. 3. Further, as the shape functions and their derivatives can only be computed for one local face number, `fn`, at a time, the face information for elements  $K^+$  with the current face number `fn` is stored in `etp1_face_block_fn`, line 7 of Fig. 3.

The rows of `dNx_p`, and `dNy_p`, correspond to the same ordering of elements as in `etp1_face_block`. The columns of `dNx_p`, and `dNy_p`, correspond to a shape function number which is selected by `i` or `j` in Fig. 1. Their computation occurs in several large matrix operations. First the Jacobian components `Jxp` and `Jyp` are computed on lines 10–11,

$$\underbrace{\begin{bmatrix} \frac{\partial x_1}{\partial \xi} & \frac{\partial x_2}{\partial \xi} & \cdots & \frac{\partial x_{nfb}}{\partial \xi} \\ \frac{\partial x_1}{\partial \eta} & \frac{\partial x_2}{\partial \eta} & \cdots & \frac{\partial x_{nfb}}{\partial \eta} \end{bmatrix}}_{Jxp} = \underbrace{\begin{bmatrix} \frac{\partial N_1}{\partial \xi} & \frac{\partial N_2}{\partial \xi} & \cdots & \frac{\partial N_{nen}}{\partial \xi} \\ \frac{\partial N_1}{\partial \eta} & \frac{\partial N_2}{\partial \eta} & \cdots & \frac{\partial N_{nen}}{\partial \eta} \end{bmatrix}}_{dNr(indx_dNr(fn, gp), :)} \underbrace{\begin{bmatrix} x_1^1 & x_1^2 & \cdots & x_1^{nfb} \\ x_2^1 & x_2^2 & \cdots & x_2^{nfb} \\ \vdots & \vdots & \cdots & \vdots \\ x_{nen}^1 & x_{nen}^2 & \cdots & x_{nen}^{nfb} \end{bmatrix}}_{coord\_xp}, \tag{32}$$

where the subscript `nfb` corresponds to the number of DG faces in the block. The Jacobian determinant and its inverse are computed in an explicit manner on lines 12–14. The global shape function derivatives for the current face, `fn`, are calculated on lines 15–16, using

$$\underbrace{\begin{bmatrix} \frac{\partial N_1}{\partial x_1} & \frac{\partial N_2}{\partial x_1} & \cdots & \frac{\partial N_{nen}}{\partial x_1} \\ \frac{\partial N_1}{\partial x_2} & \frac{\partial N_2}{\partial x_2} & \cdots & \frac{\partial N_{nen}}{\partial x_2} \\ \vdots & \vdots & \cdots & \vdots \\ \frac{\partial N_1}{\partial x_{nfb}} & \frac{\partial N_2}{\partial x_{nfb}} & \cdots & \frac{\partial N_{nen}}{\partial x_{nfb}} \end{bmatrix}}_{dNx\_p(index\_p,:)} = \underbrace{\begin{bmatrix} \frac{\partial \xi}{\partial x_1} & \frac{\partial \eta}{\partial x_1} \\ \frac{\partial \xi}{\partial x_2} & \frac{\partial \eta}{\partial x_2} \\ \vdots & \vdots \\ \frac{\partial \xi}{\partial x_{nfb}} & \frac{\partial \eta}{\partial x_{nfb}} \end{bmatrix}}_{invJxp} \underbrace{\begin{bmatrix} \frac{\partial N_1}{\partial \xi} & \frac{\partial N_2}{\partial \xi} & \cdots & \frac{\partial N_{nen}}{\partial \xi} \\ \frac{\partial N_1}{\partial \eta} & \frac{\partial N_2}{\partial \eta} & \cdots & \frac{\partial N_{nen}}{\partial \eta} \end{bmatrix}}_{dNr(indx_dNr(fn, gp), :)}. \tag{33}$$

```

1  for Block_n = 1:num_blocks % MATLAB code: lines 93–310.
2      etpl_face %defined in Figure 2b
3      etpl_face_block = etpl_face(:,bl_index)
4      for gp = 1:ngp % MATLAB code: lines 125–299.
5          for fn = 1:num_faces % MATLAB code: lines 127–170.
6              index_p = etpl_face_block(:,3)==fn;
7              etpl_face_block_fn = etpl_face_block(:,index_p);

              K+ % Elements for current SIPG face block and local face number
8              pos_el = etpl_face_block_fn(:,1);

              % Vector of x (coord_xp) and y (coord_yp)
              % coordinates for pos_el
9              [coord_xp,coord_yp] = coord(etpl(pos_el,:),:);

              % Jacobian calculation for pos_el (32)
10             Jxp = dNr(indx_dNr(fn,gp),:)*coord_xp(index_p,:);
11             Jyp = dNr(indx_dNr(fn,gp),:)*coord_yp(index_p,:);

              % Vectorised determinant and calculation [3]
12             det = (Jxp(1,:).*Jyp(2,:))-(Jxp(2,:).*Jyp(1,:));

              % Vectorised inverse Jacobian calculation
13             invJxp = [ det.*(Jyp(2,:))',-det.*(Jyp(1,:))'];
14             invJyp = [-det.*(Jxp(2,:))', det.*(Jxp(1,:))'];

              % Global shape function derivative calculation (33)
15             dNx_p(index_p,:) = invJx_p*dNr(indx_dNr(fn,gp),:);
16             dNy_p(index_p,:) = invJy_p*dNr(indx_dNr(fn,gp),:);

              % Shape functions storage
17             Np(index_p,:) = repmat(Nr(indx_Nr(fn,gp),:))...
                                ,sum(index_p,1);
18             % Calc. for dNx_n, dNy_n and Nn. MATLAB code: lines 152–168
19         end
20     end

              % Vectorised integral weight calculation
21     int_W=2.*pen./etpl_face_block(:,end);
22 end

