

High-order algorithms for vascular flow modelling

J. Peiró^{*,†}, S. Giordana, C. Griffith and S. J. Sherwin

*Biomedical Flow Group, Department of Aeronautics, Imperial College of Science, Technology and Medicine,
Prince Consort Road, London SW7 2BY, U.K.*

SUMMARY

This paper presents the development of spectral/*hp* high-order elements for vascular flows with particular attention to surface reconstruction and high-order mesh generation. Using ideas from computer visualization we apply a technique of constructing smooth implicit functions to reconstruct incomplete sectional imaging data from magnetic resonance imaging. Using this as a starting point we outline techniques to discretize and computationally solve the Newtonian flow within these geometries using high-order spectral/*hp* element methods. Finally we demonstrate the application of these ideas to anatomically correct and model distal bypass grafts. Copyright © 2002 John Wiley & Sons, Ltd.

KEY WORDS: high-order mesh generation

1. INTRODUCTION

There is growing evidence that the development of vascular diseases such as atherosclerosis and myo-intimal hyperplasia is influenced by local haemodynamics within blood vessels [1]. The need to model relevant flow features, e.g. shear stresses and residence times, has driven the use of computational methods to model vascular flows. These flows are complex as the geometry of the vessels is intricate, the flow is pulsatile and typically Reynolds number regimes are of the order where viscous and inertial effects are both significant [2].

Non-invasive magnetic resonance imaging (MRI) and CFD techniques are complementary in the study of arterial flows since MRI techniques enable measurements of the *in vivo* geometry. Further quantitative MRI velocimetry methods also permit the determination of the *in vivo* velocity field. However it is generally impractical to measure the velocity at more than a few locations *in vivo*, and the measurements are not sufficiently well resolved to determine parameters such as wall shear stresses.

Compact high-order algorithms, such as spectral element or *p*-type finite element techniques, offer the potential of high accuracy if the solution is smooth and a well behaved mapping exists between the local sub-domains and a standard region. The ability to construct suitable

*Correspondence to: J. Peiró, Biomedical Flow Group, Department of Aeronautics, Imperial College of Science, Technology and Medicine, Prince Consort Road, London SW7 2BY, U.K.

†E-mail: j.peiro@ic.ac.uk

computational meshes is currently a significant limiting factor in the development of high-order algorithms in very complex geometries.

In this paper we will address two significant issues in applying the numerical approach to vascular flow modelling. First, in Section 2, we discuss the process of reconstructing in a reproduceable fashion a surface representing an *in vivo* geometry from the incomplete information provided by MRI or ultrasound data. Second, in Section 3, we look at the application of curvature based refinement to take into account higher order information, such as surface curvature, in the mesh generation process. In Section 4, we briefly outline the spectral/*hp* element approach. Finally, in Section 5, we demonstrate the application of this technique to modelling bypass geometries.

2. IMPLICIT SURFACE RECONSTRUCTION FROM MEDICAL IMAGES

As mentioned previously, the development of non-invasive imaging techniques, such as MRI, over the last 10 years has provided a crucial contribution to simulation of the *in vivo* flow within diseased arterial vessels. However the data it provides is necessarily a series of two-dimensional sections. Further, the time available for data acquisition by non-invasive techniques such as magnetic resonance imaging (MRI) or ultrasound is restricted by patient comfort. It ranges from a few minutes to about an hour depending on the fitness of the patient and this therefore limits the number of data sections that can be acquired. Difficulties in reconstructing the arterial geometry arise mainly through the need to extrapolate information which is missing as a consequence of the limited spatial resolution of medical images. A typical example is the branch of an arterial bifurcation where the topology of the images changes between consecutive sections. The complex geometry of the apex of the bifurcation is often missing as illustrated in Figure 1.

