

13. Algorithms and arteries: Multi-domain spectral/*hp* methods for vascular flow modelling

S.J. Sherwin¹ and J. Peiró²

1. Introduction. The growing interest in the mathematical and numerical modelling of biomedical systems and, in particular, the human cardiovascular system is supported by the numerous works which have appeared on the subject in recent years, for example [2, 9, 16, 21] and the references therein. Traditionally there has been a strong focus on low dimensional models [13]. However the association of vascular disease, such as atherosclerosis, with arterial branching has promoted an interest in the application of computational fluid dynamics (CFD) to vascular flow modelling. Nevertheless the nature of the flow presents a variety of challenges. Firstly, the flow is pulsatile and in a Reynolds number regime where the viscous and inertial effects are both significant. Secondly, the geometric characteristics of the vascular system are very intricate. Finally, blood is a non-Newtonian fluid and arterial walls are distensible.

A particular focus of the CFD modelling has been to determine the wall shear stress distribution of the unsteady flow at arterial junctions and bypass grafts. The sensitivity of the wall shear stress to surface curvature therefore make the geometric representation an important factor. The flows in and around regions of stagnation and separation are also of physiological interest. In an incompressible flow we know that the normal derivative of the wall normal flow is zero near the wall and therefore requires at least a second order approximation to be resolved. Both these factors and the requirement to reproduce the unsteady flow and its derivatives in complex geometries make high-order algorithms, such as the unstructured spectral/*hp* element method, suitable from the point of view of attaining a specified error at a lower computational cost. However the problem still poses many numerical challenges which have motivated a range of developments in spectral/*hp* element methods that we shall discuss in this review article.

Furthermore we cannot completely decouple the local branching flow at an arterial junction from the full vascular system. The flow waveform observed at a given location in the vascular tree is the result of changes in sectional area of the compliant vessels to accommodate the incompressible flow of blood as it is pumped from the heart. Starting at the heart, the arterial waves are propagated and reflected at each arterial branch [25] leading to a complex waveform which changes at different locations. Although the wavelengths of these waves are much larger than the length of local arterial branches, the flow waveform can be altered by the presence of disease or surgical intervention. Therefore there is an inherent need to include a multiscale modelling to the localised CFD as discussed in [16]. Within this context, the application of simplified models has been shown to provide useful information for practitioners at a reasonable computational cost [7].

In this paper we will briefly review three topics related to the application of spectral/*hp* discretisation to vascular flow modelling. In section 2 we review the work in [19] and discuss the one-dimensional full vascular tree modelling using one-dimensional equations. In section 3 we discuss the problem of generating high-order meshes to consistently model the arterial geometries based on the work in [22, 14]. Finally in section 4 we overview a recent development in elliptic preconditioning for unstructured model spectral/*hp* methods based on a low energy numerical basis which relates to the work of [1, 18].

2. Reduced 1D modelling of the human circulation. In this section we focus on the application of a one-dimensional model of blood flow in compliant vessels to study wave propagation in the arterial tree as previously detail in [19].

¹Department of Aeronautics, Imperial College, London, U.K., s.sherwin@ic.ac.uk

²Department of Aeronautics, Imperial College, London, U.K., j.peiro@ic.ac.uk

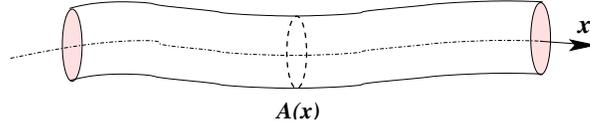


Figure 2.1: Simple compliant tube.

2.1. Governing equations. We consider a simple compliant tube, illustrated in figure 2.1, as a model of the artery. Following Brook et al. [3] we write the system of equations representing continuity of mass and momentum, for $a \leq x \leq b$ and $t > 0$, as

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x}(\mathbf{U}) = \frac{\partial}{\partial t} \begin{bmatrix} A \\ u \end{bmatrix} + \frac{\partial}{\partial x} \begin{bmatrix} Au \\ \frac{u^2}{2} + \frac{p}{\rho} \end{bmatrix} = \begin{bmatrix} 0 \\ -K_R u \end{bmatrix} \quad (2.1)$$

where the x is the axial direction, $A = A(x, t) = \int_S d\sigma$ is the area of a cross section S , ρ is the density of the blood which is taken to be constant, p is the internal pressure and $u(x, t)$ denotes the velocity of the fluid averaged across the section. The term K_R is a strictly positive quantity which represents the viscous resistance of the flow per unit length of tube. The unknowns in this system are p , A and u . Their number exceeds the number of equations and a common way to close the system is to explicitly provide an algebraic relationship between the pressure of the vessel p and the vessel area A . For example, by assuming *static equilibrium* in the radial direction of a cylindrical tube, one can derive a pressure relationship of the form

$$p = p_{ext} + \beta(\sqrt{A} - \sqrt{A_0}), \quad (2.2)$$

where

$$\beta = \frac{\sqrt{\pi} h_0 E}{(1 - \nu^2) A_0}.$$

Here h_0 and $A_0 = A_0(x)$ denote the vessel thickness and sectional area, respectively, at the equilibrium state $(p, u) = (p_{ext}, 0)$, $E = E(x)$ is the Young modulus, p_{ext} is the external pressure, assumed constant and ν is the Poisson ratio. This ratio is typically taken to be $\nu = 1/2$ since biological tissue is practically incompressible.