```

**Fig. 3.** An algorithm for generating variables for multiple DG faces simultaneously.

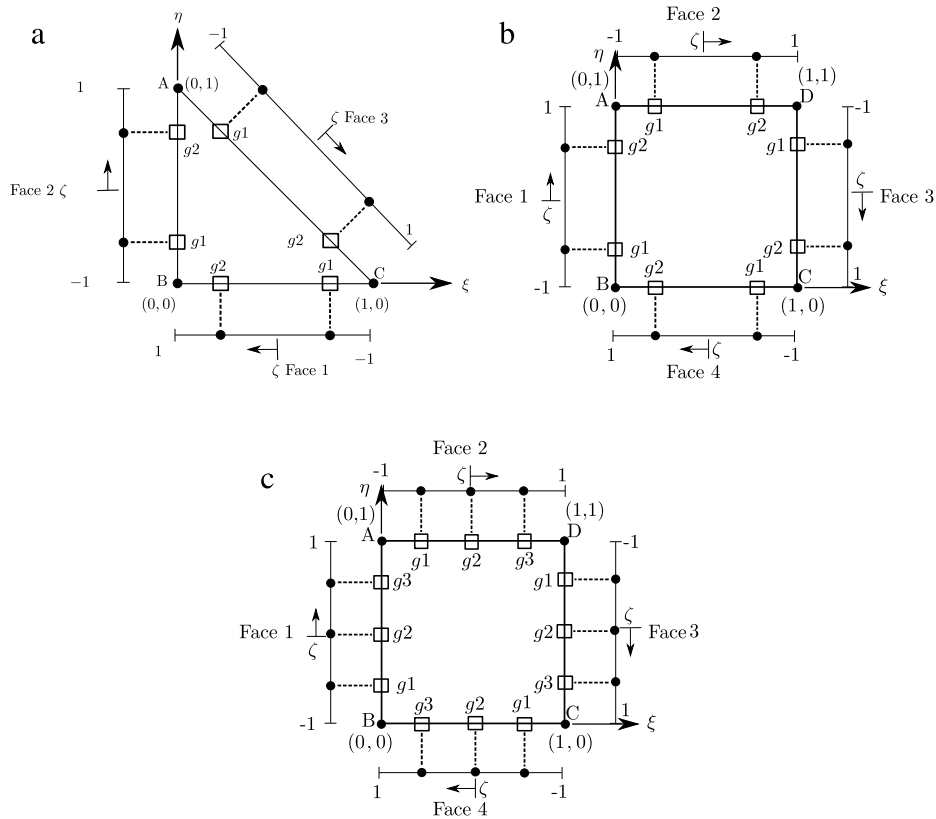
The result is stored into  $dNx\_p$  and  $dNy\_p$  with  $index\_p$ , lines 15–16, this ensures the element ordering remains consistent with  $etpl\_face\_block$ .

The shape functions are only dependent on their local position and therefore only the local value of  $Nr$ . The values of  $Nr$  for each face are stored into the matrix  $Np$  with  $index\_p$ , line 17, to ensure consistent element ordering with  $etpl\_face\_block$ .

The algorithm in Fig. 3 is only applicable to  $K^+$  elements but with a few simple changes can be for  $K^-$  elements: line 6 change  $etpl\_face\_block(:,3)$  to  $etpl\_face\_block(:,4)$ , line 8 change  $etpl\_face\_block\_fn(:,1)$  to  $etpl\_face\_block\_fn(:,2)$ , lines 15 and 16 change  $dNx\_p$  and  $dNy\_p$  to  $dNx\_n$  and  $dNy\_n$  respectively, and line 17 changes  $Np$  to  $Nn$ . Lastly change  $index\_p$  to  $index\_n$  everywhere. Gauss points along a face for  $K^-$  elements are considered in reverse order so that they align with  $K^+$  Gauss points in the global domain; MATLAB code: lines 153–154.

### 3.6. Reference Gauss point locations

For each local face,  $fn$ , the Gauss point locations in the reference frame are hard coded into the algorithm. Their positions on a face are used to generate local shape functions and their derivatives. Each face in the reference frame is numbered as



**Fig. 4.**  $K^+$  Gauss point ordering and face ordering for the constant strain triangular element (a), bi-linear quadrilateral element (b), and bi-quadratic quadrilateral element (c).  $\xi$  and  $\eta$  are the coordinates in the reference element domain,  $\zeta$  is the coordinate in the reference line domain and  $g\#$  is a Gauss point number specific to a face number. On a negative element the Gauss point ordering is reversed.

shown for a triangle element, Fig. 4a, bi-linear quadrilateral element, Fig. 4b, and bi-quadratic quadrilateral element, Fig. 4c. The Gauss points on  $F^+$  are numbered clockwise whilst on  $F^-$  they are anticlockwise. This ensures that the Gauss points for two connected elements align in the global domain.

The face integrals are performed with respect to the reference local face coordinate  $\zeta \in [-1, 1]$ . To determine the shape functions and shape function derivative values from  $\zeta$ , it is necessary to convert from the reference line domain to reference element domain, coordinates  $\xi \in [0, 1]$  and  $\eta \in [0, 1]$ , with

$$\xi = \frac{(\zeta + 1)(\xi_a - \xi_b)}{2} + \xi_b, \tag{34}$$

and

$$\eta = \frac{(\zeta + 1)(\eta_a - \eta_b)}{2} + \eta_b. \tag{35}$$

Here,  $a$  refers to the most clockwise vertex existing at the end of the face, and  $b$  the previous vertex. As an example, for face 2 of the triangular element  $a = A$ , and  $b = B$ , Fig. 4a. However, for face 1,  $a = B$  and  $b = C$ . Using the values of  $\xi$  and  $\eta$ , mapped from  $\zeta$ , the shape functions and the reference shape functions derivatives can be determined for each Gauss point location on each face. The face calculations use standard Gauss quadrature weights and locations.

### 3.7. Sparse storage

The assembly of all local face stiffness matrices forms a global stiffness matrix,  $K_F$  in (11). The algorithm which performs the summation operation is shown in Fig. 5.

On line 21 in Fig. 1, `glob_2` stores all components of  $G_2$  from (31). Equivalent storage variables exist for the remaining subscripts, see Table 1.

To store `glob_2` into the global stiffness matrix it is first rearranged into a vector form with the MATLAB function `reshape`, line 4 of Fig. 5. The new data structure of `glob_2` is described in Table 2. When steering `glob_1`, `glob_2`, `glob_3`

```

1  tot_f=size(etpl_face,2); %total number of DG faces
2  ndof=(nen*2);          % number of degrees of freedom
3  nndof=(nen*2)^2;      % number of degrees of freedom (NDOF) squared

% reshaping the global storage matrix into a vector, Table 4.
4  glob_2_rs = reshape(reshape(glob_2,tot_f,nndof)',tot_f*(nndof),1);

5  ed_p=ed(etpl_face(1,:),:); %steering matrix for + element dof
6  ed_n=ed(etpl_face(2,:),:); %steering matrix for - element dof

% steering vectors pos_i, pos_j, neg_i and neg_j, Table 4
7  pos_i = reshape(repmat(ed_p,1,ndof)',1,tot_f*nndof);
8  ed_pve = reshape(ed_p',1,ndof*tot_f);
9  pos_j = reshape(repmat(ed_pve,ndof,1),1,tot_f*nndof);
10 neg_i = reshape(repmat(ed_n,1,ndof)',1,tot_f*nndof);
11 ed_nve = reshape(ed_n',1,ndof*tot_f);
12 neg_j = reshape(repmat(ed_nve,ndof,1),1,tot_f*nndof);

% Global stiffness matrix sparse storage
13 k = k - sparse(pos_i,neg_j,glob_2_rs);
14 k=k+k'; % Completing the global stiffness matrix formulation

```

Fig. 5. Segment of MATLAB script for storing `glob_2` into a sparse matrix. MATLAB code: lines 311–350.