Turk and O'Brien [3] have demonstrated that implicit functions can be used to smoothly interpolate between topologically different sections. As explained below, this interpolation method defines the surface representation as the iso-surface corresponding to a zero value of an implicit function. This function is obtained using a radial basis interpolation through a set of given points. The choice of radial basis is determined from the minimization of a Sobolev functional chosen to maintain surface smoothness (typically in curvature) between




Figure 1. Schematic of changes in topology across sections of a branching artery.




Figure 2. Image segmentation: Area selection, edge detection and spline interpolation with boundary and interior constraints shown.

the interpolation information. Based on this idea, the process of reconstructing the geometry of the vessels from MRI data proceeds as follows.

The starting point of the reconstruction process is a set of MRI slices of the geometry. These slices are represented by matrices of 512×512 pixels. Each pixel is assigned a value of intensity on a greyscale to form the image. Image segmentation techniques for edge detection using a freeware scion image package [4] are then applied to identify the vessel wall, i.e. the surface of the object to be reconstructed. As Figure 2 shows, this results in a jagged-edge representation of the boundary that has to be suitably smoothed. This is accomplished by interpolation using the least-squares technique proposed in Reference [5].

The image segmentation process is repeated for all the available MRI sections to obtain a set of B-spline curves representing cross sections of the vessel wall. These curves are then used to fit an implicit function $f(\mathbf{x})$, where $\mathbf{x} = (x_1, x_2, x_3)$ denotes the Cartesian co-ordinates of a point, and the surface is then implicitly represented as the zero iso-contour of the function $f(\mathbf{x})$, i.e. the surface will be represented by

$$f(\mathbf{x}) = f(x_1, x_2, x_3) = 0. \quad (1)$$

We continue defining the implicit function as a linear combination of radial functions associated with N different constraints at a set of discrete points \mathbf{x}_i ; $i = 1, \dots, N$, i.e.

$$f(\mathbf{x}) = \sum_{i=1}^n c_i \phi(|\mathbf{x} - \mathbf{x}_i|) \quad (2)$$

where ϕ represents a radial basis function which depends on the distance between points at a location \mathbf{x} and the constraint point \mathbf{x}_i . Following the construction in Turk and O'Brien we consider the radial basis function proposed by Duchon in Reference [6]. In this work Duchon defined a family of semi-norms

$$\|f\|_{m,s} = \left(\int_{\mathbb{R}^n} |\mathbf{x}|^{2s} |\mathcal{F}\{D^m f(\mathbf{x})\}|^2 d\mathbf{x} \right)^{1/2}$$

where \mathcal{F} denotes the Fourier transform. Minimizing these semi-norms subject to some interpolating conditions leads to simple basis functions which are independent of co-ordinate

scaling. This family is an extension of the thin-plate interpolation ($s = 0$, $n = 2$ and $m = 2$) where the functional reduces to

$$\int_{\mathbb{R}^2} \{f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2\} \, d\mathbf{x}$$

and the minimizing basis functions lead to thin-plate functions of the form

$$\phi(|\mathbf{x} - \mathbf{x}_i|) = |\mathbf{x} - \mathbf{x}_i|^2 \log(|\mathbf{x} - \mathbf{x}_i|).$$

Motivated by Reference [7] we consider the case of ‘pseudo-polynomial splines’ which are derived from minimizing the semi-norm when $s = (n - 1)/2$, $n = 3$, and $m = 2$. These give rise to radial basis functions of the form

$$\phi(|\mathbf{x} - \mathbf{x}_i|) = |\mathbf{x} - \mathbf{x}_i|^3. \quad (3)$$

Strictly speaking there should also be an order one polynomial subspace which accounts for the linear and constant portions of f but this is not necessary ‘when considering more than a dozen constraints’ [7].

