

Parallel adaptive finite element analysis of viscous flows based on a combined compressible-incompressible formulation

Dan Givoli

*Department of Aerospace Engineering,
Technion-Israel Institute of Technology, Haifa, Israel*

Joseph E. Flaherty

*Department of Computer Science, Rensselaer Polytechnic Institute,
Troy, New York, USA, and*

Mark S. Shephard

*Scientific Computation Research Center, Rensselaer Polytechnic Institute,
Troy, New York, USA*

1. Introduction

During the last few years there has been a tremendous increase in the development and application of finite element technology for the solution of fluid flow problems, particularly on adaptively defined unstructured meshes. Some of the current finite element codes can handle large-scale three-dimensional simulations of compressible or incompressible flows with very complicated geometry, general boundary conditions, and hundred of thousands to millions of degrees of freedom[1,2].

A successful general analysis code that attempts to handle realistic problems must be equipped with four main properties: modelling flexibility, numerical stability, accuracy and efficiency. Flexibility in modelling general geometries, material properties, boundary conditions etc. is an important capability which requires special non-trivial treatment, especially in three dimensions. Stability and accuracy are essential ingredients in establishing the reliability of the numerical results. High computing speed is a critical necessity for practical reasons of obtaining results within acceptable times. Various tools have been devised in recent years in the context of finite element analysis to address each of these issues. The scheme described in this paper incorporates a number of tools which cover all four aspects, and thus leads to a powerful solution technique.

This work is partly supported by AFOSR, as part of the ARPA/AFOSR Consortium for Crystal Growth Research, under grant No. F49620-95-1-0407. The authors would also like to thank Dr Carlo Bottasso, Dr Simon Brandon, Prof. Leo Franca and Prof. Alex Ostrogorsky for their very helpful remarks.

Traditionally compressible and incompressible viscous flow problems are handled by two separate codes, since these two types of problems typically involve different concerns and numerical difficulties. However, recently Hauke and Hughes[3] showed how the compressible and incompressible cases can be combined, by approaching the incompressible limit while using either primitive (pressure-velocity-temperature) variables or entropy variables (but not conservation variables). In this work, we have adopted this approach and have developed a finite element scheme capable of solving both the compressible and incompressible Navier-Stokes equations. While a direct approach to incompressible problems may be more efficient in some cases[1], the combined compressible-incompressible approach provides a unified and convenient analysis framework, and is especially attractive when buoyancy is introduced into the incompressible Navier-Stokes equations via the Boussinesq approximation.

In the present paper, we propose a new finite element framework for the large-scale analysis of compressible and incompressible viscous flows. The scheme is based on a combined compressible-incompressible Galerkin least-squares (GLS) space-time variational formulation. Unstructured three-dimensional spatial meshes of tetrahedral elements are employed, with linear spatial interpolation in all the variables. This setup enables straightforward generalization to higher-order elements. Piecewise-constant interpolation is used in time, with local time-stepping for steady flows. The scheme incorporates automatic adaptive mesh refinement, with a choice of various error indicators. It is implemented on a distributed-memory MIMD parallel computer, and includes an automatic load-balancing procedure to distribute the load uniformly among the processors. The non-symmetric system of algebraic equations is solved using preconditioned GMRES iterations.

This computational framework includes some of the features that have been used in[2-4], as shown in Table I. It combines the capability to analyse viscous

Feature	Shakib <i>et al.</i> [4]	Bottasso <i>et al.</i> [2]	Hauke and Hughes[3]	Present
Space-time GLS	X	X	X	X
Viscous flows	X		X	X
Navier-Stokes equations + Boussinesq approximation				X
Compressible-incompressible			X	X
Three dimensions		X		X
Unstructured meshes		X		X
Parallel implementation		X		X
Adaptivity		X		X
Boundary-layer error indicator				X

Table I.
Comparison of
computational features
used in three previous
works and in the present
work

flows[4] using a combined compressible-incompressible formulation[3], with parallel adaptive capabilities in three dimensions using unstructured meshes[2]. In addition, the scheme is capable of solving the Navier-Stokes equations with the Boussinesq approximation. This allows for the analysis of buoyancy-driven flows associated with divergence-free velocity fields, which are typical, for example, in certain crystal growth applications. Also, a “boundary layer error indicator” is incorporated in the adaptive procedure, which is based on the local gradient of the velocity magnitude.

We apply the new scheme to two model viscous flow problems. The first problem concerns the Mach 3 flow of air over a flat plate, and is intended to test the performance of the scheme in analysing compressible supersonic flows. The results are compared to those presented in[4], which are obtained using a serial code with a uniform structured two-dimensional mesh.

The second problem is that of a divergence-free buoyancy-driven flow in a cavity. This is a model for the steady melt flow in a liquid encapsulated Czochralski (LEC) crystal growth process[5,6]. There have been many attempts to model Czochralski (CZ) and LEC processes with various levels of complexity. See, e.g. the local (single-phase) analyses in[7-11] and the global (multi-phase) analyses in[12-15], which use various finite volume and finite element schemes, as well as the recent review[16], and references therein. Most of the computational works on CZ and LEC assume axi-symmetric solutions, but some works are based on three-dimensional analyses[8,9-11].

Most of the finite element analyses performed in crystal growth simulations are based on a standard Galerkin weighted residual formulation. For example, the scheme in[12-14] uses continuous biquadratic interpolation for velocities and temperature, and discontinuous bilinear interpolation for pressure. This combination is successful in that it satisfies the Babuška-Brezzi stability condition[17], although it is not easily generalized to higher-order interpolation. There has been an attempt to generalize this formulation to a Petrov-Galerkin framework in[18]. Here we perform a local analysis of steady-state flow in indium phosphide (InP) melt under LEC conditions, using a rather simple model, to check the effectiveness of applying the new computational framework, including the GLS formulation and the parallel adaptive capabilities, to problems of this type. The governing equations are the Navier-Stokes equations with the Boussinesq approximation.

In the next section, we describe the flow problems under consideration. In section 3 we discuss the finite element formulation, and in section 4 we describe the computational methodology and its various ingredients in some detail. In section 5, we present some numerical results; first for the Mach 3 flow of air over a flat plate, and then for buoyancy-driven flows in InP melt. We close the paper with some remarks.

2. Statement of the problem

The compressible Navier-Stokes equations can be written in the general conservation form (see, e.g.[19]),

$$\dot{\mathbf{U}} + \mathbf{F}_{i,i}^{adv} = \mathbf{F}_{i,i}^{diff} + \mathbf{S}(\mathbf{U}) . \quad (1)$$

Parallel adaptive
finite element
analysis

Here, \mathbf{U} is the vector of conservation variables, \mathbf{F}_i^{adv} and \mathbf{F}_i^{diff} are the advective and diffusive flux vectors in the i th direction, and \mathbf{S} is the source vector. A superimposed dot indicates total differentiation with respect to time, a comma denotes partial differentiation, and the summation convention with respect to repeated indices is enforced. In the three-dimensional case (1), is a set of five equations for mass, momentum (three scalar equations), and energy conservation.

