INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2004; **46**:1211–1236
Published online 4 November 2004 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fld.797

Modelling error and constitutive relations in simulation of flow and transport

Graham F. Carey[1,*,†], William Barth[2], Juliette A. Woods[1], Benjamin S. Kirk[2], Michael L. Anderson[1], Sum Chow[3] and Wolfgang Bangerth[1]

[1]*ICES, The University of Texas at Austin, U.S.A.*
[2]*Aerospace Engineering, The University of Texas at Austin, U.S.A.*
[3]*Brigham Young University, U.S.A.*

SUMMARY

The main theme of this paper is the modelling error associated with the choice of constitutive equations and associated parameter sensitivity. These key concepts are first discussed in a general way and explored in detail subsequently for three flow application problems which share certain similarities not only in the finite element formulation but also in the nature of the non-linear constitutive laws that are investigated. The specific flow problems involve: density driven flow in porous media, the Powell–Eyring model for generalized Newtonian flows, and the related Glen's flow type of constitutive relations for glacier modelling. Copyright © 2004 John Wiley & Sons, Ltd.

KEY WORDS: modelling error; constitutive equations; adaptive strategies; finite elements

1. MODEL SELECTION AND APPROXIMATION ERROR

Continuing advances in high-speed computing permit more complex non-linear problems to be simulated. Non-linearities can arise in several common ways, perhaps the most important being non-linear couplings such as the advective term in flow and transport processes, through various non-linear reaction terms and through non-linearities associated with the applicable constitutive models. The focus of the present work is on mathematical and numerical issues associated with the constitutive model assumptions and the resulting model errors. We also study the sensitivity of models to model parameters.

*Correspondence to: G. F. Carey, ICES, The University of Texas at Austin, U.S.A.
†E-mail: carey@cfdlab.ae.utexas.edu

Contract/grant sponsor: DoD HPCMP PET Programme
Contract/grant sponsor: Sandia; contract/grant number: 56522
Contract/grant sponsor: ICES Fellowship Programme

Received 24 June 2003
Revised 22 May 2004

The increase in computer power implies also that higher resolution meshes will allow fine-scale phenomena to be captured by an approximation of a specific mathematical model. Furthermore, the very presence of such non-linearities in the mathematical model implies that there may be multiple solution states, that these states may have very different structure and that mesh resolution is important in constructing a reliable discrete non-linear approximation. If the mesh is not adequate then the discrete model may not yield the right solution multiplicity, turning points, bifurcation paths, and so on for the chosen model. Ideally, a rigorous *a posteriori* error analysis should form the foundation for an error indicator construction to ensure reliability of the mesh and approximate solution. The result then is an approximation $\tilde{u}$ to the exact solution u for a given model such that the error $e = u - \tilde{u}$ is small in a specified sense [1, 2].

We will be concerned here to a lesser extent with mesh resolution effects but our main focus is on errors directly arising from the constitutive models. In this context, a single mathematical problem is *not* specified *a priori* but rather we have in mind a family of mathematical problems corresponding to a choice of constitutive models.

The above point concerning the choice of constitutive model from a family of possible models is very germane today. This is especially the case when we have a standard linear model and an option to use a more complicated non-linear model. In particular, we need to be able to assess the modelling improvement (or modelling error) resulting from the choice of one constitutive model over another. If the associated error in the approximation, based on a simple linear constitutive model, is essentially local to some subdomain, then it may be fruitful to employ the linear model throughout most of the domain and the non-linear model only in the subdomain. Such a strategy was applied to advantage for phase change problems in melting and solidification in Reference [3] where a non-linear capacitance model was confined to the local melt/solidification zone.

These ideas are far from new. Prandtl exploited similar concepts in developing boundary layer theory a century ago [4]. Related ideas also have been applied numerically to couple exterior inviscid linear potential flows to non-linear laminar boundary layer flows; similarly elastic–plastic models in solid mechanics permit invoking the non-linear material model locally based on a local yield assessment, etc.

Multiscale problems involving non-linearities have classically been treated by using singular perturbation expansion techniques that are able to identify, in a matched asymptotic sense, those subregions where different models apply. In effect, the scaling analysis permits a simpler model in appropriate subregions such as the exterior inviscid potential flow region in the example mentioned above. The perturbation analysis is a rigorous and precise way to replace a complex mathematical problem by a sequence of simpler problems that are accurate approximations with decreasing model error specified *a priori* in a matched asymptotic sense.

Another constructive technique for assessing the non-linear model has been to approach the problem in a manner analogous to determining the discretization error. Namely, the goal is to determine an analogous modelling error. This is not a difficult concept and has been applied in an informal manner for a long time in practice. For example, in Reference [5] a flow chart and algorithm are presented which assess both discretization and model error. Of course this is still a simplistic and intuitive approach from the modelling standpoint since it implies replacement of an inferior model by a better model and re-computation with the improved model when the predictions fail an elementary model reliability test (e.g. turbulent velocity regime encountered, yield stress exceeded, etc.). More generally, one may solve the problem

with the simpler model (e.g. the linear constitutive model) until the model test is violated and then, using this solution as a starting iterate, continue to solve the problem with the more complex constitutive model. This idea of 'model continuation' has been used in semiconductor device simulations, by progressing from a potential solution to a coupled drift-diffusion model and finally, perhaps, a hydrodynamic model [6]. In this case the model continuation procedure also enhances the reliability and efficiency of the solution. Moreover, the difference in the solutions obtained with the respective models is also a reliable assessment of the modelling error. This difference in model solutions is also in many respects analogous to the approach for assessing discretization error in which the solution for a uniform refinement of a grid is computed to obtain a reliability check on a given coarser mesh solution (as advocated, for instance, in the ASME fluids journal instructions concerning papers on computational fluid dynamics [7, 8]).

In order to illustrate these concepts, let us assume for the moment that we adopt this direct approach for model comparison. Let $\tilde{u}^0$ and $\tilde{u}^1$ denote the solution to the discretized problem with the respective level 0 base model and with the level 1 improved model. For example, the base (level 0) model might be a linear problem and the improved (level 1) model might be a non-linear model corresponding to an improved constitutive relation. Further, let us assume that the discretization error has been controlled using uniform or adaptive mesh refinement (AMR) in approximating the respective exact solutions u^0 and u^1. Note that the corresponding grids will, in general, be different in this scenario. The discretization errors associated with the two respective models are $e_d^0 = u^0 - \tilde{u}^0$ and $e_d^1 = u^1 - \tilde{u}^1$. The unknown 'exact' modelling error $e_M = u^1 - u^0$ is approximately given by $\tilde{e}_M = \tilde{u}^1 - \tilde{u}^0 = -(u^1 - \tilde{u}^1) + (u^1 - u^0) + (u^0 - \tilde{u}^0) = e_M - e_d^1 + e_d^0$. Hence, provided the respective model discretization errors are controlled to a small enough error tolerance then e_d^1 and e_d^0 are negligible and to first order $\tilde{e}_M = e_M$. That is, the difference between the numerical simulations will be a reliable approximation of the modelling error. However, it is also clear that care must be exercised in controlling the discretization error in both simulations. We elaborate on this point in some of the later numerical studies.