Finally the problem of finding an implicit function through a set of constraints $f(\mathbf{x}_i) = h_i$; $1, \dots, N$ is reduced to that of calculating the expansion coefficients c_i in Equation (2). The evaluation of the coefficients c_i requires the solution of a positive semi-definite linear system equations of the form

$$\mathbf{A}\mathbf{c} = \mathbf{h} \quad (4)$$

which can be solved using LU decomposition. The entries of the matrix \mathbf{A} are

$$a_{ij} = |\mathbf{x}_i - \mathbf{x}_j|^3 \quad (5)$$

and h_i represents the value of the implicit function imposed at point \mathbf{x}_i . To make the spline interpolations within each MRI data section conform to the zero iso-contour of the function $f(\mathbf{x})$ we evaluate a discrete set of points, typically 30, within each spline section and set $f(\mathbf{x}_i) = h_i = 0$ at these points. To ensure that system (4) has a non-trivial solution we must also specify some non-zero constraints. These are chosen as interior points within each section as illustrated in Figure 2 and determined by choosing, for each point of the discrete set, a point displaced by a small value along the normal direction to the spline fit. An example of the interpolation of the implicit function through MRI data of a femoral–tibial bypass graft is shown in Figure 3.

A triangular representation of the surface as depicted in Figure 3(c) can also be obtained by using iso-surface extraction techniques [8]. In this work we have used the efficient implicit surface polygonization method proposed by Bloomenthal in Reference [9] which minimizes the number of function evaluations using appropriate range searching techniques.

The mesh obtained in this fashion is perfectly suitable for graphic display and even manufacturing, for instance using laser stereolithographic techniques. Figure 4 shows a resin cast of the bypass graft created from a triangulation similar to the one shown in Figure 3. Such an approach permits the geometry to be scaled to double its original dimensions which is very convenient when performing detailed experimental studies. The use of implicit functions for mesh generation presents several challenges since most of the required geometric operations become non-linear problems. The only exception is the operation of determining




Figure 3. Interpolation of a set of MRI images of a femoral bypass graft using the implicit function approach: (a) interpolated slice data (1500 points), (b) a shaded view of the iso-surface corresponding to the zero value of the implicit function and (c) surface triangulation generated by the surface polygonization technique proposed by Bloomenthal in Reference [9].




Figure 4. Stereolithography model of the bypass graft.

whether a point belongs to the surface which is trivial when such an implicit form is used. Further, the quality of the triangulation produced is not sufficiently good for it to be used as the starting point of the generation of a 3D volume. Consequently, further steps are required to achieve a surface mesh of reasonable quality for CFD.

One possible remedy is to reconstruct surface curvature information from the triangulation and perform mesh enhancement operations to improve its quality as proposed in Reference [10]. The degree of smoothness that can be achieved in the new triangulation is determined by the noise present in the initial triangulation and the inherent limitations of a reconstruction of curvature from discrete data.

We have followed an alternative strategy that is compatible with our mesh generation methodology and, moreover, permits better control of the smoothness of the interpolated surface. We adopt a boundary representation (B-Rep) of the domain. The B-Rep model provides a definition of the domain as the 3D region interior to a boundary surface composed by a set of trimmed surfaces. A trimmed surface is a region on a surface interior to an oriented




Figure 5. Reconstruction of a B-Rep definition of the graft by means of spline curves and surfaces.

set of edges on the surface. The edges and surfaces in the B-Rep model of the domain are represented using splines curves and surfaces. Even though the generation of bicubic spline patches that reproduce a zero level set of an implicit function is not a straightforward procedure, this process can be automated with little effort. In this work, an interactive tool that allows the user to select points on the smooth implicit surface is used to obtain Coons patches that are reasonably close to the original implicit surface. Splines are then drawn through these points and a patch is computed choosing four boundary curves. Patches obtained from this method are rough approximations of the implicit surface since they are only based on the information the user conveys by selecting a discrete number of points. These patches are then automatically projected at a higher resolution of discrete points onto the implicit surface using an algorithm proposed by Hartmann [11]. At the end of this procedure, the implicit surface is represented by a collection of points that are grouped in rectangular entities as shown in Figure 5.