Consider a transformation $\mathbf{Y} = \mathbf{Y}(\mathbf{U})$, and convert (1) to the quasi-linear form

$$\mathbf{A}_0 \dot{\mathbf{Y}} + \mathbf{A}_i \mathbf{Y}_{,i} = (\mathbf{K}_{ij} \mathbf{Y}_{,j})_{,i} + \mathbf{S}(\mathbf{Y}) , \quad (2)$$

where \mathbf{A}_0 , \mathbf{A}_i and \mathbf{K}_{ij} , $i, j = 1, 2, 3$, are \mathbf{Y} -dependent matrices. Assume that the source term \mathbf{S} can be written in the form,

$$\mathbf{S} = \mathbf{S}_0 - \mathbf{C} \mathbf{Y} , \quad (3)$$

where \mathbf{S}_0 is a constant vector, and \mathbf{C} is a \mathbf{Y} -dependent matrix. For future reference, write (2) and (3) in the form

$$\mathcal{L} \mathbf{Y} = \mathbf{S}_0 , \quad (4)$$

$$\mathcal{L} = \mathbf{A}_0 \frac{\partial}{\partial t} + \mathbf{A}_i \frac{\partial}{\partial x_i} - \frac{\partial}{\partial x_i} \left(\mathbf{K}_{ij} \frac{\partial}{\partial x_j} \right) + \mathbf{C} . \quad (5)$$

Three possible choices for \mathbf{Y} are: entropy variables[19], for which the system (2) becomes symmetric; pressure-primitive variables; and density-primitive variables[3, 20]. Hauke and Hughes[3,20] recently showed that the *incompressible limit* of (2) is well defined, and a well-behaved variational formulation of the problem can be derived when \mathbf{Y} is chosen as either the entropy variables or the pressure-primitive variables. Thus, with either of these two variable sets, the compressible and incompressible cases may be combined into a unified framework.

For fully incompressible flows, this approach seems to be less efficient than the direct one, because there is complete decoupling of the energy equation from the momentum and continuity equations. Moreover, in many cases only the velocity and pressure fields are of interest and not the temperatures. Under such circumstances, the full solution of (2) for incompressible flows entails superfluous computational effort. Indeed, traditionally compressible and incompressible viscous flow problems are handled by two separate codes[1]. However, the combined compressible-incompressible approach becomes attractive when buoyancy is considered, and introduced into the incompressible Navier-Stokes equations via the Boussinesq approximation[21]. In this case, \mathbf{S} in (2) is

HF
7,8

temperature-dependent, and all the equations are coupled, as with compressible flows. We adopt the combined compressible-incompressible approach here, using entropy variables for \mathbf{Y} .

The time-dependent Navier-Stokes equations with the Boussinesq approximation can be written in the explicit form

884

$$\rho_0 (\dot{\mathbf{v}} + \mathbf{v} \cdot \nabla \mathbf{v}) = \nabla \cdot \boldsymbol{\sigma} + \rho_0 \beta g (T - T_0) \mathbf{e}_z , \quad (6)$$

$$\nabla \cdot \mathbf{v} = 0 , \quad (7)$$

$$\rho_0 c_p (\dot{T} + \mathbf{v} \cdot \nabla T) - \nabla \cdot (\kappa \nabla T) = 0 . \quad (8)$$

Here, \mathbf{v} is the velocity vector, T is the temperature, $\boldsymbol{\sigma}$ is the stress tensor, ρ_0 is a constant reference density of the fluid, T_0 is a reference temperature (corresponding to the density ρ_0), β is the thermal expansivity of the fluid, g is the gravitational acceleration, \mathbf{e}_z is a unit vector in the z direction, c_p is the heat capacity, and κ is the heat conductivity of the fluid. Equations (6), (7) and (8) are, respectively, the momentum, continuity and energy equations. The last term on the right hand side of (6) is the buoyancy term, which is responsible for the coupling between the temperature and velocity fields. Gravity is assumed to act in the $-z$ direction. The stress-pressure-velocity relations are given, for a Newtonian fluid, by

$$\boldsymbol{\sigma} = -p\mathbf{I} + \mu(\nabla \mathbf{v} + \nabla \mathbf{v}^T) , \quad (9)$$

where \mathbf{I} is the identity tensor, and μ is the fluid viscosity.

Equations (6)-(9) are recast in the vector form (2) or (4), to fit the combined compressible-incompressible formulation. These equations are accompanied by appropriate boundary and initial conditions, to complete the statement of the problem. In the steady-state case, the time-derivative terms in (6), (8), (2) and (ref L) are dropped.

3. Variational formulation

Our starting point is the transformed time-dependent Navier-Stokes equations (2). These equations will be considered even when steady solutions are sought; in this case we shall employ the commonly-used relaxation procedure, in which we discretize the time-dependent problem (2) in space and time, and march rapidly in time to reach a steady state (see section 4 – “Interpolation in time and space”).

We define a weak form of (2) based on the time-discontinuous Galerkin least squares (GLS) method. This formulation was originally developed for two-dimensional compressible flows[4], and was also used in other configurations by [2] and [3]; see Table I. The time-discontinuous GLS method is well suited for incorporation in an automated parallel adaptive environment. The method has a firm mathematical foundation, and its stability and accuracy properties have been rigorously established (see[4] and references therein). Moreover, it has the

ability to handle moving boundary problems naturally by means of space-time deforming meshes[22]. First we introduce some notation.

Given a spatial domain Ω with boundary Γ , consider a space-time domain $\Omega \times I$, where the time interval I is divided into intervals (time steps), $I_n = (t_n, t_{n+1})$. Let $Q_n = \Omega \times I_n$ (a "space-time slab"), let $P_n = \Gamma \times I_n$ and let ν be the unit outward normal on P_n . Also, let \mathbf{W}_n^- and \mathbf{W}_n^+ denote the values of the time-discontinuous quantity \mathbf{W} as $t \rightarrow t_n$ from below and above respectively. The slab Q_n is decomposed into $(N_{el})_n$ space-time elements, denoted Q_n^e . Finally, let $\mathfrak{S}(Q_n)$ and $\mathfrak{S}_0(Q_n)$ denote the appropriate trial and test spaces respectively, defined on the slab Q_n .