It is also clear from the preceding remarks that, in many applications, one expects the modelling error $\tilde{e}_M$ to be small in all or most of the domain so that the more refined model may be needed only to improve local behaviour. This is especially the case in well behaved elliptic boundary value problems. Hyperbolic problems and convection-dominated transport processes may be more sensitive, as are problems exhibiting instabilities, with singularities, or sensitive to 'imperfections' as seen later in the Elder–Voss–Souza (EVS) problem.

Theoretical consideration of modelling error: As an illustrative case, let us introduce the balance equation

$$N(u) = -\nabla \cdot (k(\nabla u)\nabla u) = f \tag{1}$$

where the scalar material property k in the constitutive relation depends on the solution gradient. For simplicity we assume homogeneous boundary conditions for u. The balance equation may be considered as the 'true' model, i.e. one that is most physically and mathematically accurate.

Next consider a sequence of approximate models with simplified coefficient function k_ℓ:

$$N_\ell(u^\ell) = -\nabla \cdot (k_\ell(\nabla u^\ell)\nabla u^\ell) = f \tag{2}$$

Here the index ℓ is used to indicate the level of model, with the balance equation (1) being the limit of the sequence of models. The weak solutions u and u^ℓ are assumed to lie in the

same Hilbert space V equipped with the inner product denoted by $(\cdot,\cdot)$. Let $\|\nabla w\|$ denote a norm equivalent to the $H^1(\Omega)$-norm of $w \in V$ and $\|\cdot\|^*$ the corresponding dual norm. We now derive a condition that ensures that the 'model discrepancy' for the flux

$$k(\nabla u)\nabla u - k_\ell(\nabla u)\nabla u \tag{3}$$

remains small and examine a global bound for the modelling error.

If N_ℓ is a continuous, strongly monotone operator on V, then there exists a constant $a>0$ such that

$$\begin{aligned} a\|\nabla u^\ell - \nabla u\|^2 &\leqslant (k_\ell(\nabla u^\ell)\nabla u^\ell - k_\ell(\nabla u)\nabla u, \nabla u^\ell - \nabla u) \\ &= (k(\nabla u)\nabla u - k_\ell(\nabla u)\nabla u, \nabla u^\ell - \nabla u) \\ &\leqslant \|k(\nabla u)\nabla u - k_\ell(\nabla u)\nabla u\|^* \, \|\nabla u^\ell - \nabla u\| \end{aligned} \tag{4}$$

In other words, the modelling error bound is proportional to the norm of the model discrepancy. Under the assumption that ∇u is pointwise bounded, we see that if

$$\|k(\nabla u)\nabla u - k_\ell(\nabla u)\nabla u\|^* < \varepsilon_\ell \tag{5}$$

where ε_ℓ represents some acceptable tolerance, then the modelling error satisfies $\|\nabla e\| \leqslant \varepsilon_\ell / a$. Bound (5) is achieved if the material property satisfies, for example,

$$\sup_{t \geqslant 0} |k(t)t - k_\ell(t)t| \leqslant \varepsilon_\ell \tag{6}$$

If we have two solutions u^0, u^1 corresponding to two models k_0 and k_1, then

$$r^1(u^0) = -\nabla \cdot (k_1(\nabla u^0)\nabla u^0) - f = -\nabla \cdot (k_1(\nabla u^0)\nabla u^0) + \nabla \cdot (k_1(\nabla u^1)\nabla u^1) \tag{7}$$

We have from the mean value theorem that, for any $v \in V$ vanishing on the boundary,

$$\begin{aligned} (r^1(u^0), v) &= (k_1(\nabla u^0)\nabla u^0 - k_1(\nabla u^1)\nabla u^1, \nabla v) \\ &= \left(\int_0^1 (k_1(t)t)'|_{t=s|\nabla u^0| + (1-s)|\nabla u^1|} (\nabla u^1 - \nabla u^0) \, \mathrm{d}s, \nabla v \right) \end{aligned} \tag{8}$$

Consequently, for $(k_1(t)t)'$ bounded above and below by positive constants, the modelling error is bounded globally above and below by the residual error

$$C_1 \|r^1(u^0)\|^* \leqslant \|\nabla u^1 - \nabla u^0\| \leqslant C_2 \|r^1(u^0)\|^* \tag{9}$$

An analogous result involving $\|r^0(u^1)\|^*$ may similarly be derived, and for a linear model we usually have uniform bounds on $(k_0(t)t)'$. For a higher level model with coefficient such that $(k_1(t)t)'$ is not uniformly bounded, bounds (9) on the modelling error involving $\|r^1(u^0)\|^*$ fail to hold but those involving $\|r^0(u^1)\|^*$ remain valid. However, the latter is not a useful

practical result because bounds involve u^1. Computable indicators using $r^1(\tilde{u}^0)$, where $\tilde{u}^0$ is the computed approximation with the level 0 model, are considered in more detail later.

Local model error indicators: The above observation implies that one may be able to assess model adequacy by less expensive local model error indicator approaches. This appears more difficult than in the analogous idea for mesh improvement since the underlying governing equations now are different. The idea has merit since solution of the refined model may be expensive if, in fact, it has been implemented. To initiate and motivate this line of reasoning, let us draw on ideas used in the context of mesh refinement indicators.