Table 1

Storage variables and their associated row and column degree of freedom numbers. The row and column degrees of freedom,  $i$  and  $j$ , have a prefix `pos` and `neg`. `pos` and `neg` correspond to pre-or-post multiplication of,  $\mathbf{N}^+$  or  $\mathbf{B}^+$ , and,  $\mathbf{N}^-$  or  $\mathbf{B}^-$ .

Partial stiffness matrices	Face term	Row	Column
$\sum_{F \in \mathcal{F}_1} \int_F \mathbf{M}_{C1} + \mathbf{M}_{D1} + \mathbf{M}_{E1} \rightarrow$	<code>glob_1</code>	<code>pos_i</code>	<code>pos_j</code>
$\sum_{F \in \mathcal{F}_1} \int_F \mathbf{M}_{C2} + \mathbf{M}_{D2} + \mathbf{M}_{E2} \rightarrow$	<code>glob_2</code>	<code>pos_i</code>	<code>neg_j</code>
$\sum_{F \in \mathcal{F}_1} \int_F \mathbf{M}_{C3} + \mathbf{M}_{D3} + \mathbf{M}_{E3} \rightarrow$	<code>glob_3</code>	<code>neg_i</code>	<code>pos_j</code>
$\sum_{F \in \mathcal{F}_1} \int_F \mathbf{M}_{C4} + \mathbf{M}_{D4} + \mathbf{M}_{E4} \rightarrow$	<code>glob_4</code>	<code>neg_i</code>	<code>neg_j</code>

or `glob_4` into the global face matrix  $\mathbf{K}_F$ , line 13 of Fig. 5, the row numbers correspond to the finite elements' degrees of freedom that pre-multiplied the partial stiffness matrices, of `glob_1`, `glob_2`, `glob_3` or `glob_4`. The column numbers represent the degrees of freedom of finite elements that post-multiplied.

After all the stiffness matrices are stored into the global sparse matrix, the sparse matrix is transposed and summed, line 14 of Fig. 5, completing the global stiffness matrix.

#### 4. Blocking and numerical analysis

This section demonstrates the performance gain obtained when using vectorised blocked scripts to generate all the SIPG local face stiffness matrices (11). All computations were performed in a native MATLAB environment using double precision float accuracy, the backward slash operator `\` is used to solve any linear system of equations. The .m file was run from a terminal using MATLAB rather than from the MATLAB GUI. All meshes were structured and homogeneous in element type, they were constructed from either, four noded bi-linear quadrilateral elements, eight noded bi-quadratic quadrilateral elements, or three noded constant strain triangular elements. For all elements the degrees of freedom existed on the nodes. The number of Gauss points for the area and face integral is displayed in Table 3.

**Table 2**

Reshaping of `glob_2` into a vector form `glob_2_rs` for MATLAB function `sparse`, where `nf` is the number of faces in the mesh.

<code>pos_i</code>	<code>neg_j</code>	<code>glob_2_rs</code>
1	1	<code>glob_2(1,1,1)</code>
1	2	<code>glob_2(1,1,2)</code>
⋮	⋮	⋮
1	<code>ndof</code>	<code>glob_2(1,1,ndof)</code>
⋮	⋮	⋮
<code>ndof</code>	<code>ndof</code>	<code>glob_2(1,ndof,ndof)</code>
<code>ndof+1</code>	<code>ndof+1</code>	<code>glob_2(2,1,1)</code>
<code>ndof+1</code>	<code>ndof+2</code>	<code>glob_2(2,1,2)</code>
⋮	⋮	⋮
⋮	⋮	⋮
<code>tndof</code>	<code>tndof</code>	<code>glob_2(nf,ndof,ndof)</code>

**Table 3**

The number of Gauss points required for the area and face integral for different element types.

Element type	# area Gauss points	# face Gauss points
Constant strain triangle	1	2
Bi-linear quadrilateral	4	2
Bi-quadratic quadrilateral	9	3

**Table 4**

Computer specifications for blocking experiments.

Component	Computer 1	Computer 2
Family	Piledriver	Haswell
Brand	AMD A10-5800K	Intel Core i7-4790
Frequency	3.8 GHz	3.60 GHz
No. cores	4 (no multithread)	4 (no multithread)
L2 cache	2 × 2 Mb	4 × 256 kb
L3 cache	N/A	8 Mb
RAM	8 GB	16 GB
OS	Ubuntu 14.04.1	Ubuntu 15.04
MATLAB vers.	R2014a	R2014b

When testing a range of SIPG face block sizes the order of the block sizes was randomised and tested, this process was repeated. All blocking experiments were performed on both computers specified in Table 4. Numerical analysis verification and speed tests, Section 4.4, were performed only using Computer 1.

#### 4.1. Variables of vectorised multiplication

MATLAB incorporates LAPACK, which calls BLAS, to perform its mathematical computations. LAPACK is a library of numerical linear algebra routines written in Fortran [20]. Arrays in Fortran are stored in column-major order form, this section investigates the importance of the orientation of variables in MATLAB when performing large vector calculations.

A script was written to investigate the speed differences when performing vector calculations in different array orientations in MATLAB, Fig. 6. Column-major operations occurred on line 7, and row-major operations on line 19. The `for` loops on lines 1 and 13, loop through a logarithmically distributed range of vector sizes from  $10 \rightarrow 10^6$ . The loops for `i` and `j` represent the nodal loops in the face integration algorithm represented in Fig. 1. The results are shown in Fig. 7. Element-wise multiplication of arrays in column-major form is consistently and significantly faster than arrays in row-major form. The memory addresses of variables in the same column vary less than along the same row. Therefore the find and read time for a variable along a column is faster. All calculations, if possible, were therefore made to occur in this format.