With this notation, the variational formulation of (2) is stated as follows: within each slab Q_n find $\mathbf{Y} \in \mathfrak{S}(Q_n)$ such that for all $\mathbf{W} \in \mathfrak{S}_0(Q_n)$,

$$\begin{aligned} & \int_{Q_n} \left(-\dot{\mathbf{W}} \cdot \mathbf{U}(\mathbf{Y}) - \mathbf{W}_{,i} \cdot \mathbf{F}_i^{\text{adv}}(\mathbf{Y}) + \mathbf{W}_{,i} \cdot \mathbf{K}_{ij} \mathbf{Y}_{,j} - \mathbf{W} \cdot \mathbf{S} \right) dQ \\ & + \int_{\Omega} \left(\mathbf{W}_{n+1}^- \cdot \mathbf{U}(\mathbf{Y}_{n+1}^-) - \mathbf{W}_n^+ \cdot \mathbf{U}(\mathbf{Y}_n^+) \right) d\Omega \\ & + \sum_{e=1}^{(N_{el})_n} \int_{Q_n^e} (\mathcal{L}\mathbf{W}) \cdot \boldsymbol{\tau}(\mathcal{L}\mathbf{Y} - \mathbf{S}_0) dQ \\ & + \sum_{e=1}^{(N_{el})_n} \int_{Q_n^e} \delta g^{ij} \mathbf{W}_{,i} \cdot \mathbf{A}_0 \mathbf{Y}_{,j} dQ \\ & = \int_{P_n} \mathbf{W} \cdot \left(-\mathbf{F}_i^{\text{adv}}(\mathbf{Y}) + \mathbf{F}_i^{\text{diff}}(\mathbf{Y}) \right) \nu_i dP \quad (10) \end{aligned}$$

The first and last integrals in (10) are the standard Galerkin terms. The second integral weakly enforces continuity of the solution in time from one slab to the next. Thus, the first, second and last integrals together constitute the discontinuous-Galerkin method[23]. The use of finite element shape functions which are temporally discontinuous enables simple stepping in time, similar to the time marching procedure performed in finite difference schemes. The third and fourth integrals are stabilizing terms; the third is a least squares term, and the fourth is a discontinuity-capturing term. The variational form (10) constitutes the time-discontinuous GLS method (with discontinuity capturing).

The least squares (LS) term is the main stabilization mechanism. Numerical stability is always an issue in finite element formulations of the Navier-Stokes equations. A necessary condition for stability is the Babuška-Brezzi condition, and it is well known that it is not easily satisfied with standard Galerkin methods[24]. For example, equal-order interpolation in space for all the variables leads to a highly unstable numerical scheme, which typically yields non-convergent solutions even for relatively simple problems. The LS term enhances the stability of the method, and allows the use of equal-order interpolation in space, while maintaining the optimal order of accuracy.

Equal-order interpolation has the advantage that it easily accommodates high-order elements and ρ -adaptivity. Stabilization is also needed to eliminate global pollution (spurious oscillations) owing to unresolved boundary layers, large gradients, discontinuities, and advection dominance[25].

The matrix τ appearing in the LS term in (10) is designed to balance stability and accuracy in both the diffusive and advective limits. Definitions of τ for the compressible and incompressible Navier-Stokes equations can be found in[4] and[26] respectively. A way to combine both τ 's is proposed in[20].

The LS term largely eliminates global pollution effects; however, Gibbs' phenomena may still exist in the immediate vicinity of strong discontinuities such as shock waves. The fourth integral in (10) is a non-linear discontinuity capturing (DC) operator, which controls these local oscillations. For more details, as well as the definitions of δ and g^j , see[4, 20].

As can be seen from (10), the LS and DC operators are applied at the element level. The method is consistent in the sense that (10) is satisfied by the exact solution of the problem. This is ensured since both stabilization terms depend on the residual $\mathbf{L}\mathbf{Y} - \mathbf{S}_0$ (see (4)) in the element, and vanish with this residual. Note also that the weighting $\mathbf{L}\mathbf{W}$ in the LS term includes the \mathbf{Y} -dependent part of the source \mathbf{S} , but does not include its constant part \mathbf{S}_0 (see (3)).

Essential and natural boundary conditions can be applied in conjunction with (10). Essential boundary conditions are imposed directly, after expressing them in terms of the \mathbf{Y} variables. Natural boundary conditions are imposed by substituting the prescribed value for the corresponding flux in the (last) boundary integral of (10). If no boundary conditions are imposed on some part of the boundary, homogeneous natural boundary conditions result using the current values of \mathbf{Y} .

4. Computational scheme

Figure 1 summarizes the overall computational scheme. First, an initial mesh is generated. Then, the domain is decomposed into partition domains, so as to balance the initial load among the parallel processors. The ensuing main analysis loop includes five phases:

- (1) the "form phase", in which the element matrices and vectors and the global equations are formed;
- (2) the "solve phase", where the global algebraic equations are solved;
- (3) the error estimation phase, where some specified measure of the local error (the error indicator) is calculated for all the elements;
- (4) the mesh adaptation phase, in which the mesh is locally refined and/or coarsened based on the error indicator values; and
- (5) the load balancing phase, where elements are migrated between processors in order to redistribute the load in a uniform manner among the processors.

The procedures used for the mesh generation, mesh adaptation, domain partitioning, load balancing and element array formation are all linked to the parallel mesh database and attribute manager, which store, handle and supply




Figure 1.
Flowchart of the overall
solution procedure

information on mesh entities and physical data. Now we shall describe in some detail the various ingredients of this computational scheme.

Initial mesh generation and the attribute manager

Spatial meshes are generated via the finite octree technique[27], which constructs general unstructured meshes of tetrahedral elements. The first step in meshing a model region is to develop an octree; this is a multi-level tree whose basic unit on each level is a mesh cell divisible to eight subcells (octants) on a lower level. This octree is constructed so that it reflects the mesh information and is consistent with the triangulation on the boundary of the

HF
7,8

888

model region. A one-level difference of octants sharing one or more edges is enforced during this process to control smoothness of the mesh gradations. Once the octree is generated, the octants are classified as interior or boundary octants. Some interior octants are reclassified as boundary octants if they are sufficiently close to mesh entities on the boundary. The next step consists of meshing the interior octants. The last step in the mesh generation is to connect the boundary triangulation to the interior octants, using an operation called “face removal”. The finite octree mesh generation procedure has been proved to be an extremely strong tool for general three-dimensional geometries.

The attribute manager stores and handles the physical attribute information required to support the analysis, e.g. boundary conditions, initial conditions and material properties. This information is tied to the geometric model definition, rather than to the discrete model, and is defined in a general hierarchical form[28]. This is in contrast with the common procedure of defining the physical attribute information directly in terms of the discrete model. The attribute manager is especially effective in an adaptive environment.

Interpolation in time and space

The choice of finite element shape functions in space and time, in conjunction with the GLS time-space variational formulation described in section 3, determines the definition of the element arrays, and thus completes the information needed for the “form phase”.

In the present scheme, piecewise-constant interpolation is used in time. This choice is permitted in the time-discontinuous GLS method, and is very convenient and efficient for steady-state analysis. It is used for transient analysis. A single Newton iteration is performed at each time step.

In transient analysis, a time-step size is specified which may vary in time, but is uniform in space. In steady-state analysis, the time-step need not be uniform throughout the mesh, but is determined locally, e.g. to achieve a specified element Courant number to reach steady state rapidly. The Courant number measures the estimated number of elements over which the information propagates in a single time step. Using this procedure, the flow information propagates at a nearly uniform rate (a fixed number of elements per time step) throughout the spatial domain.

In space, continuous piecewise-linear spatial interpolation is used in all the variables. This convenient choice, which is never possible in standard mixed Galerkin formulations, is made possible by the LS stabilization. In addition to programming convenience, a major advantage of an equal-order interpolation is that it is amenable to easy generalization to higher-order elements. As a consequence, the present scheme can be extended, in a straightforward manner, to general *hp*-adaptive strategies.