2.2. Discontinuous Galerkin method. The wave propagation speeds in the large arteries are typically an order of magnitude higher than the average flow speeds. The characteristic speed of the system is also inherently subcritical and does not produce shock under physiological conditions. Therefore the numerical challenge is to propagate waves for many periods without suffering from excessive dispersion and diffusion errors. If the solution remains smooth then high-order methods are particularly attractive due to the fast convergence of the phase and diffusion properties with order of the scheme [17].

Following the work of Cockburn and Shu [4] we initially consider the one-dimensional hyperbolic system in conservative form (2.1) and assume that $R_K = 0$. To solve this system in a region $\Omega = [a, b]$ discretised into a mesh of N_{el} elemental non-overlapping regions $\Omega_e = [x_e^l, x_e^r]$, such that $x_e^r = x_{e+1}^l$ for $e = 1, \dots, N_{el}$. We then proceed by constructing the weak form of (2.1), i.e.

$$\left(\frac{\partial \mathbf{U}}{\partial t}, \boldsymbol{\psi} \right)_{\Omega} + \left(\frac{\partial \mathbf{F}}{\partial x}, \boldsymbol{\psi} \right)_{\Omega} = 0 \quad i = 1, 2 \quad (2.3)$$

where $(\mathbf{u}, \mathbf{v})_{\Omega} = \int_{\Omega} \mathbf{u} \cdot \mathbf{v} \, dx$ is the standard $\mathbf{L}^2(\Omega)$ inner product. Decomposing the integral into elemental regions we obtain

$$\sum_{e=1}^{N_{el}} \left[\left(\frac{\partial \mathbf{U}}{\partial t}, \boldsymbol{\psi} \right)_{\Omega_e} + \left(\frac{\partial \mathbf{F}}{\partial x}, \boldsymbol{\psi} \right)_{\Omega_e} \right] = 0. \quad (2.4)$$

Integrating the second term in (2.4) by parts leads to

$$\sum_{e=1}^{N_{el}} \left(\frac{\partial \mathbf{U}}{\partial t}, \boldsymbol{\psi} \right)_{\Omega_e} - \left(\mathbf{F}, \frac{d\boldsymbol{\psi}}{dx} \right)_{\Omega_e} + [\boldsymbol{\psi} \cdot \mathbf{F}]_{x_e^u}^{x_e^l} = 0 \quad (2.5)$$

To form the discrete approximation of our problem we choose \mathbf{U} to be in the finite space of $\mathbf{L}^2(\Omega)$ functions which are polynomials of degree P on each element. Furthermore we indicate an element of such space using the superscript δ . To attain a global solution in the domain Ω we need to allow information to propagate between the elemental regions. Information is propagated between elements by upwinding the boundary flux in the third term of equation (2.5). Denoting the upwinded flux as \mathbf{F}^u , the discrete weak formulation can now be written as

$$\sum_{e=1}^{N_{el}} \left(\frac{\partial \mathbf{U}^\delta}{\partial t}, \boldsymbol{\psi}^\delta \right)_{\Omega_e} - \left(\mathbf{F}(\mathbf{U}^\delta), \frac{d\boldsymbol{\psi}^\delta}{dx} \right)_{\Omega_e} + [\boldsymbol{\psi}^\delta \cdot \mathbf{F}^u]_{x_e^l}^{x_e^u} = 0, \quad (2.6)$$

Following the traditional Galerkin approach, we choose the test function within each element to be in the same discrete space as the numerical solution \mathbf{U}^δ . At this point if we defined our polynomial basis and choose an appropriate quadrature rule we would now have a semi-discrete scheme. However, from an implementation point of view, the calculation of the second term can be inconvenient and consequently we choose to integrate this term by parts once more to obtain

$$\sum_{e=1}^{N_{el}} \left(\frac{\partial \mathbf{U}^\delta}{\partial t}, \boldsymbol{\psi}^\delta \right)_{\Omega_e} + \left(\frac{\partial \mathbf{F}(\mathbf{U}^\delta)}{\partial x}, \boldsymbol{\psi}^\delta \right)_{\Omega_e} + [\boldsymbol{\psi}^\delta \cdot [\mathbf{F}^u - \mathbf{F}(\mathbf{U}^\delta)]]_{x_e^l}^{x_e^u} = 0. \quad (2.7)$$

We note that the information between elements is transmitted by the third boundary term as the difference between the upwinded and the local fluxes, $[\boldsymbol{\psi}^\delta \cdot [\mathbf{F}^u - \mathbf{F}(\mathbf{U}^\delta)]]_{x_e^l}^{x_e^u}$.

To complete the discretisation we also require a time integration scheme and in the current implementation we have adopted a second order Adams-Bashforth scheme. The upwind flux is calculated using a straightforward upwinding of the characteristic variables as discussed in [19]. This type of upwinding process is used to impose the characteristic boundary conditions through the flux at the ends of the global domain Ω .