#### 4.2. Optimum block size

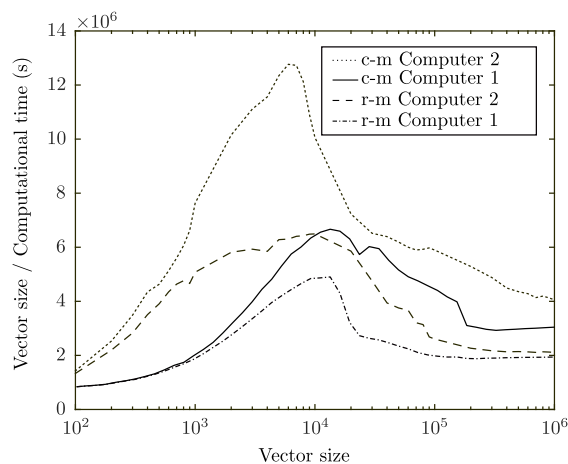
There is an optimum size of vector for an element-wise vector calculation which achieves the most floating point operations per second (flops). Managing element-wise vector operations of lengths larger than the optimum size into smaller

```

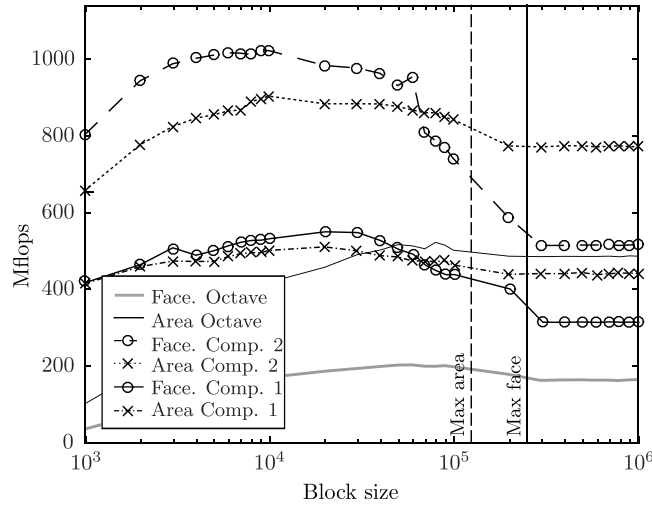
1  for s=ceil(logspace(1,6))
2      a = rand(s,1);
3      column_major = zeros(s,6,6);
4      tic
5      for i = 1:6
6          for j = 1:6
7              % Column major vector calculation
8              column_major(:,i,j)=column_major(:,i,j)+a.*a;
9          end
10         end
11     end
12     toc
13     clear column_major a
14
15     for s=ceil(logspace(1,6))
16         row_major = zeros(6,s,6);
17         b = rand(1,s);
18         tic
19         for i = 1:6
20             for j = 1:6
21                 % Row major vector calculation
22                 row_major(i,:,j)=row_major(i,:,j)+b.*b;
23             end
24         end
25     end
26     toc
27     clear row_major b
28 end

```

**Fig. 6.** A MATLAB script to investigate how the orientation of vector entry-wise multiplications are affected by array orientation.



**Fig. 7.** A computational speed comparison of performing calculations with vectors stored in a row-major or column-major form. r-m corresponds to calculations occurring in row orientated vectors and c-m corresponds to calculations operating in column orientated vector form.



**Fig. 8.** Flops performance of lines 1–5 and 6–10 of Algorithm 1 for different vector block sizes for a 2D SIPG problem with  $\approx 10^6$  degrees of freedom. Additionally Mflops performance in GNU Octave using Computer 2. Times for peak Mflops performance shown in Table 5.

**Table 5**  
Fastest computation times of the area and face integral for computers 1 and 2.

	Fastest area integral (s)	Fastest face integral (s)
Computer 1	0.25	1.04
Computer 2	0.11	0.46
Computer 2 (GNU Octave)	0.23	1.75

sizes, of optimum length, is a technique known as blocking. The vectorised SIPG code described in this paper is designed to be blocked. If the blocking algorithm for the SIPG code is effective a peak in performance corresponding to the optimum vector calculation length is expected. This section investigates whether the code has an optimum vector length, what the length is, and the number of flops achieved for this length.


To determine the performance of the optimised SIPG code, the time to calculate the linear elastic SIPG stiffness matrix  $\mathbf{K}$  was tested for different block sizes. The block size was logarithmically distributed,  $10 \rightarrow 10^6$ , and  $\mathbf{K}$  consisted of  $10^6$  degrees of freedom. The performance in terms of Mflops is presented in Fig. 8 for lines 1–5 and 6–10 of Algorithm 1, the area and face integrals. The test was performed on a 2D domain,  $\Omega$ , where  $x, y \in [0, 1]$ . The mesh distribution within  $\Omega$  is structured, constructed from bi-linear quadrilateral elements, with each element having the same area. The Young’s Modulus was set to 10 Pa and Poisson’s ratio had a value of 0.2. The tested computer architectures, OS, and MATLAB versions used, are shown in Table 4.

Computers 1 and 2 have a respective theoretical peak performance of  $3.04 \times 10^9$  and  $5.7 \times 10^9$  double precision floating point operations per second (flops). These peaks correlate to the fastest computation times achieved, shown in Table 5, and thus as their time is smallest are the optimum block sizes to perform the 2D SIPG area and face integrals.

Fig. 8 shows computer 1’s fastest performance to generate the area and face integral was  $\approx 500$  Mflops for a block size of  $\approx 3 \times 10^4$  achieving a  $\approx 16\%$  total efficiency of the theoretical peak performance. This performance was achieved through vectorisation and blocking. Whereas computer 2 achieved a higher  $\approx 1000$  Mflops for both the area and face integral corresponding to an efficiency of 17.5%. Fig. 8 demonstrates a correct implementation of a vectorised blocked algorithm to compute the SIPG global stiffness matrix with comparison to [3], since an optimum block size was achieved for both the area and face integral. The algorithm worked correctly on both computer architectures.

Using an unoptimised code, computer 1 took respectively 10.2 and 21.2 s to compute the area and face integrals, whereas Computer 2 took 5.9 and 17.91 s. Comparing the speed of the optimised code in Table 5 to the unoptimised code, computer 1 achieved a speed increase of 51 times for the area integral and 20 times for the face integral with a total speed increase of 24 times. The total speed increase from pure vectorisation is 13.7 times with blocking being 1.8 times faster than pure vectorisation. Computer 2 achieved a speed increase of 54 times for the area integral and 39 times for the face integral with a total speed increase of 41 times. The total speed increase from pure vectorisation is 23 times, a vectorised and blocked code was 1.8 times faster than pure vectorisation.

It is noted that for both computer architectures in Fig. 8 that the Mflops performance is still decreasing when the block size exceeds that of the number of area integrals and face integrals, marked by the vertical lines on Fig. 8. This suggests that for larger problems the performance gain from blocking is going to be more substantial.



**Fig. 9.** Convergence rates of 2.0 and 4.3 respectively achieved for the bi-linear (1st order), and bi-quadratic (2nd order), quadrilateral.

At the peak performance, the cache reuse is maximised. After the peak, the amount of data lying outside the CPU cache increases and so the performance decreases. A performance analysis using LIKWID [21] yielded that the combined cache throughput was maximised over all levels at the optimum block size. The performance peak for computer 1 corresponds to a larger block size compared to computer 2, Fig. 8. Since similar OS's and MATLAB versions are used, it is concluded that the shift in peak performance is caused by the L2 cache of computer 1. We found that choosing the block size as, the number of doubles that can reside in the L2 cache divided by 3 gave a performance close to the peak. The division by 3 is motivated by the fact that the typical BLAS operation requires 3 input arguments. The optimum block size can then be found through numerical experimentation from this initial point.

A speed run on Computer 2 using GNU Octave version 3.8.2 was also performed to verify that the code was effective in both MATLAB and GNU Octave. Fig. 8 demonstrated that a peak in performance was achieved for both the area and face integral. Similar to the tests performed in MATLAB, once the peak was reached the performance continued to decrease and only stabilised once the block size exceeds the number of elements and faces. The optimal performance, in comparison to MATLAB, was also slower with the area and face integrals corresponding to a loss in performance of  $\approx 2.1$  and  $\approx 3.82$  times. Despite being slower, the speed up for the vectorised blocked code when compared to an unoptimised code for the area and face integral was  $\approx 113$  and  $\approx 41$  times, much greater than that achieved with MATLAB.