3. MESH GENERATION OF HIGH-ORDER ELEMENTS

The extension of standard unstructured mesh generation technology to high-order algorithms is a non-trivial exercise. Complications arise due to the conflicting requirements to generate coarse meshes whilst maintaining good elemental properties in regions of high curvature as shown in Figure 6 where we illustrate the type of invalid elements which can arise.

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. The method is described in detail in Reference [12] but for the sake of completeness we include here a brief summary.

Given a contiguous surface representation in terms of bicubic spline surface and lines of intersection represented by cubic splines the surface is initially discretized into a coarse distribution of linear surface elements. The local topology of these linear element is influenced by the desire to include a boundary layer region or by taking into account surface curvature as described in Section 3.1. A high-order surface discretization is then generated by following a ‘bottom-up’ procedure where initially the triangular edges are discretized 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 bicubic 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 minimization procedure is adopted, as discussed in

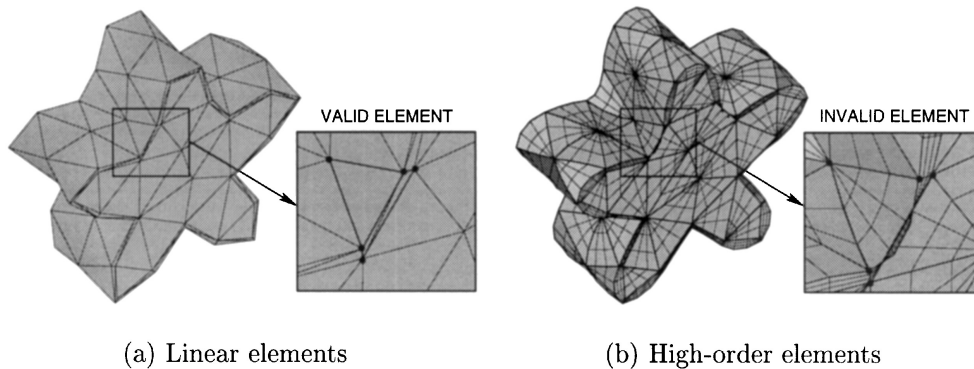


Figure 6. The subdivision of a valid mesh of linear elements (a) to generate a high-order tetrahedral mesh (b) might lead to elemental regions with singular Jacobian mappings.

Reference [12], which generates the edge and face points as geodesics of the surface with a view to minimizing the variation in the surface Jacobian.

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 Reference [13] 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.

The high-order surface definition implies that the elements adjacent to a deformed wall will also have curved internal faces which are constructed as a blend, consistent with the spectral/ hp element expansion, between the internal straight edges and the deformed surface edge (see Reference [14] 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.* [15], may be necessary in very curved domains.

3.1. Curvature based refinement for high-order elements

Curvature based refinement in which the mesh size is obtained as a function of the curvature has been proposed by several authors [10, 16] as a way to obtain an accurate piecewise linear approximation of a curved surface. In a previous paper [12] we have shown that the use of this technique 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. Here we propose an alternative method more suitable for the discretization of boundary layers.

Following the notation of Figure 7, the 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.




Figure 7. Notation for mesh spacing calculation on curves.




Figure 8. Notation for mesh spacing calculation on surfaces.

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

$$c_t \leq R \sqrt{\frac{2\delta}{R + \delta}} \tag{6}$$

The corresponding value c_q for quadrilateral elements is

$$c_q \leq \frac{2R\delta}{R + \delta} \sqrt{1 + \frac{2R}{\delta}} \tag{7}$$

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 (6) and (7) is used for the two principal directions and the corresponding mesh spacings, c_1 and c_2 in Figure 8, are calculated from the values of the principal curvatures