The 1D model of the compliant tube can be extended to handle the arterial tree by imposing suitable interface conditions at the bifurcations or branching points of the tree. At a bifurcation we have six degrees of freedom corresponding to the area and velocity conditions within each vessel. Therefore we require six equations to determine a unique solution. Applying the subsonic flow assumption we can determine the three characteristics entering the junction providing three equations. Finally, continuity of mass flux and total pressure at the bifurcation provide the three conditions required to close the system, see [19] for details.

2.3. Simulation of wave propagation in the arterial network. We have adopted the modifications proposed in [25] to the published models [26, 24] to compute the pulsatile one-dimensional blood flow through the arterial system using the discontinuous Galerkin method. The numerical values of the parameters of the arterial tree formed by the 55 main arteries is given in [19]. Figure 2.2 shows the connectivity of the arteries used in our model of the arterial network. The flow in the 55 arteries is assumed initially to be at rest. A periodic half sine wave is imposed as an input wave form at the ascending aorta (artery 1). Figure 2.2 also shows the inflow boundary conditions imposed at the ascending aorta and the time history graphs over a single cycle for three different arteries in the network: ascending aorta (artery 1), femoral artery (artery 46) and anterior tibial (artery 49).

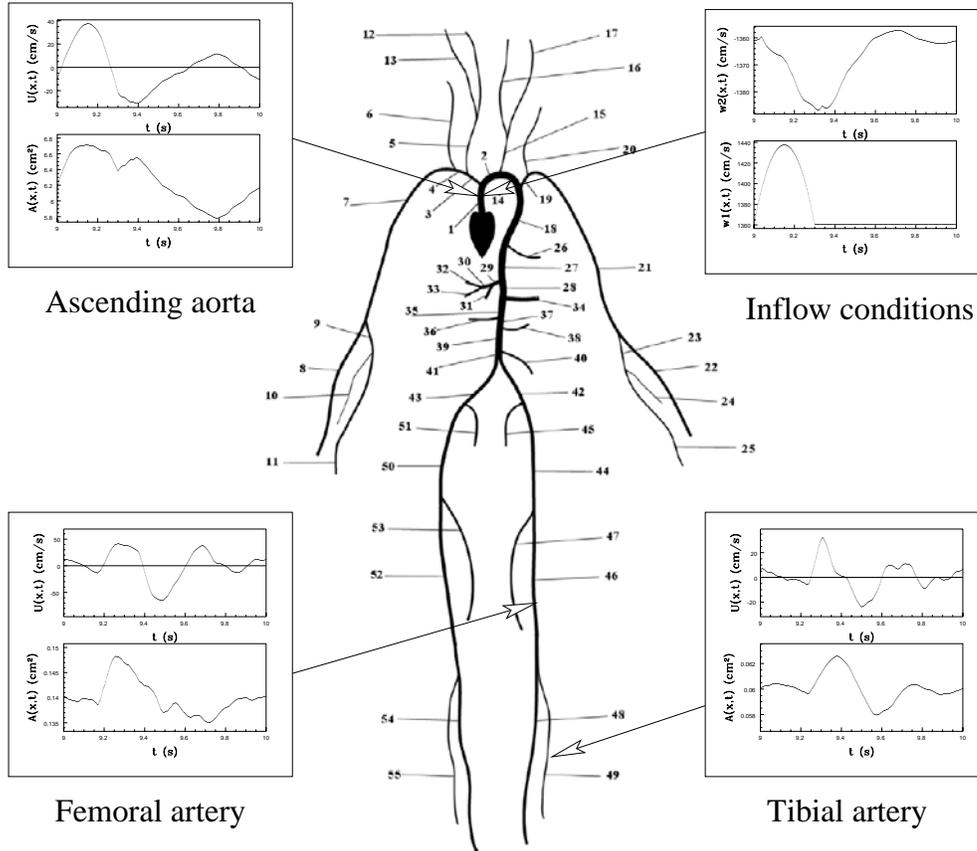


Figure 2.2: Connectivity of the 55 main arteries in the human arterial system.

The inclusion of resistance in the terminal arteries increases the number of waves in the system due to forward travelling waves being reflected at the terminal vessels and introduces backward travelling waves which are re-reflected at the bifurcations, hence a complex pattern of waves occurs in the network. Terminal resistance also creates regions of flow reversal due to the reflected velocity wave and increases in area as a result of the re-enforcing effect of the reflected pressure wave. It has also produced a waveform which includes a diastolic notch in the ascending aorta (artery 1). These results are qualitatively similar to what we would expect to see from in-vivo measurements in the human body.

3. Geometric modelling of arterial branching. The ability to construct suitable computational meshes is currently a significant limiting factor in the development of high-order algorithms in very complex geometries. In this section we will address the issues encountered in applying the high-order finite element type approach to vascular flow modelling as previously discussed in [22, 14].

3.1. Mesh generation of high-order elements. The extension of standard unstructured mesh generation technology to high-order algorithms is a not trivial exercise. Complications arise due to the conflicting requirements to generate coarse meshes whilst maintaining good elemental properties in regions of high curvature. This is shown in figure

3.1 where we illustrate the type of invalid elements which can arise.