The sparse formulation time was  $\approx 15$  s for computer 1 and  $\approx 9$  s for computer 2. A small investigation into whether blocking arrays whilst using native MATLAB function `sparse` had any influence on the storage time was performed, however it was found that only with computers with limited ram, (4 Gb), yielded any performance improvement. As highlighted in [3,10] using non-native MATLAB variations of the `sparse` command can significantly increase performance.

#### 4.3. Algorithm validation

To validate the correct implementation of FEA code for a linear set of equations, an eigenvalue convergence test can be performed. The test involves a 2D domain,  $\Omega$ , where  $x, y \in [0, 1]$ . The mesh distribution within  $\Omega$  is structured with each element having the same area. Homogeneous Dirichlet boundary conditions are applied on  $\partial\Omega \equiv \partial\Omega_D$ . The material stiffness matrix (28) is unusually defined as  $A = 1$ , and  $B, C = 0$ . This uncouples the degrees of freedom in the  $x$  and the  $y$  direction so that now the stiffness matrix contains two 2D Poisson problems added together with a doubled mass coefficient. The smallest eigenvalue of the problem is  $\pi^2$ . Last the SIPG penalty term is set as  $\beta = 10$ , see (2).

An undamped dynamic system of equations for linear elasticity modelled using SIPG is


$$(\mathbf{K} - \lambda_1^2 \mathbf{A}) \mathbf{U}_n = 0, \quad (36)$$

where  $\mathbf{K}$  is the global stiffness matrix,  $\mathbf{A}$  is the mass matrix, [15], and  $\lambda_1^2$  is the first natural frequency squared. The computed convergence rates were close to the analytical convergence rates for all elements, as shown in Fig. 9, [14].

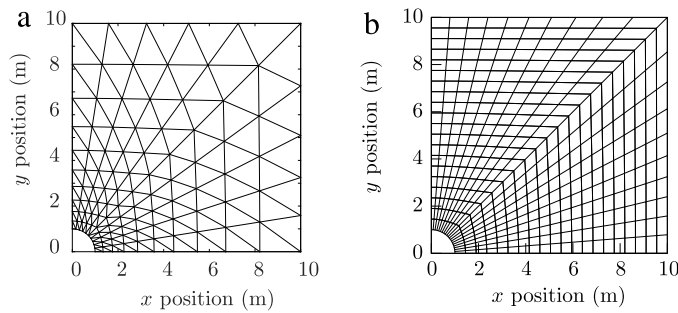
#### 4.4. Hole in plate verification

To demonstrate that the optimised 2D DG algorithm converges to correct solutions for linearly elastic problems, as well as to demonstrate the performance gains using an optimised SIPG code, a plane stress analysis of an infinite plate with a hole





**Fig. 10.** Schematic of the computational experiment. Given that the problem is symmetric, roller boundary conditions exist at  $x = 0$  and  $y = 0$ .  $a$  is the hole radius,  $\sigma_{xx}$ ,  $\sigma_{yy}$  and  $\sigma_{xy}$  are the plane and shear stresses, and  $\sigma_\infty$  is the uniform far field stress of the infinite plate.  $r$  and  $\theta$  are polar coordinates.



**Fig. 11.** The element mesh distribution for the problem in Fig. 10 with triangle elements (a) and quadrilateral elements (b).

at its centre subjected to a uniaxial tensile stress is now considered, [22]. Here the combined performance improvements, from blocking and vectorisation, of an optimised area integral as present by [3] is also analysed in conjunction with the optimised SIPG face integral presented here.

The solution to the infinite problem is provided by [23]. The infinite problem is truncated at the boundary by using the stress solution as a Neumann boundary conditions on  $\partial\Omega$ . The reduced problem setup is provided by Fig. 10. The analytical stress solution is

$$\sigma_{xx} = \sigma_\infty \left[ 1 - \frac{a^2}{r^2} \left( \frac{3}{2} \cos(2\theta) + \cos(4\theta) \right) + \frac{3a^4}{2r^4} \cos(4\theta) \right], \tag{37}$$

$$\sigma_{yy} = \sigma_\infty \left[ -\frac{a^2}{r^2} \left( \frac{1}{2} \cos(2\theta) - \cos(4\theta) \right) - \frac{3a^4}{2r^4} \cos(4\theta) \right], \tag{38}$$

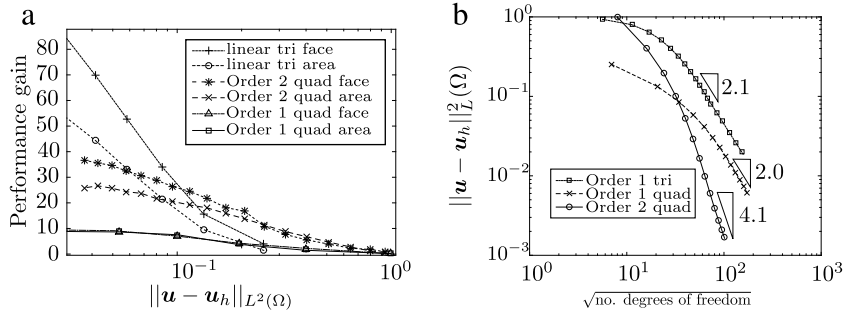
and,

$$\sigma_{xy} = \sigma_\infty \left[ -\frac{a^2}{r^2} \left( \frac{1}{2} \sin(2\theta) + \sin(4\theta) \right) + \frac{3a^4}{2r^4} \sin(4\theta) \right], \tag{39}$$

where  $\theta$  and  $r$  are polar coordinates and  $a$  is the hole radius see Fig. 10.

A schematic of the problem is shown in Fig. 10, with sides of length  $l = 10$  m and hole radius  $a = 1$  m. The material is modelled in plane stress with a Young’s modulus of  $10^3$  Pa, Poisson’s ratio of 0.2, with an applied far field stress of  $\sigma_\infty = 10^2$ .

For this problem all three element types are used: The constant strain triangle, the bi-linear quadrilateral, and the bi-quadratic quadrilateral. An example of their respective meshes is shown in Fig. 11a and b. For the constant strain triangle element and bi-linear quadrilateral element the meshes have the same number of nodes along the radius and the circumference. The bi-quadratic quadrilateral element has the same number of element vertex nodes along the circumference and radius. Along the circumference the nodes are equally spaced in terms of  $\theta$ . Along the radius,  $r$ , a scaling factor,  $sf$ , is applied to prevent distorted elements.



**Fig. 12.** (a) A performance comparison between optimised and non-optimised scripts against error for the hole in an infinite plate problem. (b) L2 convergence rates for linear triangle, and, linear and quadratic quadrilateral.