Let again $N_0(u^0) - f = 0$ be the level 0 model equation and $N_1(u^1) - f = 0$ be the level 1 model equation. One possibility here is to first use a simple residual-based indicator to guide the mesh improvement for the base (level 0) model and compute a numerical approximation $\tilde{u}^0$ on some given mesh. Introducing the approximation into the level 0 equation, the corresponding computable residual is $r^0(\tilde{u}^0) = N_0(\tilde{u}^0) - f$. Next, substituting $\tilde{u}^0$ in the governing equation for model level 1 we get a new residual (on the same mesh) of $r^1(\tilde{u}^0) = N_1(\tilde{u}^0) - f$. This is simply an inexpensive post-processing calculation and $r^1(\tilde{u}^0)$ represents the amount by which $\tilde{u}^0$ fails to solve the more complex model. Since $f = N_1(u^1)$ we have $r^1(\tilde{u}^0) = N_1(\tilde{u}^0) - N_1(u^1)$. Where $\tilde{u}^0$ is close to u^1, we expect $r^1(\tilde{u}^0)$ to be small. Elsewhere the approximation is not adequate and the model error may be large. Let us assume for illustrative purposes that $N_0 = L_0$ is a linear operator and $N_1 = L_0 + n_1$; that is, N_1 is the linear operator plus a new non-linear term or terms n_1. This decomposition is often encountered in practice. Then the residual for the refined model becomes $r^1(\tilde{u}^0) = N_1(\tilde{u}^0) - f = L_0\tilde{u}^0 + n_1(\tilde{u}^0) - f = r^0(\tilde{u}^0) + n_1(\tilde{u}^0)$. This implies that evaluating the contribution of the additional terms n_1 at the model 0 solution will give a guideline to the local model error and determine if and where the model needs to be refined to the level 1 model. The term $r^0(\tilde{u}^0)$ is an indicator of the discretization error for the model 0 approximation $\tilde{u}^0$. Clearly, this model adaption can be automated as part of the simulation together with adaptive mesh refinement.

Of course, we are interested in how good an approximation our numerical solution $\tilde{u}^0$ is for the exact solution u^1 of the better model described by the operator N_1. For this, let $\delta = u^1 - \tilde{u}^0$ be the error. To first order in the error δ, we then have

$$0 = N_1(u^1) - f = N_1(\tilde{u}^0 + \delta) - f \approx N_1(\tilde{u}^0) - f + N_1'(\tilde{u}^0)\delta \tag{10}$$

or, equivalently, to first order

$$N_1'(\tilde{u}^0)\delta = -r^1(\tilde{u}^0) \tag{11}$$

There are several ways to exploit this equation, all of which are related to methods that are also commonly used in the literature on residual-based *a posteriori* error estimates for AMR. The first is to explore whether it is useful to continue the solution process with the more complicated model at all: if we can show that the error δ for a given model residual $r^1(\tilde{u}^0)$ is smaller than a given tolerance, then we can stop computations with the approximation $\tilde{u}^0$ we have obtained using the simpler and presumably cheaper model. For this, assume that $N_1'(\tilde{u}^0)$ satisfies a stability estimate of the form $\|N_1'(\tilde{u}^0)^{-1}\| \leqslant C_S$ for an appropriate norm; then $\|\delta\| \leqslant C_S\|r^1(\tilde{u}^0)\|$. Such stability estimates are available, in particular for elliptic problems, and are extensively used both in the analysis of partial differential equations and the *a posteriori* estimation of finite element errors [9].

Next, let us assume that instead of the global error we are, for example, interested in the solution in a certain part of the domain, the heat flux across a boundary, or the drag on a body immersed in a fluid. In this case, assume the quantity we are interested in is a functional $J(u)$, and we ask whether $J(\tilde{u}^0)$ is good enough, or whether we need to use the more complicated model. To treat this, we use the framework presented in Reference [1]. Assuming for a moment that J is a linear and bounded functional, then it can be represented by $J(\varphi)=(j,\varphi)$ with an inner product $(\cdot,\cdot)$; in view of the linearity, the question of reliability of $\tilde{u}^0$ is equivalent to asking whether $J(\delta)$ is small enough. To answer this, define z to be the solution of the dual problem $N_1'(\tilde{u}^0)^* z = j$. Then,

$$J(\delta)=(j,\delta)=(N_1'(\tilde{u}^0)^* z,\delta)=(z,N_1'(\tilde{u}^0)\delta)=-(z,r^1(\tilde{u}^0)) \tag{12}$$

In practice, the solution z of the dual problem can also be approximated numerically. Again, if $J(\delta)$ is already small, then there is no need to do any further computations with the more expensive model N_1.

Both the formulas for $\|\delta\|$ and $J(\delta)$ can also be used to drive an adaptive procedure. Since they can be split up into local contributions for each cell of a finite element or finite difference discretization, we can use these indicators for either local mesh refinement, or for model refinement. In particular, if the problem allows one to choose different models in different parts of the domain, this can be used to select those cells where the more expensive model is preferable.

Finally, (11) can also be used to reduce the error since it is of defect correction type. It allows us to either do an approximate global solve for δ, or, for example, solve local problems in analogy to inexpensive methods for *a posteriori* error estimation in finite element methods. In both cases, a new approximation $\tilde{u}$ is obtained; since in Equation (11) we have not assumed anything about the origin of $\tilde{u}^0$, we may then apply the above observations to $\tilde{u}$ again.

Subgrid models: The following porous media flow problem indicates the relevant ideas. In porous media flows, the Darcy flow law for a homogeneous medium is the standard constitutive flow model. However, in oil reservoir simulation and similar groundwater applications involving wells, the flow rate near wells may be large, and more complex non-linear constitutive models such as Forchheimer's flow model have been proposed. Both models are actually simplifications of Navier–Stokes formulations of viscous flow problems.

Porous media flows involve scales ranging from the macroscopic behaviour associated with the Darcy flow constitutive relation at the coarse scale to flow through the fine scale pore structure of the medium. Recent homogenization and multiscale averaging approaches attempt to bridge the scale effects either at the level of the mathematical model or by subgrid approaches in an associated numerical scheme. The usual approach is to assume the same flow model and constitutive relation apply at both scales. This homogenization to the macroscale then is designed to treat the effect of variation in material properties. However, one may have a coarse grid model based on Darcy flow as the level 0 model and a multiscale model for level 1 that attempts to encapsulate the key contributions arising from the fine scale flow behaviour due to a different model. The previous strategy can be modified to assess model reliability in this scheme as follows: assume a calculation at a given timestep is made with the simpler model using the macroscopic Darcy flow law (level 0) but the local behaviour at this coarse cell block level has high flow rates; then a local solution on a subdomain using the subscale model (level 1) can be used to assess the adequacy of the Darcy model. Of course,

one could also apply this approach to every coarse grid cell in level 0, but the computational overhead is a factor that must be considered. It is well known, however, that in porous media that contain fractures or fine scale channels, simple homogenization with the usual Darcy flow law can under-predict fluid transit times by an order of magnitude or more. Hence, in certain applications such as risk assessment of radionuclide transport to the water supply, an improved capability to assess model error and invoke more reliable models is desirable.

Note that one can also extend the previous local residual argument to this multiscale subgrid situation—the approximate level 0 solution $\tilde{u}^0$ for the Darcy coarse grid model is substituted in the subgrid model equations to construct a subgrid residual. This subgrid model would involve a different constitutive model, and the associated subgrid residual can also be projected against the coarse grid finite element basis to obtain a coarse grid indicator of model error effect from the subgrid enhanced model.