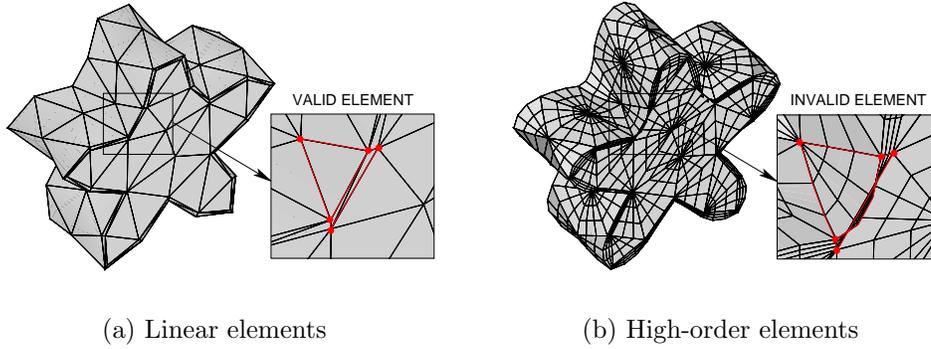


Figure 3.1: The subdivision of a valid mesh of linear elements (a) to generate a high-order tetrahedral mesh (b) might lead to elemental regions with singular Jacobian mappings.

In our approach, the generation of an unstructured mesh of high-order spectral/*hp* elements is accomplished through the subdivision of a coarse mesh of linear elements. Given a surface representation in terms of cubic splines, the surface is initially discretised into a coarse distribution of linear surface elements. The local topology of these linear elements is influenced by the desire to include a boundary layer region or by taking into account surface curvature as described in section 3.3. The mesh generation then proceeds in a manner consistent with standard linear mesh generation process. Our current approach is based upon the method of advancing layers described in [15] but alternative mesh generation techniques can also be used. In this method the vertices of the original linear triangulation in the near-wall regions are assigned a direction and new interior vertices are created in successive layers up to a prescribed boundary layer thickness. These points are then linked to form a mesh of tetrahedral or prismatic elements, known as the boundary layer mesh. The rest of the domain is finally filled with a mesh of linear tetrahedra which, in our case, is generated by means of the advancing front technique.

A high-order surface discretisation is generated by following a “bottom-up” procedure where initially the triangular edges are discretised into $P+1$ points for a P^{th} order polynomial mesh. Subsequently the $(P-3)(P-2)/2$ points internal to the triangular faces are generated to complete the polynomial representation. The high-order point generation is typically performed in the parametric space of the surface splines which may have a non-isometric mapping to the physical space. In order to optimise the high-order element point distribution, a non-linear minimisation procedure is adopted, as discussed in [22], which generates the edge and face points as geodesics of the surface with a view to minimising the variation in the surface Jacobian.

3.2. Optimizing surface representation. To address the problem of obtaining an optimal distribution of points, consider a quadrature with N integration points and associated normalised weights z_i ; $i = 1, \dots, N$ ($-1 \leq z_i \leq 1$) in a 1D interval $a \leq x \leq b$. It is known that the optimal positions x_i ; $i = 1, \dots, N$ of the points are given by

$$x_i = \frac{a}{2}(1 - z_i) + \frac{b}{2}(1 + z_i) \quad i = 1, \dots, N. \quad (3.1)$$

This leads to an isometric mapping and therefore a constant Jacobian. The extension to elements with straight sides and faces is straightforward. However, a different strategy for

curved edges and faces is required to account for the distortion introduced by surface curvature.

A parametric surface is a mapping between a 2D parametric space and the 3D space. Mesh generation is considerably simplified if performed in the parametric space. However, approximating the element edges using the point distribution (3.1) along straight lines in the parametric plane might lead to highly deformed or invalid elemental regions if the surface mapping induces severe distortion.

An “optimal” point distribution can be obtained by recasting the problem as that of minimising the potential energy of a set of springs linking adjacent points. It is easily shown that the optimal distribution (3.1) is a minimum of the potential energy of such system of springs given by

$$\mathcal{J}(x_2, \dots, x_{N-1}) = \sum_{i=1}^{N-1} \frac{(x_{i+1} - x_i)^2}{z_{i+1} - z_i} = \sum_{i=1}^{N-1} \frac{\delta_i^2}{z_{i+1} - z_i}. \quad (3.2)$$

This approach, unlike (3.1), is directly applicable to curved edge and faces on surfaces. A more detailed description of the procedure could be consulted in [22].

The high-order surface definition implies that the elements adjacent to a deformed wall will also have curved internal faces. These are constructed as a blend, consistent with the spectral/*hp* element expansion, between the internal straight edges and the deformed surface edge (see [11] for more details). In general, high-order elements allow for all internal face and edges to be deformed which, as discussed in the work of Dey *et al.* [5], may be necessary in very curved domains.

3.3. Curvature based mesh refinement for high-order elements. Mesh refinement as a function of the curvature has been proposed by several authors [8, 12] as a way to obtain an accurate piecewise linear approximation of a curved surface. In [22] we have shown that the use of curvature based refinement enhances the quality of the high-order meshes generated from linear tetrahedral and prismatic meshes. However, this criterion on its own is not sufficient to guarantee validity of all high-order elements as it does not account for the possible intersection of the boundary sides and faces with those on the interior. In [14] we have proposed an alternative method more suitable for the discretisation of boundary layers which we detail below.