$\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\Omega)}$  is the error between the analytical solution for the displacement  $\mathbf{u}$  and the computed displacement  $\mathbf{u}_h$  and is calculated from,

$$\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\Omega)} = \sqrt{\sum_{K \in \Omega} \int_K |\mathbf{u} - \mathbf{u}_h|^2}. \tag{40}$$

This error is used in Fig. 12a and b to validate the convergence rates for different element types, [15], as well as to compare performance gains between the optimised and non-optimised codes.

Convergences rates of 2.1, 2.0 and 4.1 were achieved for the constant strain triangle, and for the bi-linear and bi-quadratic quadrilaterals using the optimised SIPG code. They are very similar to their analytical counterparts 2, 2 and 4, [15]. This demonstrates correct implementation of the optimised code for multiple elements types for a linear elastic problem.

For the speed investigation computer 1 was used, the block size of  $3 \times 10^4$  was used for all computations. For all elements the performance gain, from combined blocking and vectorisation, of the optimised code against the non-optimised code improves with problem size, Fig. 12a. The initial poor performance improvement was because at a low element number the number of BLAS calls was similar for both codes. The highest performance gain of  $\approx 90$  was achieved by the constant strain triangle elements, the lowest performance gain of  $\approx 9$  was achieved by the bi-quadratic quadrilateral elements. For both quadrilateral elements the rate of performance gain increase, decreases with problem size. For the triangular elements the rate of performance gain increase remains constant.

The ratio between the number of matrix calculation BLAS calls between the optimised and non optimised codes is similar to that of the performance gain for large problems. For the bi-quadratic quadrilateral element, an error of  $\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\Omega)} \approx 4 \times 10^{-2}$  has a BLAS call ratio  $\approx 10$  and a performance gain of  $\approx 10$  for both the area and face integral. The same can be said for the bi-linear quadrilateral when considering an error of  $\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\Omega)} \approx 4 \times 10^{-2}$ . The SIPG algorithm performance gain for the area and face integral, from blocking and vectorisation, was  $\approx 36$  and  $\approx 27$  with a corresponding BLAS call ratio of  $\approx 36$  and  $\approx 27$ .

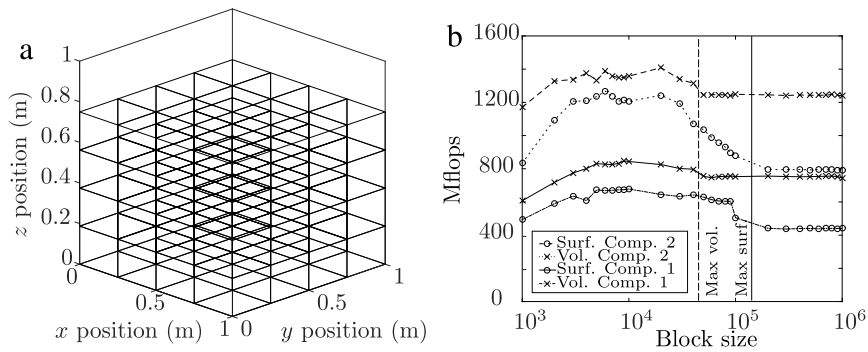
However for the triangular element this correlation breaks down. For a  $\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\Omega)} \approx 3 \times 10^{-2}$  the BLAS call ratios for the area and surface are 24.3, and 12 respectively. This is far below the performance gain in Fig. 12a. This is likely due to the BLAS call number for the triangle element being 500, which is significantly smaller than the number of BLAS calls for both quadrilateral elements which exceed 5000. For the triangle code there is no correlation between the performance gain and BLAS call number, this would indicate that the speed of the optimised for the triangle codes is longer dictated by the BLAS overhead, whereas the optimised quadrilateral codes are.

#### 4.5. 3D verification

Here a unit sided cube exists in a reference 3D Cartesian coordinate system where  $[x, y, z] \in \mathbb{R}^3$ . The cube is modelled using the SIPG method described in Section 2.1. Roller boundary conditions exist on all faces except where  $z = 1$ ; here a displacement of  $dw = -0.25$  m is applied in the  $z$  direction. The material has a Young's modulus of 1 Pa and Poisson's ratio of 0.2.

The displacements  $u, v, w$  correspond to the directions of  $x, y, z$ . The analytical solution to the constant stress problem is  $u, v = 0$  and  $w = z \times dw$ . Any mesh discretisation would achieve the correct solution to machine precision. Here a homogeneous  $5 \times 5 \times 5$  mesh of tri-linear hexahedral elements was used to show the algorithm running correctly, the result is shown in Fig. 13a.

To show that the blocked script could also be applied to 3D SIPG problem, a problem consisting of  $10^6$  degrees of freedom is run. The problem consists of a unit cube constructed from a uniform distribution of linear hexahedral elements. The volume



**Fig. 13.** (a) Compression of a unit cube with roller boundary conditions. (b) Mflops performance of lines 1–5 and 6–10 of Algorithm 1 for different block sizes for a 3D SIPG problem with  $\approx 10^6$  degrees of freedom.

and surface integrals require 8 and 4 Gauss points respectively. Fig. 13b shows that it is possible to block vectorised 3D SIPG code for both volume and surface integrals. For both computers there is a peak in performance at a block size of  $\approx 10^4$ . The peak corresponds to the cache reuse being maximised. The performance drop after the peak corresponds to not all data, required for a BLAS operation, residing in the cache. Similar to Section 4.2 the performance is still decreasing once the block size becomes larger than the number of volume and surface integrals, highlighting the importance of blocking for larger problems.

The fastest runtime achieved for the volume and surface integral by computer 1 was respectively 4.1 s and 9.2 s, computer 2 achieved a run time of 1.4 s and 3.9 s. The total runtime to generate the global stiffness matrix was 29.3 s and 20.2 s for computers 1 and 2. Profiling revealed the MATLAB function `sparse`, necessary to generate the global stiffness matrix, took the majority of the time 18 s and 13 s for computers 1 and 2.

## 5. Conclusion

This paper has presented for the first time an efficient blocked vectorised algorithm for producing SIPG face stiffness terms in a native MATLAB environment for linear elasticity for a range of elements in both 2D and 3D. Optimisation was achieved through: (i) vectorising the local face stiffness matrix equations; (ii) optimising the vectorised calculations by maximising the CPU cache reuse; (iii) storing vectors in a column-major form; (iv) ensuring all matrix calculations were as large as possible; and (v) reducing the number of calculations by only considering symmetric terms.

The block length optimisation results demonstrate a clear optimal block length, the performance peak coincides with a maximisation of the cache reuse. For all elements types, both 2D and 3D, the optimised codes were able to compute the global stiffness matrix for a system comprising of  $10^6$  DOF system in under 30 s. It was found when testing the optimum block size for  $10^6$  DOF, in Section 4.2, that vectorisation was the main contributor to the performance improvements. Vectorisation alone achieved performance achievements between 13.7 and 23 times, whilst blocking improved the vectorised algorithm by a further 1.8 times. The theoretical peak performance ranged from 16% to 17.5% across the computer architectures, similar to results found in the literature [3]. Additionally a number of different verification techniques have been used to demonstrate correct implementation of both linear systems in both 2D and 3D.