Dial-an-operator capability: The ability to explore model error and the effect of different choices of constitutive models is facilitated in our subsequent studies by the use of a dial-an-operator capability in the software framework that permits changing the structure of a differential equation and also adding more equations. Hence we can make certain changes to the constitutive models with relative ease. This facilitates comparison studies of different models and we have written additional post-processing software for detailed comparison of results from different models and graphical display of the model error. Graphical displays describing the sensitivity of the solution to changes in key parameters are also developed. The dial-an-operator implementation has not yet been extended to dynamically adapt the model during a simulation. It is, however, a relatively straightforward matter for us to implement a simple residual-type post-processing capability for local model reliability assessment.

2. PARAMETER ESTIMATION AND MODEL SELECTION

In the preceding discussion we have focused primarily on model error associated with different choices or forms of constitutive relations. However, given a specific constitutive law such as the Darcy flow relation, a density–concentration relation, the Powell–Eyring viscosity relation for a generalized Newtonian flow, or the Glen's flow law that are all considered later, various parameters are needed to define the specific problem. Typically, such parameters are obtained from experimental studies. In groundwater and reservoir flow cases they are a result of averaging approaches for reservoir characterization using, for instance, seismic data and a sparse set of drilling 'core' samples. Even in well-controlled industrial manufacturing procedures there will be statistical variations in material parameters. If the model is 'robust' or, let us say, 'well-conditioned' with respect to parameter variation then the response of the model will generally be insensitive to small changes in the material parameters. On the other hand, if the model is not robust with respect to parameter changes or to other 'imperfections' then we refer to the model as 'imperfection sensitive' or an 'ill-conditioned model'. Indeed it may be the case that the underlying physical problem exhibits instabilities and is also 'imperfection sensitive'. Furthermore, adding dissipation to stabilize computations may make the simulation more robust but obscure the solution structure and the real behaviour of the underlying physical process.

There are several ways to assess sensitivity to parameter variation. For instance, one can analyse the effect of varying the parameter in a constitutive model as seen for the Powell–

Eyring viscosity model in Section 3.2 and in the sensitivity study for Glen's flow law in the glacier application considered in Section 3.3. In a related way, one can explore the numerical sensitivity to local variations in the parameter, the idea being that this may also generally be a local problem that is important, say, near singularities or in regions of high solution gradients. For example, the local variation of the residual due to a perturbation of the material parameter, reaction rate, etc. is an indicator, and a more formal solution of a corresponding local boundary value problem could be applied. Likewise, the parameters that enter the viscosity models in the generalized Newtonian flow models are often determined by curve fit to experimental data and the sensitivity of the solution to these parameters is of interest. In viscoelastic flow problems, constitutive models such as the Oldroyd-B fluid, Giesekus, and PTT models are more complex than the Powell–Eyring model considered here, but the same ideas are applicable.

Above, we have assumed that the 'superior' model N_1 and a lower level model N_0 are known. In practice, this is often not the case, as we will show in the next section for certain porous media and viscous flow models. For example, in density-driven flow, a number of models for the relationship between salt concentration and fluid density have been proposed and identification of the superior model is in question. Correspondingly, the models relating stress and strain in non-Newtonian flows may generate significantly different outcomes. Moreover, such models usually contain parameters that are not always easy to determine in a controlled laboratory experiment, such as the parameters in the glacier flow example in Section 3.3.

If no measurements for a given system are available, then model and parameter selection must necessarily be driven by modelling assumptions and physical insight. However, if measurements are available, then the best model can be chosen by a two-stage process from a number of proposed models. The first step is to take each model and adapt the parameters in it so that the resulting solution of the model best matches the experimental data. In practice, this is most often done using a minimization procedure, where the variables are the parameters in the model, and the quantity to be minimized is the misfit between what we compute as the result of the model for a set of given parameters on the one hand and the measured data on the other hand (see, e.g. References [10, 11]). This optimization is usually performed as a least squares problem. In order to formulate it in a mathematical framework, let us denote the set of parameters of model k by q_k, and the corresponding operator by $N_k = N_k(q_k)$. Note that the cardinality of the sets q_k may vary between models so that more complicated models may have more parameters than simpler ones. The solution of this model is then $u^k = N_k(q_k)^{-1} f$. To this solution, we apply the operator M that extracts measurement data; for example M might return the time-averaged heat flux through the boundary of a convection cell, or it may be the drag coefficient of a body. On the other hand, we have measured this value in an experiment, with the measurement being denoted by d. The misfit between prediction and measurement is then the residual

$$\Phi_k(q_k) = \tfrac{1}{2} \|Mu^k - d\|^2 = \tfrac{1}{2} \|MN_k(q_k)^{-1} f - d\|^2 \tag{13}$$

The goal of parameter estimation is then to find those values of q_k that minimize the misfit $\Phi_k(q_k)$, i.e. for which the prediction $MN_k(q_k)^{-1} f$ and measurement d match each other best. We will denote these 'best' parameters for model k by $q_k^* = \arg\min_{q_k} \Phi_k(q_k)$. The problem of finding q_k^* is in general not exactly solvable, due to the non-linearity of the parameter-to-output map $MN_k(\cdot)^{-1} f$. In addition, this problem requires significant computational resources, since the PDE model and its linearized adjoint have to be solved several times. However, this

is still a relatively straightforward problem if the number of parameters in q_k is small (for example less than 5 or 10).

The result of this first step is that we have obtained an optimal instance q_k^* for each parameter-dependent model. However, in general, none of the models will be able to reproduce the measurements exactly, i.e. $\Phi_k(q_k^*) \neq 0$. For each model, there will always be a residual between the best approximation and the target measurements. The second step in model selection is then to compare these residuals $\{\Phi_k(q_k^*)\}_k$ and pick the model for which the residual between measurements and prediction is the smallest, i.e. we want to know the number $k^* = \arg\min_k \Phi_k(q_k^*)$ of the best model. Unfortunately, this determination of the 'best' model is not often done in the literature, and certainly is done much less often than the proposal of a new model. Proposing new models without this model reliability check exacerbates the difficulty since it leads to a proliferation of models of which it is not known which are better or worse than others. One of the reasons for this may also be that the result is not canonical: it may be that one model describes one situation better than another model in a specific application, but it is the other way round in a different situation or application. This situation is illustrated by the density models in Section 3.1. For a study comparing different models in different situations, see Reference [12].