A curve is locally approximated by a circle of radius R , the radius of curvature. We assume that the mesh spacing can be represented by a chord of length c in the circle and a spacing δ in the normal direction. In the modelling of viscous flows, the value of δ is usually prescribed to achieve a certain boundary layer resolution. The value of c is therefore chosen to guarantee that the osculating circle representing the curve does not intersect the interior sides of the elements, i.e. $\theta \geq 90^\circ$ for the triangular element. The value of c , which should be considered as a maximum mesh spacing, can now be obtained as a function of R and δ . Its value c_t for triangular elements is

$$c_t \leq R \sqrt{\frac{2\delta}{R + \delta}}. \quad (3.3)$$

The corresponding value c_q for quadrilateral elements is

$$c_q \leq \frac{2R\delta}{R + \delta} \sqrt{1 + \frac{2R}{\delta}}, \quad (3.4)$$

where the boundary displacement is assumed to be the same on either side of the rectangle. It is interesting to notice that, for a given δ , the quadratic element allows for a mesh spacing c_q which is about twice the value of spacing c_t for the triangular element.

The extension of this method to surfaces is straightforward. The refinement criterion given by formulas (3.3) and (3.4) is used for the two principal directions and the corresponding mesh spacings, c_1 and c_2 , are calculated from the values of the principal curvatures $k_{1,2} = 1/R_{1,2}$.

An example of a hybrid mesh generated for the geometry previously considered in figure 3.1 and using the criterion (3.4), is shown in figure 3.2(a). This high-order mesh does not contain singular elements. However, the refinement applied here does not account for the sign of the surface curvature and the use of criterion (3.4) to ensure element validity is too restrictive in those regions where the domain is locally convex. For a convex region, the less restrictive criterion $\delta < R$ suffices to guarantee element validity. This is highlighted in figure 3.2(b) where the refinement criterion (3.4) has been selectively applied to concave regions only. The result is a valid mesh with fewer elements.

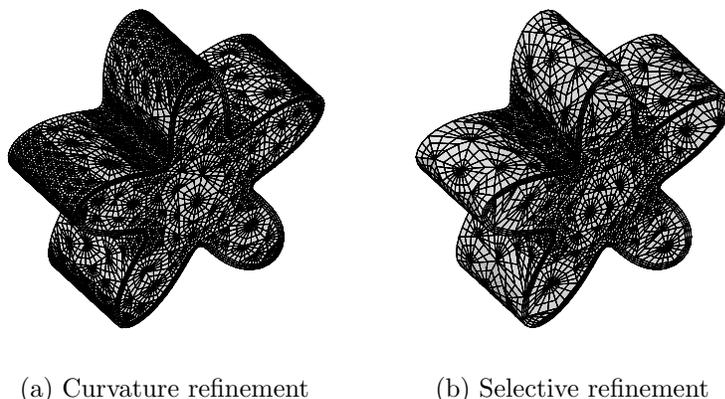


Figure 3.2: Curvature based mesh refinement for prismatic elements: (a) Refinement according to equation (3.4), (b) Refinement is applied to concave regions only.

As previously discussed our area of interest is the surgical intervention required when an artery becomes blocked, typically due to vascular disease, and the blockage is circumvented by an anastomosis. This procedure typically requires the construction of an alternative path normally using an autologous vein. A high percentage of long term failures of arterial bypass grafts are observed at the downstream, or distal, end of the bypass loop. Understanding the nature of this failure has made the geometric features of the bypass junction a particular focus of three-dimensional computational modelling. An example of a high-order mesh for an anatomically realistic geometry is shown in figure 3.3. This mesh has 1624 prismatic elements and 3545 tetrahedral elements. Also shown is the distribution of wall shear stresses calculated using a fourth-order polynomial approximation.

4. Low energy preconditioning for spectral/ hp discretisations. To solve the fluid flow problem at arterial branches, as shown in figure 3.3 we have applied a three-dimensional unstructured spectral/ hp element solver [20] with a high order splitting scheme of the Navier-Stokes equations [10]. The splitting scheme requires the solution of a Pressure Poisson equation and three Helmholtz problems. The iterative inversion of the discrete elliptic problems is currently the limiting factor in computational speed.

Building on the work of Bica [1], we have developed an efficient preconditioning strategy for substructured solvers based on a transformation of the closed form expansion basis to a “low energy” basis [18]. Following this approach, the strong coupling in the matrix system between two different degrees of freedom of the original basis is significantly reduced by

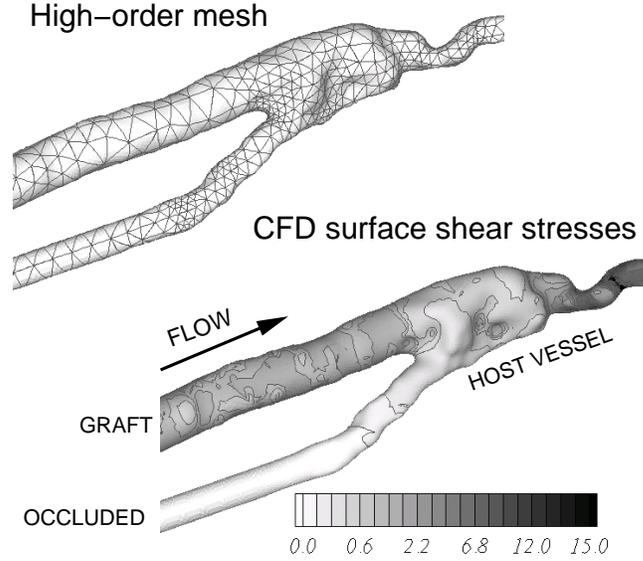


Figure 3.3: High-order mesh and distribution of surface shear stresses obtained using a fourth-order polynomial approximation in the spectral/*hp* CFD solver. The values of the shear stress have been normalized so that the inflow wall shear stress (Hagen-Poiseuille flow) is 1.

introducing a degree of orthogonality between degrees of freedom. The transformed matrix system is then amenable to block diagonal preconditioning.