In the linear elastic 2D performance gain study, in Section 4.4, it was shown that the performance gain, between an optimised and an unoptimised code, continues to increase with problem size. It was also shown that the performance gains were dependent on element type, with triangular elements achieving gains excess of 50 times. There was also a correlation between the performance gain and the ratio of BLAS calls, between the unoptimised and optimised codes, for the quadrilateral elements. This suggests that the optimised quadrilateral code is still subject to a bottleneck from the BLAS call overhead. This was not the case for the triangular code which had significantly fewer BLAS calls in the optimised code.

The script could be optimised further by using the MATLAB's parallel function `parfor` and incorporating GPUs into the calculation. The final scripts are designed to be a black box, taking in element topology and outputting the global stiffness matrix for a SIPG problem.

## Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/K502832/1]. All data created during this research is openly available at <https://doi.org/10.15128/r1xs55mc04f>.

## Appendix. List of variables

Name	Dimensions	Description
Ai	1	The degree of freedom row positioning of entry (1,1) of partial stiffness matrices $M_{C1}^r, M_{C2}^r, M_{C3}^r, M_{C4}^r, M_{E1}^r, M_{E2}^r, M_{E3}^r$ , and $M_{E4}^r$ , for current $i$ .
Aj	1	The degree of freedom column positioning of entry (1,1) of partial stiffness matrices $M_{C1}^r, M_{C2}^r, M_{C3}^r, M_{C4}^r, M_{E1}^r, M_{E2}^r, M_{E3}^r$ , and $M_{E4}^r$ , for current $j$ .
Bi	1	The degree of freedom row positioning of entry (1,1) of partial stiffness matrices $M_{D1}^r, M_{D2}^r, M_{D3}^r$ , and $M_{D4}^r$ , for current $i$ .
Bj	1	The degree of freedom column positioning of entry (1,1) of partial stiffness matrices $M_{D1}^r, M_{D2}^r, M_{D3}^r$ , and $M_{D4}^r$ , for current $j$ .
bl_index	[1, nel_block]	An index for selecting rows of etpl_face and glob_2 which are in the current SIPG face block loop.
Block_n	1	Current SIPG face block number.
C2t_11	[ nel_block,1]	Vector of components of entry (1,1) of $M_{C2}^r$ which vary with $i$ , for all faces in the current SIPG face block loop.
coord	[ nnodes,2]	Coordinates of all nodes in the mesh.
coord_xp	[sum(index_p),1]	Nodal $x$ -coordinates of all elements in the current SIPG face block with local face number $fn$ . Their order is dictated by etpl_face(:, 1).
coord_yp	[sum(index_p),1]	Nodal $y$ -coordinates of all elements in the current SIPG face block with local face number $fn$ . Their order is dictated by etpl_face(:, 1).
D2t_12	[ nel_block,1]	Vector of components of entry (1,2) of $M_{D2}^r$ , which vary with $j$ , for all faces in the current SIPG face block loop.
det	[sum(index_p),1]	Jacobian determinant for all $K^+$ elements in the current SIPG face block loop with local face number $fn$ .
dNr	[ nf* ngp*2, nen]	Reference shape function derivatives for all Gauss points for all local element faces.
dNx_p	[ nel_block, ndof]	Global shape function derivatives, with respect to $x$ , for all $F^+$ faces in the current loop.
dNy_p	[ nel_block, ndof]	Global shape function derivatives, with respect to $y$ , for all $F^+$ faces in the current loop.
ed	[ nels, ndof]	Steering matrix $\forall K \in \mathcal{T}$ . Row number corresponds to element, column number to the global degree of freedom.
ed_n	[ tot_f, ndof]	Steering matrix to steer the local face stiffness matrices into the global matrix. Corresponds to degrees of freedom of $K^-$ , ordered as in etpl_face(:, 2).
ed_p	[ tot_f, ndof]	Steering matrix to steer the local face stiffness matrices into the global matrix. Corresponds to degrees of freedom of $K^+$ , ordered as in etpl_face(:, 1).
etpl	[ nels, nen]	Element topology matrix of all elements in the mesh.
etpl_face	[ tot_f,7]	Description of the SIPG faces in the mesh.
etpl_face_block	[ nel_block,7]	Columns of etpl_face for SIPG faces in the current block number.
etpl_face_block_fn	[ sum(index_p),7]	Columns of etpl_face_block for SIPG faces with local positive face number $K^+$ .
fn	1	Current face number.
glob_2	[ tot_f, ndof, ndof]	Storage variable for $G_2$ .
glob_2_rs	[ tot_f*ndof,1]	glob_2 reshaped into a vector form.
glob_pn	[ nel_block, ndof, ndof]	Temporary storage variable for all SIPG faces matrices that are respectively pre-and-post multiplied by a $K^+$ and $K^-$ element.
gp	1	Current Gauss point number in the loop.
i	1	Row node number for partial stiffness matrices: $M_{C1}^r, M_{C2}^r, M_{C3}^r, M_{C4}^r, M_{E1}^r, M_{E2}^r, M_{E3}^r$ , and $M_{E4}^r$ .

(continued on next page)

(continued)