Parallel procedures

The finite element scheme used here is designed to work in an automatic parallel adaptive environment. One important aspect of parallelization is the partitioning of data among the processors, given the irregular communication patterns that

characterize unstructured meshes. Domain decomposition is used for partitioning the computational domain into subdomains, each being assigned to a processing node. The domains exchange data with each other through their boundaries.

The partitioning of the initial mesh is done using a tree-based moment-of-inertia recursive bisection (RB) algorithm. It is performed by loading the mesh onto one processor and recursively bisecting it in a direction normal to the region's principal axis of inertia. Successive halves are sent to other processors to be further split in parallel. The initial mesh can also be loaded onto multiple processors in any order and partitioned with a parallel bisection procedure[29].

The message passing interface (MPI) is used to implement all parts of the software. Parallel implementation of the form phase is straightforward, with no communication among the processors. The solve phase is based on the use of the preconditioned GMRES algorithm (see section 4 – “Equation solving, for the details”).

The parallel mesh database (see Figure 1) stores, handles and supplies information on mesh entities. It has three main roles: to provide a common interface and a single library for all the mesh related application, namely mesh generation, mesh refinement/coarsening and finite element analysis; to provide information on adjacency relations among shared mesh entities on different processors; to support a mesh migration algorithm which facilitates arbitrary mobility of mesh entities on processors. The mesh database defines the mesh in terms of the topological entities of the mesh (vertices, edges, faces and regions) with pointers between them.

The data structures used in the scheme provide fast query of domain boundary information. Such information includes adjacencies for entities located on more than one partition, the number and list of adjacent processors given an entity type adjacency, and so on. In addition to these queries, fast update procedures are used for the refinement/coarsening and element migration/load balancing components of the solver. The fast query and update are achieved via a topological-entity hierarchy data structure, which provides a two-way link between the mesh entities of consecutive order, i.e. regions \leftrightarrow faces, faces \leftrightarrow edges and edges \leftrightarrow vertices. Using this hierarchy, any entity relationship can be derived by local traversals[30]. The entities on domain boundaries are augmented with links, which indicate the processor holding the entity and point to the location of the corresponding entity on the neighbouring processor.

Scalable parallel procedures are used to migrate elements between processors for the purpose of redistributing the mesh in order to achieve *load balance*[31, 32]. The migration routines are implemented in three stages: the element mesh and its attribute data is packed into messages and sent; packed elements are received and unpacked; and the inter-processor links are updated. Procedures allow each processor to send and receive multiple migrations of elements. An element-based iterative dynamic load balancing scheme is used to determine those elements to be migrated from heavily loaded to less loaded processors. The hierarchic load request is represented by a tree, on which a scan

operation is performed. The complexity of this scan is $O(\log T)$, where T is the number of nodes in the tree balance[31].

These parallel procedures result in a scalable scheme (see section 4 – “Overall performance”) that runs on a distributed-memory MIMD parallel computer. The simulations described in section 5 are performed on 36- and 512-processor IBM SP-2 computers at Rensselaer Polytechnic Institute and Cornell University respectively.

Error estimation

Local error control of non-linear hyperbolic and parabolic systems based on rigorous mathematical foundations, i.e. on local a posteriori error estimates, is very difficult in general (as opposed to the linear elliptic case, for example). Traditionally, error indicators are used in place of error estimators when the former are not available. Although error indicators do not necessarily measure the true discretization error, they do indicate, if chosen carefully, local regions of the computational domain associated with high errors, and thus can be very useful in an adaptive environment.

Two types of error indicators are currently implemented in the scheme described here; both can be applied to one or more of the physical variables (typically density, pressure or temperature). The first error indicator is the magnitude of the gradient of the variable chosen. This indicator is not based on solid mathematical ground, but it is useful in practice when relatively small local regions of large gradients of the solution are present. Error indicators based on the magnitude of density or pressure gradients are usually effective in tracking shocks in compressible supersonic flows. A nodal error indicator which is appropriate when thin boundary layers are present is

$$e_I = A_I (|\nabla v|) \quad , \quad v \equiv |v| \tag{11}$$

where $A_I(\cdot)$ is the average value of a discontinuous quantity at node I. This error indicator measures the local change in the magnitude of the nodal velocity. The magnitude of the velocity v is assumed to vary linearly on each element. In the simulations of section 5 – “LEC melt flows”, we use a combination of the “boundary layer error indicator” (11) and one which measures temperature gradient.

The second error indicator which is implemented in the scheme is related to the magnitude of the second derivatives of the variable chosen[33]. This indicator is appropriate for linear elements, since it attempts to measure the error by estimating the leading (quadratic) term truncated. It has the form

$$e_I = \frac{h^2 H_I(\phi)}{h A_I(|\nabla \phi|) + \epsilon M_I(|\phi|)} \tag{12}$$

where h is the mesh size parameter, ϕ is the solution variable being monitored, $H_I(\phi)$ is a measure of the nodal magnitude of the second derivatives of ϕ , $M_I(\cdot)$ is the weighted average of a quantity over all the nodes which surround node I , ϵ

is a tuning parameter, and $A_i(\cdot)$ is defined as before. The second derivative measure H_i is computed using a variational recovery technique[33].

For each edge, the value of the error indicator is computed by averaging the associated two nodal values of the indicator. These edge values are then used to control mesh adaptation. Maximum and minimum thresholds, M and m , are supplied for the error indicators, so that the edge is refined if the error associated with at least one of the selected physical variables is higher than M , and is collapsed if the error is smaller than m for all selected variables. In some cases, especially if sudden changes in the character of the solution are expected, m is specified as zero, to turn off mesh coarsening. The threshold M is chosen so that only about 5-15 per cent of the elements in the mesh are refined, to prevent an excessive computational effort in the subsequent analysis cycle. Parallel implementation of the error indicator calculation is straightforward.

The adaptive process starts from an initial mesh (see Figure 1), which is generated using the finite octree technique. This mesh should be fine enough to capture the main features of the solution, but ideally not much finer than this. Starting from an initial mesh which is too crude is potentially dangerous, because it may lead to situations where a certain feature of the exact solution (e.g. a weak shock or local oscillations), which occurs at a scale much smaller than that of the initial mesh size parameter, may remain undetected by the error estimation phase.

In transient analysis, an adaptive step, which includes error estimation and mesh adaptation, is invoked every specified number of time steps. In a steady-state analysis, the procedure adopted in the present work is as follows: first, an initial mesh is generated. Then, local time-stepping is performed with this mesh for a sufficient number of time steps, until the residual norm is reduced below a certain specified level, indicating that a steady state is reached. Then, one step of mesh adaptation is performed, using one of the error indicators mentioned above. After this step, another cycle of local time stepping is performed, until steady state is again reached. This procedure is repeated for a desired number of cycles. A typical behaviour of the residual norm is: possible initial oscillations, then reducing monotonically until the first mesh adaptation, then slightly increasing, and again reducing monotonically until the next adaptation step, and so on. Cases where a more oscillatory behaviour is observed typically indicate that a steady state was not reached successfully, possibly because one does not exist or is unstable for the specified parameters.