Both parameter estimation and model selection crucially depend on the fact that different values of parameters and different models generate different output. In other words, parameter estimation depends on the fact that Φ_k depends in a noticeable way on its argument q_k, and model selection on the fact that $\Phi_k(q_k^*)$ is significantly different from $\Phi_j(q_j^*)$ for two different models k and j. The first of these conditions is most conveniently expressed by the sensitivity matrix: since q_k^* is a minimizer of Φ_k, we can expand the misfit in the form

$$\Phi_k(q_k^* + \delta q_k) \approx \Phi_k(q_k^*) + \tfrac{1}{2}\, \delta q_k^{\mathrm{T}} H_k \delta q_k \tag{14}$$

with the matrix

$$H_k = \partial^2 \Phi_k / \partial q_k^2 \tag{15}$$

H_k is called the sensitivity matrix since small eigenvalues indicate that even large changes in parameters may induce only small changes in the output. Conversely, this also means that small changes in the measurements may induce large changes in the optimal parameter, making the problem ill-posed. In particular, the size of $(H_k)_{ii}$ is a measure for the sensitivity of changes in $(q_k)_i$, and $1/\sqrt{(H_k)_{ii}}$ for the variance of the resulting optimal value of it. Likewise, the off-diagonal elements are covariances. Note that the entries of H_k have physical units, so terms such as 'large' and 'small' need to be understood in the context of typical sizes of the respective parameters. These ideas are explored for the Laplace–Young equation in Reference [13].

If the optimal values of $\Phi_k(q_k^*)$ do not differ significantly between models, then there is no method that can distinguish the predictive quality of these models. Often, models will be over-parameterized in these cases; i.e. they contain terms and parameters that are not important for the prediction given by this model. The sensitivity matrix will then have very small entries for the parameters that control the size of these terms, a fact that can be used to simplify a model by removing terms that do not affect the output. Often, however, the output is indeed sensitive to variations in parameters and models, and we will show below several examples of the dramatic changes induced by varying a parameter or the model. In these

cases, there is a good potential for identifying optimal parameters and subsequently selecting the best model.

3. APPLICATION STUDIES

We consider three related application classes in the model studies and numerical results described in this section. Each considers model error associated with choice of constitutive relation and/or parameter selection. The first application concerns density-driven porous media flow that involves plume evolution from a surface boundary zone with high contaminant concentration. Several popular constitutive models for the density/concentration constitutive relation are considered using the dial-an-operator framework and the sensitivity of plume structure to these models is investigated. This problem is further complicated by the fact that it is inherently an 'imperfection-sensitive' problem and there is also a singularity in the boundary and initial data. The effect of variations in the porous media material properties is also briefly examined and we comment on sensitivity to distortion of elements in the computational mesh.

The dial-an-operator framework used for the first application is also exploited in the second problem class that deals with the solution of Powell–Eyring-generalized Newtonian models for dilute suspensions. This viscosity model exhibits shear thinning and such effects are naturally more pronounced near walls due to the no-slip boundary condition that induces strong local shearing. In addition to application simulations, and parametric sensitivity studies, we compare the models with the simpler Newtonian flow model and plot the associated model error using simulations at high mesh resolution on a distributed parallel Beowulf cluster. The very high resolution implies that the discretization error is negligible and thereby isolates the modelling error. The ideas in Section 1 are utilized to demonstrate a simple error indicator for the fluid models.

The third application is for a class of simplified non-Newtonian flow models corresponding to Glen's flow law that are popular in modelling the movement of glaciers [14–16]. Here we focus on the sensitivity of the solution to the parameters of Glen's flow law, but also consider model errors associated with choice of boundary conditions.

3.1. Variable-density flow and transport in porous media

The EVS problem: Flow and transport in porous media is of great practical interest in groundwater pollution modelling. Often the concentration of solute is high enough that it significantly alters the fluid density: in such cases, the governing equations for flow and transport become non-linear, and a relation between density and solute concentration is required. This type of flow has been explored numerically using the EVS problem, which is a contaminant transport problem adapted by Voss and Souza [17] from a heat-flow study of Elder [18]. The EVS problem has been investigated in many published studies, e.g. see Reference [19]. In the present study, it is used as a test problem to investigate modelling approximation related to the choice of constitutive equation for density as a function of concentration and to examine sensitivity to variations in the porous media material properties and mesh perturbations.

The domain for the EVS problem is a two-dimensional rectangular vertical cross-section of length 600 m and depth 150 m with no change assumed normal to the plane. Centred along the top boundary is a 300 m long segment with a constant concentration boundary condition

normalized to 1 unit, representing a source of solute such as that arising in a landfill site or salt lake. The bottom boundary is held at a constant concentration of zero, under the same scaling. No-flux boundary conditions apply on the vertical sides and on the remainder of the surface. Initially, the concentration throughout the region is zero and the pressure field is hydrostatic. The total time simulated in the computations is 20 years.

While various governing equations have been used to describe the physics of the EVS problem, we consider here those adopted for salt water invading an aquifer in Reference [17]:

$$\varepsilon \frac{\partial \rho}{\partial c} \frac{\partial c}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{16}$$

$$\varepsilon \rho \frac{\partial c}{\partial t} + \rho \mathbf{v} \cdot \nabla c - \nabla \cdot (\varepsilon \rho D \nabla c) = 0 \tag{17}$$

where t is time, $\varepsilon = 0.1$ is porosity, ρ is fluid density and is a function of the concentration c, $D = 3.565 \times 10^{-6}\,\mathrm{m^2\,s^{-1}}$ is molecular diffusivity and $\mathbf{v}$ is the Darcy flux. The Darcy flux is given by the relation

$$\mathbf{v} = -\frac{k}{\mu}(\nabla p - \rho \mathbf{g}) \tag{18}$$

where p is pressure, $k = 4.845 \times 10^{-13}\,\mathrm{m^2}$ is the permeability, $\mu = 10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}$ is dynamic viscosity, and $\mathbf{g}$ is the vector of gravitational acceleration. To complete this set of equations, a constitutive relation describing the dependence of density on concentration is required.

Density as a function of concentration: Here we consider constitutive relations that have been used to model sodium chloride solutions. This is important for numerical studies of salt lakes, groundwater flow near salt domes and regions of saline intrusion into coastal aquifers. Laboratory studies have shown that the saturation concentration of sodium chloride in water is $0.26\,\mathrm{kg\,kg^{-1}}$ at 20°C and that such a solution has a density of $1197.2\,\mathrm{kg\,m^{-3}}$ while freshwater density is $998.23\,\mathrm{kg\,m^{-3}}$. The change of density with concentration has been modelled in various ways [20]. Ségol [21] uses a linear function $\rho(c) = 998.23 + 700c$, where ρ is density and c is mass fraction concentration. According to Holzbecher [20], Cussler [22] also uses a linear function but with a different 'slope' parameter, $\rho(c) = 998.23 + 756c$. Intraval [23] adopts an exponential model: $\rho(c) = 998.23 \exp(0.6923c)$, while Stuyfzand [24] uses a quadratic equation for density that depends on both temperature and concentration (of total dissolved solids rather than sodium chloride alone). At 20°C it reduces to $\rho(c) = 998.336 + 759.24c - 314.6c^2$. Other similar models have been studied or proposed in References [25–27]. The various constitutive equations agree fairly well for concentrations less than seawater ($c = 0.035$). However, at concentrations approaching saturation there is a variation of almost 15% of the density difference between freshwater and saturated brine. (Note that concentrations near saturation do occur in the field, e.g. near salt domes and in naturally saline discharge zones [28].)