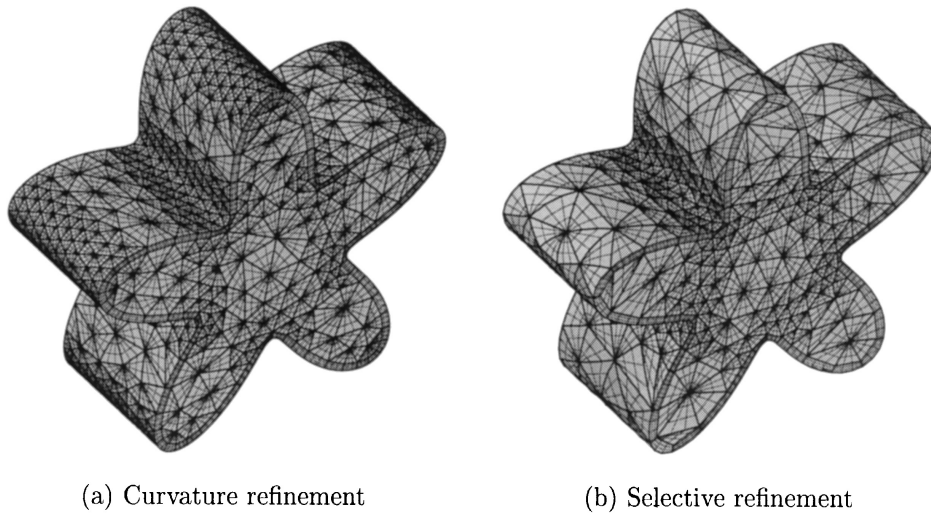


Figure 9. Curvature based mesh refinement for prismatic elements: (a) Refinement according to Equation (7), (b) Refinement is applied to concave regions only.

$k_{1,2} = 1/R_{1,2}$. Expressions for the curvature of curves and surfaces in implicit form can be found in Reference [11]. Reference [17] provides the equivalent expressions for the parametric form.

An example of a hybrid mesh generated for the geometry previously considered in Figure 6 and using criterion (7), is shown in Figure 9(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 (7) to ensure element validity is too restrictive in those regions where the domain is locally convex. For a convex region, the less restrictive criteria $\delta < R$ and $c < 2R$ suffice to guarantee element validity. This is highlighted in Figure 9(b) where the refinement criterion (7) has been selectively applied to concave regions only. The result is a valid mesh with fewer elements.

4. SPECTRAL/*hp* FLOW SIMULATION

Highly accurate algorithms, for example as Fourier methods, have proved successful in analysing fundamental fluid mechanics through the use of direct numerical simulations on simple configurations such as channel flows. However many problems of interest, such as the area of vascular flow modelling, involve complex geometries. It is the ability of finite element and finite volume methods to handle complex geometries which has led to their widespread application. Nevertheless, the low accuracy of these methods makes them relatively inefficient from the point of view of *accuracy per unit of computational power*. As illustrated in Figure 10, by applying higher order polynomial expansions within a series of elemental sub-domains, spectral/*hp* element methods [14] combine the high accuracy of Fourier techniques with the geometric flexibility of finite elements thereby providing an efficient method of achieving the highest possible accuracy for a given computational cost.




Figure 10. Illustration of the spectral/*hp* element discretization.

Figure 11 presents a comparison between standard finite element and spectral/*hp* finite element approaches for the test case of potential flow around a circular cylinder where the potential is given as $\phi(r, \theta) = U_0(r + a^2/r) \cos \theta$. The standard finite element converges as the characteristic size of each element in the mesh, denoted by h , is reduced. Therefore a series of six mesh refinements was considered as shown in Figure 11(a). In the spectral/*hp* finite element approach, a fixed resolution mesh was considered and convergence was achieved by increasing the order of the polynomial expansion within each element. The piecewise spectral approximation considered here is continuous across elements.