Mesh adaptation

The mesh level adaptive scheme combines local refinement, coarsening and triangulation optimization using local retriangulations. The refinement step makes use of division patterns (templates). All possible patterns are considered and implemented to allow for speed and avoid possible over-refinement. The coarsening step is based on the edge collapsing technique. This procedure does not require storage of any history information and therefore does not depend on the refinement procedure. Triangulation optimization is necessary to prevent mesh quality degradation. It is based on the iterative local retriangulation of

well defined polyhedra, using the two dual techniques of edge removal and multi-face removal[29].

Since refinement uses templates, its parallel implementation presents no difficulties. First, mesh faces on the partition boundary are triangulated. Face level interprocessor links are set up for the child faces of the mesh faces on the partition boundary. Then, mesh regions are triangulated without communication involved. Any mesh edge carrying minimal inter-processor links transfers link information to its two child edges.

The efficient parallelization of local adaptive mesh refinement and coarsening is discussed in [29]. Efficient parallel retriangulation of polyhedra is performed in three steps:

- (1) retriangulate polyhedra which are fully interior to the partition;
- (2) shift the partition boundary using the element migration techniques;
- (3) retriangulate polyhedra that are now fully accessible owing to the shift.

Multiple processors can request the same off-processor region, and in such cases the processor with the lowest identification number is the one which is granted priority. This makes triangulation optimization an iterative process and, naturally, shifting the partition boundary constantly in one direction would quickly create a load imbalance. Therefore, a load balancing step (see section 4 – “Parallel procedures”) is applied after each adaptive step. A way of anticipating and avoiding these imbalances is described in[34].

Equation solving

The non-linear system of algebraic equations obtained from the finite element discretization in space and time is linearized using Newton iterations. The resulting non-symmetric linear system is solved iteratively using preconditioned GMRES. Preconditioning is performed by means of a nodal block-diagonal scaling transformation. Given the non-symmetric system

$$\mathbf{T}\mathbf{p} = \mathbf{F} , \tag{13}$$

the GMRES algorithm[35] starts from an initial guess \mathbf{p}_0 , and attempts to find an approximate solution $\mathbf{p}_0 + \mathbf{z}$ to (13), where \mathbf{z} is a vector in the Krylov space $\mathbf{K} = (\mathbf{r}_0, \mathbf{T}\mathbf{r}_0, \dots, \mathbf{T}^{k-1}\mathbf{r}_0)$ and $\mathbf{r}_0 = \mathbf{F} - \mathbf{T}\mathbf{p}_0$. The vector \mathbf{z} is found by solving the minimization problem

$$\min_{\mathbf{z} \in \mathbf{K}} \|\mathbf{F} - \mathbf{T}(\mathbf{p}_0 + \mathbf{z})\| , \tag{14}$$

using the QR algorithm. An orthogonal basis for \mathbf{K} is obtained by a Gram-Schmidt procedure. This procedure involves matrix-vector multiplication and vector dot products, which constitute the computationally-intensive part of the algorithm.

The matrix-vector multiplications are implemented in parallel and necessitate the exchange of data through the interprocessor boundaries. In order to overlap communication and computation for efficiency reasons, these operations are realized following a four-step procedure on each processor:

-
- (1) sending information on the inter-processor boundaries to each neighbouring processor;
 - (2) performing computations involving only data associated with nodes that lie within the internal volume of the partition;
 - (3) receiving information on the inter-processor boundaries from all the neighbouring processors; and
 - (4) performing computations involving only data associated with nodes that lie on the inter-processor boundaries.

For the implementation of the vector dot product operations, nodes that lie on the inter-processor boundaries are randomly split, so that two partitions that share an internal boundary are assigned only a set of the nodes of that boundary. Then, each processor performs the dot product involving nodes contained in its internal volume and its subset of nodes on the partition boundaries. The sum of all the processor results yields the global dot product.

The minimization problem (14) is written in terms of an upper Hessenberg matrix, whose dimension is the same as the dimension of the space \mathbf{K} , and whose entries are the results of the dot products performed during the orthogonalization procedure. At the end of the Gram-Schmidt procedure, each processor has complete knowledge of the upper Hessenberg matrix and is therefore able to find a solution of (14) without inter-processor communication. Similarly, once convergence is achieved in the GMRES solver, the Newton update (i.e. addition of the incremental solution to the solution from the previous Newton iteration) is performed by each processor independently with no communication.

Overall performance

With the features described in the previous sections, the proposed scheme contains the four desired properties mentioned in section 1: flexibility (mesh and physical attribute handling), stability (LS and DC operators), accuracy (adaptivity, ability to handle high-order elements), and speed (parallel environment, local time-stepping).

The parallel adaptive procedure enables the solution of large problems, with CPU times which are much smaller, for the same level of accuracy, than those required in a serial code using a uniform mesh. A single steady-state simulation of those reported in section 5 – “LEC melt flows”, with an initial mesh of about 140,000 elements, with two adaptive steps leading to a final mesh of about 400,000 elements, on 16 processors, took about 11 hours of computing, including the initial mesh generation and partitioning (also see section 5 – “Mach 3 flow over a plate”).

The parallel environment is associated with an additional memory requirement compared to a serial code, but it is usually marginal. This additional storage is handled by the parallel mesh database, which defines the mesh in terms of its topological entities with pointers between them. The mesh on each processor is stored, and in addition, special treatment is given to those mesh

HFF
7,8

entities which are on the boundary of the partition. For them, an entity neighbour is identified by indicating the processor it is on and its local memory pointer on that processor. There is also a need for a linked list of the entities on the interprocessor boundaries to support the gather/scatter operations performed by the GMRES solver. This comprises the added memory needed in the parallel scheme, and it is typically small compared to the total memory volume needed. The size of the added memory is proportional to the number of mesh entities on the interprocessor boundaries.

894

The scheme has a high degree of scalability. This is demonstrated by Figure 2, which shows the speed-up gained relative to a serial run as a function of the number of processors used.




Figure 2.
Scalability of the
numerical scheme

These results correspond to one of the simulations reported in section 5 – “LEC melt flows”. When the number of processors increases from one to two, a speed-up of about two is obtained. When 16 processors are used, a speed-up of 14 is obtained. Thus, the slope of the scalability curve is about 0.875 (vs. the ideal unit slope). The 12 per cent reduction in the slope is owing to communication overhead. Communication cost is mainly owing to the load balancing procedure, which is dominated by interprocessor communications. However, this cost is maintained at a reasonable level (as Figure 2 shows) owing to the special parallel procedures mentioned in section 4 – “Parallel procedures”. The communication cost is especially small for relatively large problems, where a large number of finite elements is used per processor.