We now investigate model sensitivity to the choice of density relation for the EVS problem. While the physical parameters given by Voss and Souza [17] do not explicitly mention a sodium chloride solution, its density variation of $1000–1200\,\mathrm{kg\,m^{-3}}$ closely matches that of brine ($998.23–1197.23\,\mathrm{kg\,m^{-3}}$). Moreover, the EVS problem has been simulated in the literature using various constitutive equations, including linear [17], exponential [19] and the

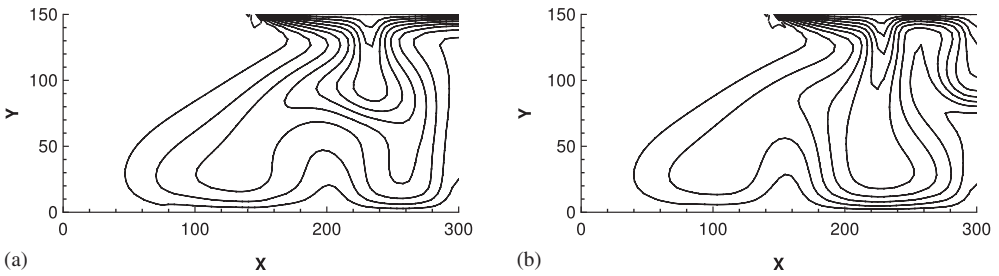

Figure 1. Concentration contours after 8 years for simulations employing the following equations for density as a function of concentration: (a) Equation (19); and (b) Equation (20) (contour intervals 0.1).

fractional form of Lever and Jackson [25]. The EVS problem is usually discussed in terms of scaled concentration $c' = c/c_s$, where c_s is the concentration at saturation, so that $c' \in [0, 1]$. Applying this scaling we obtain (for References [21–24], respectively):

$$\rho(c') = 998.23 + 184.87c' \tag{19}$$

$$\rho(c') = 998.23 + 199.66c' \tag{20}$$

$$\rho(c') = 998.23 \exp(0.1828c') \tag{21}$$

$$\rho(c') = 998.336 + 200.51c' - 21.94c'^2 \tag{22}$$

Note that no other variables are scaled in the usual presentation of the EVS problem in the literature.

The finite element scheme used to solve the EVS problem is described in detail in Reference [27]. Only the left half of the domain is simulated as the problem is taken to be symmetric with respect to the line $x = 300$. Twenty uniform timesteps are used for each simulated year and the mesh consists of 32×16 quadratic elements. Following the literature, we refer to descending plumes as 'downwelling' into the intervening low density regions of similar shape, and reversed orientation as 'upwelling'.

Concentration contours of the results are given in Figure 1 for models (19)–(20). Even though differences in the density are small, it can be seen that the choice of constitutive equation can substantially alter the numerical results. Note that (19) yields concentration contours in which a single main plume of higher salinity extends through the half-domain, with an upwelling of less saline fluid at $x = 300$ m, whereas (20) yields two plumes, one of which represents a downwelling of fluid and solute at $x = 300$ m. We remark that the results for (21) and (22) are very similar to those shown for (20) and (19), respectively, and therefore are not graphed here. Note that (19) and (22) both considerably underestimate fluid density at high concentration values. The concentration contours generated by (20) and (21) are visually indistinguishable.

Through this example, it can be seen that variable-density flow and transport in a porous medium can be sensitive to the choice of constitutive equation for density. This suggests that model errors, as distinct from numerical errors, are important in explaining discrepancies in

simulation results for the EVS problem. However, other modelling sensitivity issues such as those associated with material properties are relevant and considered next.

Variations in porous media property: To test the sensitivity of EVS simulations to small variations in the physical parameters, the permeability is locally perturbed so that small inhomogeneities are present in the matrix material. While in many field applications permeability may vary by one or two orders of magnitude, here we consider only very slight changes, $k = \bar{k} \pm \delta$ (at the Gauss points of each element) where $\bar{k}$ is the mean permeability and δ is a random function of uniform distribution and maximum magnitude $\bar{k}/200$ or $\bar{k}/2000$. Constitutive relation (20) for density was chosen for the numerical test. As before, a 32×16 quadratic element mesh and 20 timesteps per year were used.

It was found that the random variation in the permeability of $\bar{k}/2000$ had little discernible impact on concentration contours when compared to the appropriate uniform permeability case. That is, simulations yielded a downwelling solution that was indistinguishable from the uniform permeability results of Figure 1(b). However, a $\bar{k}/200$ variation in permeability did have a significant impact on results. In a series of 10 realizations in which each case differed only in the seed value given to the random number generator, both downwelling and upwelling results were observed, even though the uniform-permeability case for (20) yields a downwelling solution (i.e. fluid movement for $x = 300$ m is predominantly downwards). The $c' = 0.2$ contours are plotted for the eighth year in Figure 2 to give an impression of the variation in results. In field applications, it is often important to consider outlier scenarios in which contaminant plumes spread quickly, and in a similar plot of $c' = 0.6$ contours there is one realization in which one downwelling contour extends more than twice the depth of the others near $x = 200$, illustrating the importance of small variations in permeability to the spread of solute.

The mean concentration averaged over all 10 realizations has plumes which are truncated when compared to the uniform-permeability case plotted in Figure 1(b). This is not surprising, as it averages over both downwelling and upwelling solutions. The standard deviation is highest along the line $x = 300$ m, which is precisely where the largest discrepancies (upwelling vs downwelling) in the series of simulations occur.