The computational cost per degree of freedom in the spectral/*hp* element method is proportional to $N_{\text{el}} P^{M+1}$ where N_{el} is the number of elemental subdomains and M is the spatial dimension of the approximation. Although the computational cost per piece of information is higher using the spectral/*hp* element method for smooth solutions, the rate of convergence is exponential in the polynomial order. The error per unit of computational work is accordingly less. This point is illustrated in Figure 11(b) which shows the error in the solution as a function of computational work. The dotted lines show the error per unit of computational work for the standard finite element approach (*h*-type convergence) using a first and second order polynomial approximation. For these cases we have considered computational work to scale as the number of degrees of freedom. The solid lines show the error per unit of computational work when the higher polynomials order is increased on fixed meshes A and C in Figure 11(b) (*p*-type convergence). For these cases we have scaled the computational work as $N_{\text{el}} P^3$. We see that for larger errors it is impossible to distinguish between the two approaches since the exponential convergence is only observed when the solution is captured

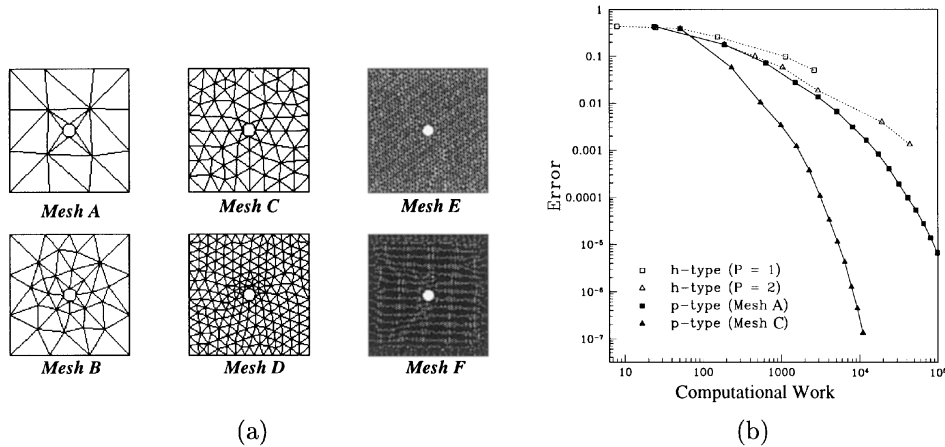


Figure 11. A model problem of potential flow around a circular cylinder was considered on a series of six meshes as shown in (a). The error in the solution as a function of idealized computational cost for the standard finite element approach (dotted lines) and the spectral/*hp* finite element approach (solid lines) is shown in (b).

spatially. However if we require a lower error then the spectral/*hp* element method can achieve this requirement at a lower computational cost as the exponential rate of convergence is realized.

These methods provide an efficient algorithm to perform accurate time dependent simulations [18, 19] and have been used to solve the unsteady Navier–Stokes equations with great success, particularly in the study of fundamental fluid dynamics through the use of direct numerical simulations [20, 21].

5. APPLICATION TO VASCULAR FLOWS

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 configurations considered in this section (Figures 12–14) a particular focus of three-dimensional computational modelling.

Vascular modelling of this type can roughly be split into two categories. The first is to consider model geometries, as shown in Figure 12, under either average or pulsatile waveforms. The second is to consider more anatomically correct situations as illustrated in Figure 14. The advantage of the first approach is that a more fundamental understanding can be obtained in order to build up a better picture of the flow dynamics in the anatomically correct example. A direct analysis of the anatomical model can be very difficult especially under physiological pulsatile conditions when the artery walls may also be distensible.




Figure 12. Advection of two lines of constant vorticity in a model 45° bypass graft.




Figure 13. High-order mesh and distribution of surface shear stresses obtained using a fourth order polynomial approximation in the *hp*/spectral CFD solver. The values of the shear stress have been normalized so that the inflow wall shear stress (Hagen–Poiseuille flow) is 1.




Figure 14. Advection of a vorticity ring through an anatomically correct porcine bypass graft.

In Figure 12 we consider a model 45° bypass junction where the inflow is a fully developed Hagen–Poiseuille flow at the location shown in Figure 12(a) at a Reynolds number of $Re = 250$. The computational domain contains 1742 tetrahedral elements and has been computed up to a polynomial order of $P = 8$ implying 287 430 local degrees of freedoms. This calculation was first reported and validated in Reference [22].