5. Numerical examples and results

Mach 3 flow over a plate

To demonstrate the capability of the scheme to handle compressible viscous flows, we consider steady flow of air over a flat plate, with inflow conditions

corresponding to $M = 3$ and $Re = 1,000$ (see [4, 20]). From the leading edge of the plate, a thick boundary layer and a shock develop. The solution is two-dimensional, but here it is obtained using the three-dimensional computational domain $-0.2 \leq x \leq 1.2$, $0 \leq y \leq 0.8$, $0 \leq z \leq 0.1$, where the "thickness" 0.1 in the z direction is arbitrary, and is chosen for numerical convenience. The boundary of this domain is divided into seven faces, with the following boundary conditions:


$x = -0.2$ (inflow):	$\rho = 1, v_1 = 1, v_2 = v_3 = 0, T = 2.769 \cdot 10^{-4}$.
$y = 0.8$ (inflow):	same as above.
$y = 0, x \leq 0$ (plate surface):	no-slip, $T = 7.754 \cdot 10^{-4}$.
$y = 0, x < 0$ (symm.):	slip condition, zero normal heat flux.
$z = 0$ (symm.):	same as above.
$z = 0.1$ (symm.):	same as above.
$x = 1.2$ (outflow):	no boundary conditions.

The gas is assumed to be ideal, with Sutherland's viscosity law, $\mu = 0.0906 T^{1.5} / (T + 0.0001406)$, and $Pr = 0.72$. The initial mesh used is shown in Figure 3(a). It consists of 6,486 tetrahedral elements and 1,770 nodal points (with 5 degrees of freedom per node). The figure shows the trace of this mesh on the xy plane. Two cycles of adaptive mesh refinements were applied (see section 4 – "Error estimation"). The magnitude of the density gradient served as an error indicator. The final mesh thus obtained is shown in Figure 3(b). Mesh refinement at and near the boundary layer and the shock is clear.


Figures 4(a)-4(d) show the contours of the finite element solution for the density, Mach number, temperature and pressure respectively. The results agree well with those obtained in [4] and [20]. The latter results were obtained using a serial, two-dimensional code, with a uniform structured mesh which is as fine as the finest level in the mesh of Figure 3(b). Thus, we obtain a similar solution to that of [4, 20] while achieving significant saving in the number of degrees of freedom. It is apparent from Figures 3(b) and 4(a)-4(d) that the strong features in the solution, namely the boundary layer and the shock, are captured very effectively by the adaptive scheme.

The final mesh is not uniform in the "thickness direction" z of the plate; elements in the fine regions are small in all three dimensions, since the tetrahedral elements generated during the adaptive step are well-proportioned. The coarse regions of the mesh consist of two elements in the z direction, whereas in the fine regions there are up to eight elements in this direction.

This analysis was performed with eight and 16 processors of an IBM SP/2 computer. Figures 5(a) and 5(b) show, respectively, the division of the initial mesh and of the final mesh, in the eight-processor case. As more elements are concentrated in the regions of the shock and the boundary layer, the processor subdivision is changed by the dynamic load balancing procedure discussed in section 4 – "Parallel procedures".



(a)



(b)

Figure 3.
Flow over plate: finite
element meshes: (a)
initial mesh; (b) final
mesh, after two steps of
adaptive refinement

LEC melt flows

The finite element scheme is now applied to a problem of divergence-free buoyancy-driven flow in a cavity. This is a simplified single-phase model for the steady melt flow in a LEC crystal growth process (see section 1). The governing equations are the Navier-Stokes equations with the Boussinesq approximation, (6)-(9). We use this model problem to check the effectiveness of applying the new computational framework to problems of this type.

We consider buoyancy-driven flow in InP melt contained in a small cylindrical crucible with a curved bottom, with radius r_c and height H . The upper surface of the melt (including the melt-crystal and melt-encapsulant interfaces) is assumed to be flat. We introduce a cylindrical system of co-ordinates (r, θ, z) , with the z axis coinciding with the crucible axis and pointing upwards. We denote the part of the boundary defined by the crucible walls and bottom by Γ_c






Figure 4.
Flow over plate: finite
element results. Contour
lines are shown for (a)
density; (b) Mach
number; (c)
temperature;
(d) pressure



(a)



(b)

Figure 5.
Flow over plate:
partitioning of the mesh
to eight parallel
processors: (a) initial
mesh partitioning; (b)
final mesh partitioning

The melt-solid interface is defined by the circular domain $z = H, 0 \leq r \leq r_s$, and is denoted Γ_s . The melt-encapsulant interface is defined by the annular domain $z = H, r_s < r \leq r_c$ and is denoted Γ_f .

The boundary conditions considered are as follows. A no-slip condition is applied on the whole boundary. On Γ_c a given constant temperature T_c is prescribed. On Γ_s the temperature is equal to the melting-point temperature T_{mp} . On Γ_f zero normal heat flux is assumed. The no-slip condition at the melt-encapsulant interface is reasonable since the encapsulant is much more viscous than the melt. The values of the parameters are: $r_c = 3.4\text{cm}$, $r_s = r_c/2$, $H = r_c$, $\rho_0 = 5.05 \text{ g/cm}^3$, $\mu = 0.0081\text{g/cm-s}$, $\kappa = 0.23\text{W/cm-K}$, $c_p = 0.42\text{J/g-K}$, $\beta = 4.44 \cdot 10^{-4} \text{ K}^{-1}$, $T_0 = T_{mp} = 1,335\text{K}$. For later reference, we define the Grashof number, which is the ratio of buoyancy to viscous forces,

$$\text{Gr} = \rho_0^2 \beta g r_c^3 (T_c - T_{mp}) / \mu^2 . \quad (15)$$

We control the value of Gr by assigning an appropriate value to the crucible wall temperature, T_c . It should be noted that although the length scale of the problem is realistic for InP crucibles (they are much smaller than silicon growth crucibles), the Gr values considered here are smaller than the realistic ones in a LEC process[6, 36]. This simplification is not so uncommon in crystal growth modelling, since CZ and LEC systems with high Gr numbers are associated with several complicated interacting flow instability phenomena[16].

The flow problems are solved using eight, 16 and 32 parallel processors. As initial conditions, we set the velocity equal to zero and the temperature equal to the crucible wall temperature T_c everywhere. The adaptive procedure described in section 4 – “Error estimation is used”, with an error indicator based on a combination of the boundary layer indicator (11) and the magnitudes of the temperature gradient. After each adaptive mesh-refinement step, time marching is performed, with a local time step which is determined by setting the algorithm Courant number equal to 10, until the residual norm is reduced by three orders of magnitude. The dimension of the Krylov space used in the GMRES solver is chosen to be 25. In these simulations the use of the DC term is suppressed.

Figure 6(a) shows a three-dimensional view of the initial mesh used, which contains about 140,000 elements. Figure 6(b) shows a vertical cross-section through this mesh. Figures 6(c) and 6(d) show the meshes obtained after the first and second adaptive refinement steps.

These meshes contain, respectively, about 230,000 and 400,000 elements, and similar numbers of degrees of freedom. The latter is the final mesh used in this example. In each refinement step, about eight of the elements in the current mesh associated with the largest error indicator values are refined.