Similar behaviour was noted in a series of simulations using (22) as the constitutive relation for density: a variability of $\bar{k}/2000$ in permeability again made no discernible difference when

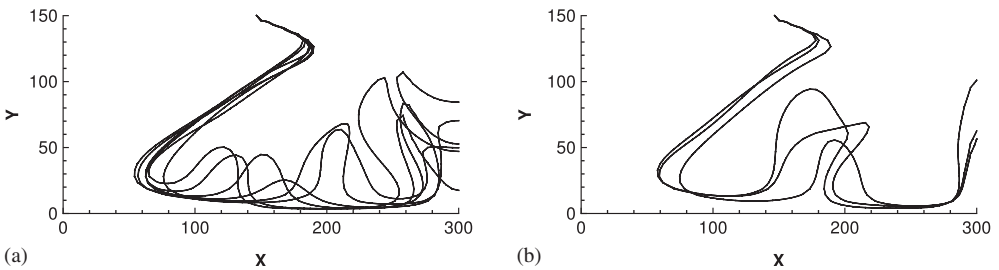

Figure 2. Concentration contours after 8 years for 10 different realizations of a $\pm 0.5\%$ variability in permeability using model Equation (20): (a) $c' = 0.2$ for downwelling solutions; and (b) $c' = 0.2$ for upwelling solutions.

compared to the uniform-permeability case, whereas a $\bar{k}/200$ variability did lead to major differences in results such that some realizations produce downwelling results and some produce upwellings. It was found that the maximum standard deviation in nodal concentration values after 8 years was higher at 0.3 for the model (22), as compared to 0.275 for the model (20).

Mesh effects: Local element distortion by moving each node by a random amount, up to some percentage of the mean element width h has a similar effect to that observed above. For example, a 1% h perturbation leads to significant changes in plume development for computations with (20) on a 32×16 mesh. All results indicate that the EVS problem is sensitive to element distortion for this choice of surface boundary segment having $c' = 1$.

Small numerical oscillations emanating from the singularity at $(150, 150)$ may also influence model behaviour. Accordingly, parallel computations for uniformly and adaptively refined meshes were carried out. Here we present a result from these simulations, which were run with the `libMesh` [29] finite element library. The error indicator chosen to guide the adaptive refinement process was based on the concentration flux jump across element faces [30]:

$$\eta_k^2 = \sum_i \frac{h_i}{24} \int_{\Gamma_i} \left[\frac{\partial c}{\partial n} \right]^2 \mathrm{d}\Gamma \tag{23}$$

where $[\partial c/\partial n]$ denotes the jump in the normal component of the concentration gradient across element edge Γ_i of cell k and h_i is the length of the edge. Only the concentration field was considered in the error indicator.

Adaptive mesh refinement provides the high resolution necessary to accurately model the behaviour near the singularity. The adaptive algorithm was applied to an initial, uniform mesh of 32×16 bi-quadratic finite elements. At the end of each timestep, indicator (23) was employed to select the elements with the highest error for refinement. The top 10% of elements, sorted in order of increasing error, were refined and the time step was repeated. To avoid over-refinement of the singularity only four levels of refinement over the initial grid were allowed. Figure 3 depicts part of the mesh and the concentration field near the singular point. The concentration field is virtually identical to that of Figure 1(b), which implies that the discretization error associated with the uniform 32×16 grid is sufficiently

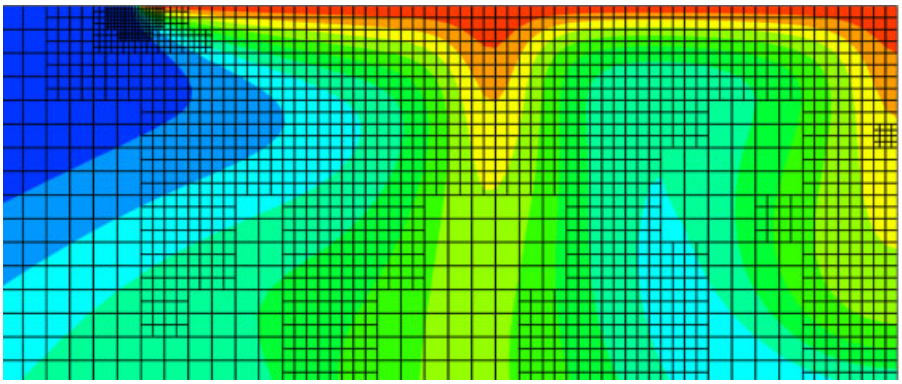

Figure 3. Adaptively refined mesh and concentration contours after 8 years for the subregion near the singular point.

small to accurately predict the behaviour of the plume structures and allow for objective model comparisons. It also shows that the under-resolved behaviour near the singular point does not introduce substantial errors in the global plume structure.

Hence the EVS problem is to some degree physically sensitive to constitutive model material imperfections, and mesh imperfections. This, in turn, implies that imperfections introduced in numerical discretization and even through the choice of convergence tolerances in the algorithms may also have significant effects.

As shown in Reference [27], however, this plume sensitivity is related to the length of the zone with $c' = 1$ on the top boundary. The length chosen by Elder and Voss–Souza is close to a critical value that separates downwelling and upwelling plume structures. As a consequence, the problem is particularly sensitive to perturbations.

3.2. *Incompressible non-Newtonian fluids*

Many viscous flow applications can be adequately modelled using a Newtonian constitutive model assumption but in some applications, generalized Newtonian models are needed. The increasing complexity of these models often implies a significant increase in degree of difficulty of the associated computations. Here we consider the Powell–Eyring constitutive model and investigate modelling error by comparison to the Newtonian assumption.

The equations of interest are the incompressible Navier–Stokes equations:

$$\rho\left(\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u}\cdot\nabla\boldsymbol{u}\right) = \nabla\cdot\boldsymbol{\sigma} \tag{24}$$

$$\nabla\cdot\boldsymbol{u} = 0 \tag{25}$$

for momentum and mass conservation, respectively, where $\boldsymbol{u}$ is the velocity vector, $\boldsymbol{\sigma}$ the stress tensor, and ρ the density. The constitutive law for the generalized Newtonian fluids considered here is

$$\sigma_{ij} = -pI_{ij} + 2\mu(s(\boldsymbol{u}))D_{ij}(\boldsymbol{u}) \tag{26}$$

where I_{ij} is the identity tensor, $D_{ij}(\boldsymbol{u})$ is the strain rate tensor given by

$$D_{ij}(\boldsymbol{u}) = \tfrac{1}{2}(u_{i,j} + u_{j,i}) \tag{27}$$

and $\mu = \mu(s(\boldsymbol{u}))$ is the viscosity with

$$s(\boldsymbol{u}) = \sqrt{2D_{ij}(\boldsymbol{u})D_{ij}(\boldsymbol{u})} \tag{28}$$

The case $\mu =$ constant in (26) corresponds to the familiar Newtonian fluid model.