A method for calculating particle trajectories on unstructured and hybrid meshes using a high-order approximation of the velocity field has been proposed in Reference [23]. The calculation of the particle trajectory is based on a standard Runge–Kutta integration in time but using a novel hybrid approach that advances a particle in both the physical and the parametric space without the need for non-linear iterations.

Following the arguments in Reference [24] we can idealize the Hagen–Poiseuille inflow as a lumped vorticity ring of aligned normal to the axis of the inflow branch. Considering the advection of this vorticity ring through the junction as shown in Figure 12(a)–12(c) we can interpret how secondary flow motions are established as the vorticity ring transports through the junction. Secondary flow motions can be very significant in the particle transport and mixing and therefore represent an important factor in disease modelling. These secondary motions establish two vortices known as Dean vortices which are also illustrated in Figure 14 using a technique known as coherent structure identification [25, 24].

The two next CFD calculations involve anatomically realistic geometries. The first example correspond to the femoral–tibial bypass graft described in Section 2. The upstream end of the host vessel was assumed to be completely occluded. The velocity profile at the inflow section of the bypass graft was taken to be a steady Hagen–Poiseuille profile corresponding to a Reynolds number $Re = 400$. Figure 13 shows the high-order mesh, with 1624 prismatic elements and 3545 tetrahedral elements, and the distribution of wall shear stresses calculated using a fourth-order polynomial approximation. The solution shows high values of wall shear stress in the regions of the host vessel where the attachment of the graft has caused a narrowing of its cross-sectional area.

Figure 14 shows a series of time slices from a steady computation at $Re = 100$ in a reconstruction of an anatomically correct porcine coronary bypass graft. Once again an analytic inflow is imposed based on the Hagen–Poiseuille flow. The computational domain is a hybrid mesh consisting of 750 prismatic elements and 1719 tetrahedral elements and the computation was performed at a polynomial order of $P = 7$. What is interesting about this case is that whilst the overall geometry is close to having a plane of symmetry, as shown in Figure 12, the flow characteristics are distinctly different once the vorticity ring has transported through the junction. However the asymmetry of the slight narrowing (stenosis) at the junction is sufficient to completely distort the orientation of the secondary recirculation cells causing one to finally dominate. This modification is illustrative of the sensitivity of the flow dynamics to the geometry and therefore to the reconstruction process.

6. CONCLUSIONS

We have discussed the application of implicit variational functions to reconstruct *in vivo* data from arterial bypass grafts. Using high-order algorithms to construct a computational mesh and solve the Newtonian fluid mechanics we have demonstrated the application of these techniques to anatomically correct and model arterial bypass grafts. Steady flow results in these two configurations have illustrated the sensitivity of the flow to features of the geometric reconstruction.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support for this work from the Smiths' Charity, the Bupa Foundation and the Clothworkers' Foundation. The Imperial College centres for Biomedical Visualization and Parallel Computing provided computational resources. The authors would also like to acknowledge Dr. Bal Sanghera of the Department of Mechanical Engineering at Imperial College and Yannis Papaharilaou for producing the stereolithographic model of the bypass graft depicted in

Figure 4. We would also like to acknowledge Dr. Denis Doorly for his invaluable contribution to the discussion on vorticity transport in bypass junctions.