The mesh refinement near the walls of the crucible, and especially the top boundary, where thin boundary layers are present, is apparent. The asymmetric refinement seen in the final mesh (Figures 6(c) and 6(d)) is owing to the fact that the original mesh, being an unstructured mesh of tetrahedral elements, is not axisymmetric in nature. The error indicator calculation is slightly sensitive to the local shape and size of the elements, and thus leads to small perturbations in the error indicator values around the specified threshold value, and hence to a slightly asymmetric refinement. However, this asymmetry is reduced as the mesh is further refined, and it does not lead to significant asymmetry in the solution itself.

Figures 7(a) and 7(b) show, respectively, the initial and final mesh partitioning, when eight processors are used. The differences between the two are owing to the dynamic load balancing procedure employed. It is clear that the final processor distribution is such that more effort is concentrated near the top surface of the melt, where most of the refinement has taken place.

Figures 8-10 show the finite element results. Figure 8 is the temperature contour plot in the vertical yz plane, for $Gr = 3.3 \cdot 10^4$. Figures 9(a), 9(b) and 9(c) show the velocity pattern in a vertical plane for $Gr = 3.3 \cdot 10^4$, $3.3 \cdot 10^5$ and $3.3 \cdot 10^6$ respectively. Figure 10 shows the velocity distribution in the horizontal




Figure 6.
Melt flow of InP: finite
element meshes: (a)
initial mesh; (b) vertical
section through initial
mesh; (c) mesh after one
adaptive refinement; (d)
mesh after a second
adaptive refinement

plane $z = 3.2$, slightly below the top surface of the melt. These results are similar in principle to those presented in [7,8,13]. At the heated vertical wall of the crucible the melt rises owing to buoyancy, and then turns radially inwards at the melt surface. This gives rise to the Rayleigh-Bénard toroidal cell flow seen in Figures 9(a)-9(c). The flow cell is not centred; its centre is closer to the crucible wall, as in the “bulk-flow” model in[13]. The radial flow shown in Figure 10, which is the upper part of the toroidal cell, exists below the boundary layer adjacent to the top surface of the melt. In CZ flows, where no encapsulant is used, this radial flow appears on the free surface of the melt.




Figure 7.
Melt flow of InP: mesh
partitioning, with eight
processors: (a) initial
mesh partitioning; (b)
final mesh partitioning

We note, however, that these results are expected to differ significantly from those obtained in a fully-coupled LEC simulation. This has been demonstrated in [13], where the results of a simple CZ “bulk-flow” model and of a coupled thermal-capillary model, were compared in various configurations. The significant differences in the results of the two models, especially at high Gr numbers, show that flow coupling effects are important, and thus the uncoupled model is too simple to yield accurate information.







Figure 8.
Melt flow of InP:
temperature contour
lines in the vertical yz
plane, for $GR = 3.3 \cdot 10^4$



(a)



(b)



(c)

Figure 9.
Melt flow of InP:
velocity pattern in the
vertical yz plane:
(a) $Gr = 3.3 \cdot 10^4$;
(b) $Gr = 3.3 \cdot 10^5$;
(c) $Gr = 3.3 \cdot 10^6$




Figure 10.
Melt flow of InP:
velocity distribution in
the horizontal plane
 $z = 3.2$, slightly below
the top surface of the
melt

Although our finite element scheme solves three-dimensional problems, the steady-state results shown here are all axisymmetric. Both experimental evidence and numerical calculations with realistic value of Gr , show that CZ and LEC processes always involve flows which are time-dependent and fully three-dimensional. For silicon with buoyancy-driven flows, transition to time-dependent three-dimensional solutions typically occurs at $Gr \approx 5 \cdot 10^6$ [7, 8].

The extension of the current methodology to the three-dimensional regime of realistic Gr values is currently underway. Time-dependent simulations and some preliminary three-dimensional results for large Gr are reported in [37]. These results show that transition to unsteady three-dimensional flows occurs at about $Gr = 10^7$, and that these flows are characterized by short transient, followed by periodic oscillations with period of about 3.5 seconds. The simulations in [37] also include flows in silicon and InP melts with the combined effects of natural convection owing to buoyancy and forced convection owing to crystal and crucible rotation.

6. Concluding remarks

In this paper we have described a new parallel adaptive finite element scheme for the large-scale analysis of viscous flows. We demonstrated how this scheme can be applied to both compressible and incompressible flows, and to problems governed by the Navier-Stokes equations with the Boussinesq approximation.

Further developments are currently underway. First, the spatially linear finite elements used here may be replaced by higher-order elements. This extension is rather straightforward, and would lead to a highly accurate and stable scheme. Stability is guaranteed owing to the GLS formulation. Incorporating a general hierarchy of high-order finite elements would enable

the use of *hp*-adaptive strategies, which have known advantages over *h*-adaptive schemes in many cases.

Second, the capability of efficient time-dependent analysis is being added to the present scheme. Effective time integration, which is essential in a large-scale non-linear analysis, will be achieved using high-order singly implicit Runge-Kutta (SIRK) methods[38].

Finally, we intend to extend and apply the methodology described in this paper to the coupled multiphase transient analysis of InP high-pressure crystal growth systems. The difficulties entailed in the realistic analysis of the full LEC system, including all the important effects, are enormous and we believe that the current methodology serves as an excellent framework for such simulations.

References

1. Mittal, S. and Tezduyar, T.E., "Massively parallel finite element computation of incompressible flows involving fluid-body interactions", *Comp. Meth. Appl. Mech. Engng.*, Vol. 112, 1994, pp. 253-82.
2. Bottasso, C.L., de Cougny, H.L., Dindar, M., Flaherty, J.E., Ozturan, C., Rusak, Z. and Shephard, M.S., "Compressible aerodynamics using a parallel adaptive time-discontinuous Galerkin least-squares finite element method", in *Proceedings of 12th AIAA Applied Aerodynamics Conference*, no. 94-1888, Colorado Springs, CO, June 1994, AIAA.
3. Hauke, G. and Hughes, T.J.R., "A unified approach to compressible and incompressible flows", *Comp. Meth. Appl. Mech. Engng.*, Vol. 113, 1994, pp. 389-95.
4. Shakib, F., Hughes, T.J.R. and Johan, Z., "A new finite element formulation for computational fluid dynamics: X. The compressible Euler and Navier Stokes equations", *Comp. Meth. Appl. Mech. Engng.*, Vol. 89, 1991, pp. 141-219.
5. Hurle, D.T.J. and Cockayne, B., "Czochralski growth", in Hurle, D.T.J. (Ed.), *Handbook of Crystal Growth*, Vol. 2A, Chap. 3, North-Holland, Amsterdam, 1994, pp. 99-212.
6. Müller, G. and Ostrogorsky, A., "Convection in melt growth", in Hurle, D.T.J. (Ed.), *Handbook of Crystal Growth*, Vol. 2B, Chap. 13, North-Holland, Amsterdam, 1994, pp. 709-819.
7. Crochet, M.J., Wouters, P.J., Geysling, F.T. and Jordan, A.S., "Finite element simulation of Czochralski bulk flow", *Journal of Crystal Growth*, Vol. 65, 1983, pp. 153-65.
8. Bottaro, A. and Zebib, A., "Three dimensional thermal convection in Czochralski melt", *Journal of Crystal Growth*, Vol. 97, 1989, pp. 50-8.
9. Milhelčić, M. and Wingerath, K., "Instability of the buoyancy driven convection in Si melts during Szochralski crystal growth", *Journal of Crystal Growth*, Vol. 97, 1989, pp. 42-9.
10. Leister, H.-J. and Perić, M., "Numerical simulation of a 3D Czochralski melt flow by a finite volume multigrid algorithm", *Journal of Crystal Growth*, Vol. 123, 1992, pp. 567-74.
11. Xiao, Q. and Derby, J.J., "Three-dimensional melt flows in Czochralski oxide growth: high-resolution, massively parallel, finite element computations", *Journal of Crystal Growth*, Vol. 152, 1995, pp. 169-81.
12. Sackinger, P.A., Brown, R.A. and Derby, J.J., "A finite element method for analysis of fluid flow, heat transfer and free interfaces in Czochralski crystal growth", *International Journal of Numerical Methods for Fluids*, Vol. 9, 1989, pp. 453-92.
13. Xiao, Q. and Derby, J.J., "Bulk-flow versus thermal-capillary models for Czochralski growth of semiconductors", *Journal of Crystal Growth*, Vol. 129, 1993, pp. 593-609.