The efficiency of the preconditioner is maintained by developing a new low energy basis on a symmetric reference element and ignoring, in the preconditioning step, the role of the Jacobian of the mapping from the reference to a global element. By applying an additive Schwarz block preconditioner to the low energy basis combined with a coarse space linear vertex solver we have observed up to six fold reductions in execution time for our complex geometry Navier-Stokes solver.

4.1. Overview. In this section we outline the key concepts behind the preconditioner. Full details of the formulation can be found in [18]. A representative elliptic boundary value problem is

$$\nabla^2 u(x, y, z) + \lambda u(x, y, z) = f(x, y, z) \quad (4.1)$$

which is discretised into spectral/*hp* elements by decomposing the solution domain into non-overlapping subdomains within which a polynomial expansion is applied [11]. The Galerkin formulation of equation (4.1) leads us to a matrix problem of the form

$$\mathbf{H}\hat{\mathbf{u}} = \mathbf{f}$$

where \mathbf{H} is the weak Helmholtz operator, \mathbf{f} is the inner product of the forcing term and $\hat{\mathbf{u}}$ represent the expansion coefficients of the original closed form basis. In a spectral/*hp* element approach the expansion basis is normally decomposed into interior and boundary modes where the interior modes have zero support on the element boundaries and the boundary modes make the expansion complete. Such a decomposition lends itself to substructuring [23] where we construct the boundary degrees of freedom Schur complement \mathbf{S} of \mathbf{H} . This is

essentially an orthogonalisation of the boundary degrees of freedom from the interior degrees of freedom and may also be considered as a basis transformation. This is attractive because it leaves a block diagonal matrix corresponding to the interior modes which is easily invertible. At this stage we still need to invert the positive definite Schur Complement \mathbf{S} and this can be achieved using a preconditioned conjugate gradient technique.

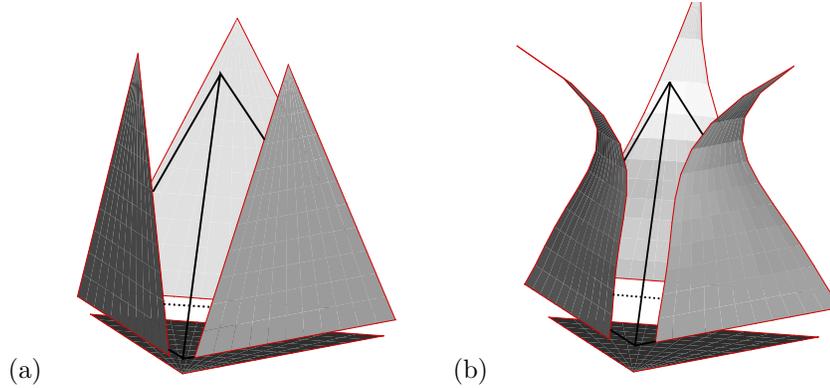


Figure 4.1: Projected mode shape of vertex 4 (a) original and (b) low energy basis. The polynomial order was $P = 5$.

The choice of the preconditioner therefore defines the efficiency of the numerical algorithm. For two-dimensional hierarchical spectral/ hp type discretisations the block diagonal preconditioner proposed by Dryja et al [6] leads to the attractive property of polylogarithmic conditioning. However for a three-dimensional hierarchical expansion this approach is not so effective [1, 18]. A significant factor is the coupling between the face expansion modes (i.e. the modes which have zero support on all edges and vertices) with the “wire-basket” space containing expansion modes which have support along the edges and at the vertices. The low energy preconditioning strategy transforms the original closed form bases to a numerically defined basis which decouples the degrees of freedom associated with each face from the vertex and edge degrees of freedom. In doing so the new basis has low energy in the sense that the inner product in the bilinear energy norm of two boundary modes is small or at least significantly reduced.

The formal details of transforming the basis are dealt with in [18]. However to illustrate the concepts we consider the shape of a vertex mode in the original and low energy basis at a polynomial order of $P = 5$ as shown in figure 4.1. The closed form original vertex mode is identical to the standard linear finite element mode and it can be appreciated that the energy associated with the inner product of this mode with any other mode in the energy norm will be reasonably high due to its high magnitude throughout the subdomain. Not too surprisingly, the shape of the low energy vertex modes decays rapidly away from the vertex where it is required to have the same magnitude as the original basis. The rapid decay is consistent with the concept of low energy in the energy norms.

From an implementation point of view the numerical orthogonalisation of each of the face boundary modes from the wire-basket modes would be as difficult as inverting the full matrix. Nevertheless the important feature of the low energy basis can be captured by defining the new basis on a rotationally symmetric region. This inherently ignores the mapping from the symmetric region to the local element within the computational domain but maintains the computational efficiency of the standard implementation.