Name	Dimensions	Description
i2	1	Row node number for partial stiffness matrices: $M_{D1}^r$ , $M_{D2}^r$ , $M_{D3}^r$ , and $M_{D4}^r$ .
indx_Nr	1	Index to select row of Nr for a specific fn and gp.
index_n	[nel_block],1	Logical variable to select columns of etp1_face_block, for $K^-$ elements, with the face number fn. 1 indicates same fn, 0 otherwise.
index_p	[nel_block],1	Logical variable to select columns of etp1_face_block, for $K^+$ elements, with the face number fn. 1 indicates same fn, 0 otherwise.
indx_dNr	[1,2]	Index to select rows of dNr for a specific fn and gp.
int_W	[nel_block],1	Gauss face integral weight and Jacobian determinant for all SIPG faces in the current block loop.
invJxp	[sum(index_p),2]	Row 1 of inverse Jacobian matrix for all $K^+$ elements in the current SIPG face block loop with local face number fn.
invJyp	[sum(index_p),2]	Row 2 of inverse Jacobian matrix for all $K^+$ elements in the current SIPG face block loop with local face number fn.
j	1	Column node number for partial stiffness matrices: $M_{C1}^r$ , $M_{C2}^r$ , $M_{C3}^r$ , $M_{C4}^r$ , $M_{E1}^r$ , $M_{E2}^r$ , $M_{E3}^r$ , and $M_{E4}^r$ .
j2	1	Column node number for partial stiffness matrices: $M_{D1}^r$ , $M_{D2}^r$ , $M_{D3}^r$ , and $M_{D4}^r$ .
jxp	[2,sum(index_p)]	Column 1 of Jacobian matrix for all $K^+$ elements in the current SIPG face block loop with local face number fn.
jyp	[2,sum(index_p)]	Column 2 of Jacobian matrix for all $K^+$ elements in the current SIPG face block loop with local face number fn.
K	[max(ed(:)), max(ed(:))]	Global stiffnessmatrix.
ndof	1	Total number of degrees of freedom for one element.
neg_i	[tot_f*nndof],1	Row degrees of freedom for all local stiffness matrices pre-multiplied by a $K^-$ element. Corresponds to global storage vectors glob_3_rs and glob_4_rs.
neg_j	[tot_f*nndof],1	Column degrees of freedom for all local stiffness matrices post-multiplied by a $K^-$ element. Corresponds to global storage vectors glob_2_rs and glob_4_rs.
nel_block	1	Total number of faces in the block.
nels	1	Total number of elements in the mesh.
nen	1	Number of nodes for one element.
nf	1	Number of faces on an element.
ngp	1	Number of face Gauss points.
nndof	1	Total number of entries in local element matrix (ndof*ndof).
nnodes	1	Total number of nodes in the mesh.
Np	[nel_block],1	Face shape functions for all $K^+$ element faces in the current block.
Nr	[nf*ngp,nen]	Local face shape functions.
num_blocks	1	Number of SIPG face blocks.
num_faces	1	Number of SIPG faces in a block.
nx	[nel_block],1	Normal x component to all interior faces in the block.
ny	[nel_block],1	Normal y component to all interior faces in the block.
pen	1	SIPG penalty values for linear elasticity.
pos_el	[sum(index_p),1]	List of $K^+$ elements in the SIPG face block loop with face number fn.
pos_i	[tot_f*nndof],1	Row degrees of freedom for all local stiffness matrices pre-multiplied by a $K^+$ element. Corresponds to global storage vectors glob_1_rs and glob_2_rs.
pos_j	[tot_f*nndof],1	Column degrees of freedom for all local stiffness matrices post-multiplied by a $K^+$ element. Corresponds to global storage vectors glob_1_rs and glob_3_rs.
tndof	1	Total number of degrees of freedom in mesh.
tot_f	1	Total number of interior faces.

## References

- [1] W.M. Coombs, R. Crouch, C. Augarde, 70-line 3D finite deformation elastoplastic finite-element code, Proc. Numerical Methods in Geotechnical Engineering (NUMGE), Trondheim, Norway, 2010, pp. 151–156.
- [2] O. Sigmund, A 99 line topology optimization code written in MATLAB, Struct. Multidiscip. Optim. 21 (2) (2001) 120–127.
- [3] M. Dabrowski, M. Krotkiewski, D. Schmid, MILAMIN: MATLAB-based finite element method solver for large problems, Geochem. Geophys. Geosyst. 9 (4) (2008).
- [4] M. Krotkiewski, M. Dabrowski, Parallel symmetric sparse matrix–vector product on scalar multi-core CPUs, Parallel Comput. 36 (4) (2010) 181–198.
- [5] T. Rahman, J. Valdman, Fast MATLAB assembly of FEM stiffness- and mass matrices in 2D and 3D: nodal elements, Appl. Math. Comput. 219 (13) (2013) 7151–7158.
- [6] I. Anjam, J. Valdman, Fast MATLAB assembly of FEM matrices in 2D and 3D: Edge elements, Appl. Math. Comput. 267 (2015) 252–263.
- [7] E. Andreassen, A. Clausen, M. Schevenels, B.S. Lazarov, O. Sigmund, Efficient topology optimization in MATLAB using 88 lines of code, Struct. Multidiscip. Optim. 43 (1) (2011) 1–16.
- [8] F. Cuvelier, C. Japhet, G. Scarella, An efficient way to perform the assembly of finite element matrices in MATLAB and Octave, 2013, In MATLAB and Octave, Preprint, University of Paris 13 and INRIA.
- [9] F. Cuvelier, C. Japhet, G. Scarella, An efficient way to assemble finite element matrices in vector languages, BIT Numer. Math. (2015) 1–32.
- [10] S. Engblom, D. Lukarski, Fast MATLAB compatible sparse assembly on multicore computers, Parallel Comput. 56 (2014) 1–17.
- [11] W.H. Reed, T. Hill, Triangular mesh methods for the neutron transport equation, Los Alamos Report LA-UR-73-479, 1973.
- [12] G.R. Richter, The discontinuous Galerkin method with diffusion, Math. Comp. 58 (198) (1992) 631–643.
- [13] F. Bassi, S. Rebay, A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier–Stokes equations, J. Comput. Phys. 131 (2) (1997) 267–279.
- [14] S. Giani, E.J. Hall, An a posteriori error estimator for hp-adaptive discontinuous Galerkin methods for elliptic eigenvalue problems, Math. Models Methods Appl. Sci. 22 (10) (2012) 1250030.
- [15] N.S. Ottosen, H. Petersson, Introduction to the Finite Element Method, Prentice Hall International, 1992.
- [16] D. Arnold, F. Brezzi, B. Cockburn, L. Marini, Unified analysis of discontinuous Galerkin methods for elliptic problems, SIAM J. Numer. Anal. 39 (5) (2002) 1749–1779.
- [17] F. Frank, B. Reuter, V. Aizinger, P. Knabner, FESTUNG: A MATLAB/GNU Octave toolbox for the discontinuous Galerkin method, Part I: Diffusion operator, Comput. Math. Appl. 70 (1) (2015) 11–46.
- [18] R. Bird, W. Coombs, S. Giani, Vectorised SIPG matrix formulation for MATLAB and Octave, [https://github.com/robertbirdurham/MATLAB\\_OPTIMISED.git](https://github.com/robertbirdurham/MATLAB_OPTIMISED.git), 2017. Accessed: 28.03.17.
- [19] P. Hansbo, M.G. Larson, Energy norm a posteriori error estimates for discontinuous Galerkin approximations of the linear elasticity problem, Comput. Methods Appl. Mech. Engrg. 200 (45) (2011) 3026–3030.
- [20] C. Moler, MATLAB Incorporates LAPACK, <http://uk.mathworks.com/company/newsletters/articles/matlab-incorporates-lapack.html>, 2015. Accessed: 25.09.15.
- [21] J. Treibig, G. Hager, G. Wellein, LIKWID: A lightweight performance-oriented tool suite for x86 multicore environments, in: Proceedings of PSTI2010, the First International Workshop on Parallel Software Tools and Tool Infrastructures, San Diego CA, 2010.
- [22] Z. Stowell, Elbridge, Stress and strain concentration at a circular hole in an infinite plate, 1950.
- [23] G. Kirsch, Die Theorie der Elastizität und Die Bedürfnisse der Festigkeitslehre, Springer, 1898.