REFERENCES

1. Wootton DM, Ku D. Fluid Mechanics of vascular systems, diseases and thrombosis. *Annual Review in Biomedical Engineering* 1999; **1**:299–329.
2. Pedley TJ. The fluid mechanics of large blood vessels. Cambridge Monographs on Mechanics and Applied Mathematics. Cambridge University Press: Cambridge, 1980.
3. Turk G, O'Brien JF. Shape transformation using variational implicit surfaces. *Computer Graphics Proceedings, Annual Conference Series, SIGGRAPH99*, ACM, Los Angeles, CA, USA, August 8–13, 1999; 335–342.
4. Scion Image (based on NIH Image). www.scioncorp.com.
5. Dierckx P. *Curves and Surface Fitting Using Splines*. Oxford Science Press: Oxford, 1993.
6. Duchon J. Spline minimizing rotation-invariant semi-norms in Sobolev spaces. In *Constructive Theory of Functions on Several Variables*, Schempp W, Zeller K. (eds). Lecture Notes in Mathematics, vol. 571. Springer-Verlag: Berlin, 1977; 85–99.
7. Turk G, Dinh H, O'Brien J, Yngve G. Implicit surfaces that interpolate. In *Shape Modelling International*. Genova, Italy, May 7–11, 2001.
8. Gallagher RS. *Computer Visualization: Graphics Techniques for Scientific and Engineering Analysis*. CRC Press: Boca Raton, 1995.
9. Bloomenthal J. An implicit surface polygonizer. In *Graphics Gems*, Heckbert PS. (ed.). vol. IV. New York: Academic Press, 1994; 324–349.
10. Frey P, George P-L. *Maillages*. Editions Hermes, 1999.
11. Hartmann E. On the curvature of curves and surfaces defined by normalforms. *Computer Aided Geometric Design* 1999; **16**:355–376.
12. Sherwin SJ, Peiró J. Mesh generation in curvilinear domains using high-order elements. *International Journal for Numerical Methods in Engineering* 2001; **53**:207–224.
13. Peiró J, Sayma AI. A 3-D unstructured multigrid Navier–Stokes solver. In *Numerical Methods for Fluid Dynamics*, Morton KW, Baines MJ (eds). vol. V. Oxford University Press: Oxford, 1995.
14. Em Karniadakis G, Sherwin SJ. *Spectral/hp Element Methods for CFD*. Oxford University Press: Oxford, 1999.
15. Dey S, Shephard MS, Flaherty JE. Geometry representation issues associated with p-version finite element computations. *Computational Methods in Applied Mechanics and Engineering* 1997; **150**:39–55.
16. Lee CK. On curvature element-size control in metric surface generation. *International Journal for Numerical Methods in Engineering* 2001; **50**:787–807.
17. Do Carmo MP. *Differential Geometry of Curves and Surfaces*. Prentice-Hall: Englewood Cliffs, NJ, 1976.
18. Sherwin SJ, Em Karniadakis G. A triangular spectral element method: Applications to the Incompressible Navier–Stokes equations. *Computer Methods in Applied Mechanics and Engineering* 1995; **123**:189.
19. Sherwin SJ, Em Karniadakis G. Tetrahedral *hp* finite elements: Algorithms and flow simulations. *Journal of Computational Physics* 1996; **124**:14–45.
20. Darekar R, Sherwin SJ. Flow past a square-section cylinder with a wavy stagnation face. *Journal of Fluid Mechanics* 2001; **426**:263–295.
21. Henderson RD, Barkley D. Secondary Instability in the wake of a circular cylinder. *Physics of Fluids* 1996; **8**(6):1683.
22. Sherwin SJ, Shah O, Doorly DJ, Peiró J, Papaharilaou Y, Watkins N, Caro CG, Dumoulin CL. The influence of out-of-plane geometry on the flow within a distal end-to-side anastomosis. *ASME Journal of Biomechanics* 2000; **122**:1–10.
23. Coppola G, Sherwin SJ, Peiró J. High-order particle tracking. *Journal of Computational Physics* 2001; **172**: 353–386.
24. Doorly DJ, Sherwin SJ, Franke PT, Peiró J. Vortical flow structure identification and flow transport in arteries. *Special Issue on CFD in Computer Methods in Biomechanics and Biomedical Engineering*, 2001.
25. Joeng J, Hussain F. On the identification of a vortex. *Journal of Fluid Mechanics* 1995; **285**:69–94.