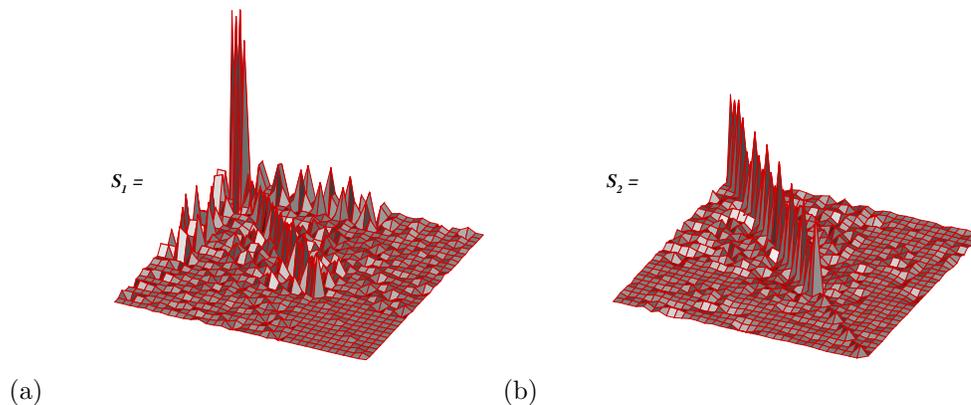


Figure 4.2: Scatter plot of Schur complement matrices of a $P = 5$ polynomial expansion: (a) Original basis (b) Low energy basis (scaled by a factor of 4).

We conclude the summary by considering a scatter plot of the Schur complement systems arising from the original and low energy basis discretisation of a Poisson equation as shown in figure 4.2. In this figure the boundary modes were ordered so that the vertex modes were followed by the edges modes which, in turn, were followed by the face boundary modes. From this plot we see that original basis has a high magnitude/energy in the vertex modes even in the off-diagonal component. There is also a significant energy between the edges and vertices. Furthermore, we see that the coupling between the face and wire-basket modes is larger than the coupling between the face modes with themselves. The low energy basis on the other hand has a more diagonally dominated structure which makes it suitable for block diagonal preconditioning.

4.2. Result. Tests of regular elements [18], where the effect of ignoring the mapping of the elements is not significant, have demonstrated that a polylogarithmic scaling of the condition number is recovered when using the low energy basis preconditioner.

In figure 4.3(a) we shown a geometrically complex computational domain of practical interest. This problems originated from the reconstruction of the downstream junction of a porcine arterial bypass [14]. The domain consists of an unstructured triangular surface discretisation from which prismatic elements are constructed by extruding the triangular surface elements in the surface normal direction. The interior region is then discretised using tetrahedral subdomains. The discretisation shown in figure 4.3 consists of 749 prismatic and 1720 tetrahedral elements.

In this domain, we have solved a Poisson equation with Dirichlet boundary conditions corresponding to the solution $u(x, y, z) = \sin x \sin y \sin z$. The condition number of the diagonal and low energy preconditioned systems are shown in figure 4.3(b). This improvement in the condition number also reflected in the speed up of the back solve of the low energy preconditioner over the diagonal preconditioner. We have observed speed-ups of approximately 6 at a polynomial order of $P = 8$ and the break-even polynomial order was approximately $P = 3$.

5. Acknowledgments. This work was partially supported by the Smiths' Charity, and the Clothworkers' Foundation. The Imperial College centres for Biomedical Visualization and Parallel Computing provided computational resources.

The authors would like to thank Prof. Kim Parker of the Department of Bioengineering

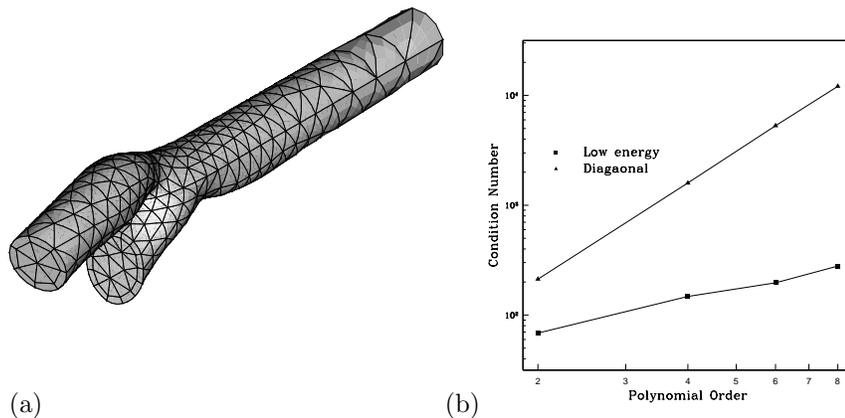


Figure 4.3: (a) Hybrid domain of a downstream arterial bypass graft. (b) Condition number as a function of polynomial order of the diagonal and low energy basis for a Poisson problem.

at Imperial College (London) and Prof. Luca Formaggia of EPFL (Lausanne) for many fruitful discussions on the modelling of wave propagation in the vascular system. The work described here has benefitted from the collaboration with Victoria Franke, Sergio Giordana and Dr. Denis Doorly of the Department of Aeronautics at Imperial College.