14. Derby, J.J., Brandon, S., Salinger, A.G. and Xiao, Q., "Large-scale numerical analysis of materials processing systems: high-temperature crystal growth and molten glass flows", *Comp. Meth. Appl. Mech. Engng.*, Vol. 112, 1994, pp. 69-89.
15. Zhang, H. and Prasad, V., "A multizone adaptive process model for low and high pressure crystal growth", *Journal for Crystal Growth*, Vol. 155, 1995, pp. 47-65.
16. Dupret, F. and van den Bogaert, N., "Modelling Bridgman and Czochralski growth", in Hurlle, D.T.J. (Ed.), *Handbook of Crystal Growth*, Vol. 2B, Chap. 15, North-Holland, Amsterdam, 1994, pp. 875-1010.
17. Brezzi, F. and Fortin, M., *Mixed and Hybrid Finite Element Methods*, Springer-Verlag, New York, NY, 1991.
18. Adornato, P.M. and Brown, R.A., "Petrov-Galerkin methods for natural convection in directional solidification of binary alloys", *International Journal of Numerical Methods for Fluids*, Vol. 7, 1987, pp. 761-91.
19. Chalot, F., Hughes, T.J.R. and Shakib, F., "Symmetrization of conservation laws with entropy for high-temperature hypersonic computations", *Comp. Sys. Engng.*, Vol. 1, 1990, pp. 495-521.
20. Hauke, G., "A unified approach to compressible and incompressible flows and a new entropy-consistent formulation of the $k - \epsilon$ model", PhD thesis, Stanford University, 1995.
21. Gebhart, B., Jaluria, Y., Mahajan, R.L. and Sammakia, B., *Buoyancy-induced Flows and Transport*, Hemisphere Publishing, Washington, DC, 1988.
22. Tezduyar, T.E., Behr, M., Mittal, S. and Liou, J., "A new strategy for finite element computations involving moving boundaries and interfaces – the deforming spatial domain/space time procedure: I. the concept and the preliminary tests", *Comp. Meth. Appl. Mech. Engng.*, Vol. 94, 1992, pp. 353-71.
23. Johnson, C., *Numerical Solutions of Partial Differential Equations by the Finite Element Method*, Cambridge University Press, Cambridge, 1987.
24. Hughes, T.J.R., *The Finite Element Method*, Prentice-Hall, Englewood Cliffs, NJ, 1987.
25. Hughes, T.J.R., Franca, L.P. and Hulbert, G., "A new finite element formulation for computational fluid dynamics: VIII. The Galerkin/least-squares method for advective-diffusive equations", *Comp. Meth. Appl. Mech. Engng.*, Vol. 73, 1989, pp. 173-89.
26. Franca, L.P. and Frey, S.L., "Stabilized finite element methods: II. the incompressible Navier-Stokes equations", *Comp. Meth. Appl. Mech. Engng.*, Vol. 89, 1992, pp. 141-219.
27. Shephard, M.S. and Georges, M.K., "Automatic three-dimensional mesh generation by the finite octree technique", *International Journal of Numerical Methods for Engineering*, Vol. 32, 1991, pp. 709-49.
28. Shephard, M.S., "The specification of physical attribute information for engineering analysis", *Eng with Comp.*, Vol. 4, 1988, pp. 145-55.
29. Shephard, M.S., Flaherty, J.E., de Cougny, H.L., Ozturan, C., Bottasso, C.L. and Beall, M.W., "Parallel automated adaptive procedures for unstructured meshes", in *Parallel Computing in CFD*, Vol. R-807, AGARD, Neuvilly-sur-Seine, France, 1995, pp. 6.1-6.49.
30. Beall, M.W. and Shephard, M.S., "A general topology-based mesh data structure", *International Journal of Numerical Methods for Engineering*, 1997.
31. de Cougny, H.L., Devine, K.D., Flaherty, J.E., Loy, R.M., Ozturan, C. and Shephard, M.S., "Load balancing for the parallel solution of partial differential equations", *Applied Numerical Mathematics*, Vol. 16, 1994, pp. 157-82.
32. Ozturan, C., "Distributed environment and load balancing for adaptive unstructured meshes", PhD thesis, Rensselaer Polytechnic Institute, Troy, NY, 1995.
33. Lohner, R., "An adaptive finite element scheme for transient problems in CFD", *Comp. Meth. Appl. Mech. Engng.*, Vol. 61, 1987, pp. 323-38.

HF
7,8

906

34. Flaherty, J.E., Loy, R.M., Ozturan, C., Shephard, M.S., Szymanski, B.K., Teresco, J.D. and Ziantz, L.H., "Parallel structures and dynamic load balancing for adaptive finite element computation", SCOREC Rep. No. 22-1996, Scientific Computation Research Center, Rensselaer Polytechnic Institute, Troy, 1996.
35. Saad, Y. and Schultz, M., "A generalized minimum residual algorithm for solving nonsymmetric linear systems", *SIAM J. Sci. Stat. Comp.*, Vol. 7, 1986, pp. 856-69.
36. Iseler, G.W., "Advances in LEC growth of InP crystals", *Journal of Electronic Materials*, Vol. 13, 1984, pp. 989-1,011.
37. Givoli, D., Flaherty, J.E. and Shephard, M.S., "Simulation of Czochralski melt flows using parallel adaptive finite element procedures", *Modelling Simul. Mater. Sci. Eng.* (forthcoming).
38. Moore, P.K. and Flaherty, J.E., "High-order adaptive finite element-singly implicit Runge-Kutta methods for parabolic differential equations", *BIT*, Vol. 33, 1993, pp. 309-31.