We also would like to acknowledge Mario Casarin for his direct contribution in the development of the low energy preconditioning methods.

REFERENCES

- [1] I. Bica. *Iterative substructuring algorithms for the p-version finite element method for elliptic problems*. PhD thesis, Courant Institute, NYU, 1997.
- [2] R. Botnar, G. Rappitsch, M. Scheidegger, D. Liepsch, K. Perktold, and P. Boesiger. Hemodynamics in the carotid artery bifurcation: A comparison between numerical simulations and in-vitro measurements. *J. of Biomech.*, 33:137–144, 2000.
- [3] B. Brook, S. Falle, and T. Pedley. Numerical solution for unsteady gravity-driven flows in collapsible tubes: evolution and roll-wave instability of a steady state. *J. Fluid Mech.*, 396:223–256, 1999.
- [4] B. Cockburn and C. Shu. TVB Runge-Kutta projection discontinuous Galerkin finite element methods for conservation laws II general framework. *Math. Comm.*, 52:411–435, 1989.
- [5] S. Dey, M. S. Shephard, and J. E. Flaherty. Geometry representation issues associated with p-version finite element computations. *Comp. Meth. Appl. Mech. Engng.*, 150:39–55, 1997.
- [6] M. Dryja, B. F. Smith, and O. B. Widlund. Schwarz analysis of iterative substructuring algorithms for elliptic problems in three dimensions. *SIAM J. Numer. Anal.*, 31(6):1662–1694, December 1994.
- [7] L. Formaggia, F. Nobile, A. Quarteroni, and A. Veneziani. Multiscale modelling of the circulatory system: a preliminary analysis. *Computing and Visualisation in Science*, 2:75–83, 1999.
- [8] P. Frey and P.-L. George. *Maillages*. Editions Hermes, 1999.

- [9] T. Hughes, C. Taylor, and C. Zarins. Finite element modeling of blood flow in arteries. *Comp. Meth. Appl. Mech. Eng.*, 158:155–196, 1998.
- [10] G. Karniadakis, M. Israeli, and S. Orszag. High-order splitting methods for incompressible Navier-Stokes equations. *J. Comp. Phys.*, 97:414, 1991.
- [11] G. E. Karniadakis and S. Sherwin. *Spectral/hp Element Methods for CFD*. Oxford University Press, 1999.
- [12] C. K. Lee. On curvature element-size control in metric surface generation. *Int. J. Numer. Meth. Engng.*, 50:787–807, 2001.
- [13] T. J. Pedley. *The fluid mechanics of large blood vessels*. Cambridge monographs on mechanics and applied mathematics. Cambridge University Press, 1980.
- [14] J. Peiró, S. Giordana, C. Griffith, and S. Sherwin. High-order algorithms for vascular flow modelling. *Int. J. Num. Meth. Fluids*, 39, 2002.
- [15] J. Peiró and A. I. Sayma. A 3-D unstructured multigrid Navier-Stokes solver. In K. W. Morton and M. J. Baines, editors, *Numerical Methods for Fluid Dynamics V*. Oxford University Press, 1995.
- [16] A. Quarteroni, M. Tuveri, and A. Veneziani. Computational vascular fluid dynamics: Problems, models and methods. *Computing and Visualisation in Science*, 2:163–197, 2000.
- [17] S. Sherwin. Dispersion analysis of the continuous and discontinuous Galerkin formulations. In *International Symposium on Discontinuous Galerkin Methods*, 1999. Newport, RI.
- [18] S. Sherwin and M. Casarin. Low-energy basis preconditioning for elliptic substructured solvers based on unstructured spectral/hp element discretization. *J. Comp. Phys.*, 171:394–417, 2001.
- [19] S. Sherwin, L. Formaggia, J. Peiró, and V. Franke. Computational modelling of 1D blood flow with variable mechanical properties and its application to the simulation of wave propagation in the human arterial system. *Int. J. Num. Meth. Fluids*, 2002. under review.
- [20] S. Sherwin and G. Karniadakis. Tetrahedral hp finite elements: Algorithms and flow simulations. *J. Comp. Phys.*, 124:14–45, 1996.
- [21] S. Sherwin, O. Shah, D. Doorly, J. Peiró, Y. Papaharilaou, N. Watkins, C. Caro, and C. Dumoulin. The influence of out-of-plane geometry on the flow within a distal end-to-side anastomosis. *ASME J. Biomech.*, 122:1–10, 2000.
- [22] S. J. Sherwin and J. Peiró. Mesh generation in curvilinear domains using high-order elements. *Int. J. Numer. Meth. Engng.*, 53:207–223, 2002.
- [23] B. Smith, P. Bjorstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, 1996.
- [24] N. Stergiopoulos and D. Young. Computer simulation of arterial flow with applications to arterial and aortic stenoses. *J. Biomech.*, 25(12):1477–1488, 1992.
- [25] J. Wang and K. Parker. Wave propagation in a model of the arterial circulation. Submitted to *J. Biomech.*, 2002.
- [26] N. Westerhof, F. Bosman, C. D. Vries, and A. Noordergraaf. Analog studies of the human systemic arterial tree. *J. Biomech.*, 2:121–143, 1969.