Powell–Eyring constitutive relation: The three-parameter Powell–Eyring apparent viscosity model for non-Newtonian fluid behaviour is an interesting model for suspensions of polymers in solvents and polymer melts with low elasticity. However it has received relatively little attention in the modelling and simulation areas despite its favourable performance [31] and thermodynamic background. The Powell–Eyring model is based on Eyring reaction-rate theory

giving it a strong thermodynamic underpinning [32]. The corresponding viscosity relation for the model is given by

$$\mu(s(\boldsymbol{u})) = \mu_\infty + (\mu_0 - \mu_\infty)\,\frac{\sinh^{-1}(\lambda s(\boldsymbol{u}))}{\lambda s(\boldsymbol{u})} \tag{29}$$

where the parameters μ_0 and μ_∞ are the limiting viscosities at zero and infinite strain rates and λ is a characteristic relaxation time for the fluid.

For $\mu_\infty > 0$, it is easy to check that the viscosity function satisfies

$$m_0 \leqslant \mu(s) \leqslant M_0 \quad \text{for all } s \geqslant 0 \tag{30}$$

$$m_1 \leqslant \frac{\mathrm{d}}{\mathrm{d}s}(\mu(s)s) \leqslant M_1 \quad \text{for all } s \geqslant 0 \tag{31}$$

for some positive finite constants m_0, m_1, M_0, and M_1. Consequently, for the stationary problem with homogeneous Dirichlet boundary conditions specified for the velocity, the natural function space $V \times Q$ containing a weak solution $(\mathbf{u}, p)$ is

$$V = [H_0^1(\Omega)]^3, \quad Q = \left\{ q \in L^2(\Omega);\ \int_\Omega q\,\mathrm{d}x = 0 \right\} \tag{32}$$

We shall denote, respectively, the norms over V and Q by $\|\cdot\|_V$ and $\|\cdot\|_Q$ and the dual spaces by V' and Q'.

Finite element approximation: A mixed Galerkin finite element method with a continuous piecewise quadratic velocity basis and a continuous piecewise linear pressure basis is implemented using a Newton–Raphson algorithm. The resulting linear Jacobian systems are solved with a Krylov solver (in this case parallel ILU preconditioned element-by-element BiCGStab).

If $(\boldsymbol{u}_h, p_h) \in V^h \times Q^h$ is the mixed finite element approximate solution to the incompressible stationary Stokes equations (i.e. (24)–(25) without the convective term), then it may be shown by standard arguments that the non-linear operator is monotone and there exists a constant C such that the finite element approximation exhibits optimal order convergence

$$\|\boldsymbol{u} - \boldsymbol{u}_h\|_V^2 \leqslant \inf_{(\mathbf{v}_h, q_h)} C(\|\boldsymbol{u} - \mathbf{v}_h\|_V^2 + \|p - q_h\|_Q(\|\boldsymbol{u} - \mathbf{v}_h\|_V + \|\boldsymbol{u} - \boldsymbol{u}_h\|_V)) \tag{33}$$

If, in addition, the discrete inf–sup condition for the Stokes problem holds, then

$$\|p - p_h\|_Q \leqslant \inf_{q_h \in Q_h} C(\|\boldsymbol{u} - \boldsymbol{u}_h\|_V + \|p - q_h\|_Q) \tag{34}$$

Furthermore, *a posteriori* finite element error estimates may also be obtained using the techniques described in Reference [33]. Let $r_u(\mathbf{u}_h, p_h)$ denote the residual of the momentum equation, then

$$\|\mathbf{u} - \mathbf{u}_h\|_V \leqslant C\|r_u(\mathbf{u}_h, p_h)\|_{V'} \leqslant \left(C \sum_{k \in \mathscr{T}_h} \eta_k^2 \right)^{1/2} \tag{35}$$

with the local error indicator η_k given by

$$\eta_k^2 = h_k^2 \|\mathbf{f}\|_{L^2(k)}^2 + \sum_i h_i \int_{\Gamma_i} \left[\mu(s(\mathbf{u}_h)) \frac{\partial u_h}{\partial n} \right]^2 \mathrm{d}\Gamma \tag{36}$$

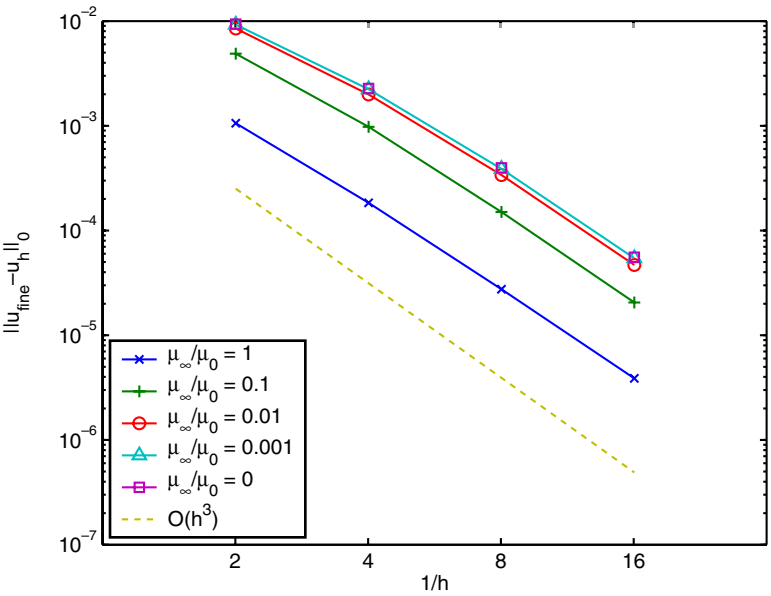

Figure 4. Error as a function of the mesh width for Stokes flow in a square channel.

where h_k is the diameter of element k in the triangulation $\mathscr{T}_h$, h_i is the length of the edge Γ_i of k and $[\cdot]$ is the jump across an edge. A similar result holds for the pressure residual.

Figure 4 shows the results of a mesh refinement study for two-dimensional Stokes flow in a channel of height H, for a Powell–Eyring fluid at viscosity ratios, μ_∞/μ_0, of 1 (the Newtonian case), 0.1, 0.01, and 0.001, with dimensionless relaxation time, $\lambda^* = \lambda U/H = 100$. Also plotted is the expected $O(h^3)$ rate line for comparison (where h is the mesh width). (To compute the error for this non-linear problem where an exact solution is not known, the solution u_{fine} on a fine grid is used.)

In order to investigate the modelling error made by treating a non-Newtonian fluid of this class as Newtonian, we consider flow in a 2-D channel of height H, with uniform inflow velocity U (smoothed by C^1 piecewise quadratics near the wall), zero-stress boundary condition and parallel flow at the outflow, at Reynolds number $Re = \rho UH/\mu_0$ of 1, 10, and 100, for fluids with μ_∞/μ_0 chosen with five logarithmically spaced values in the interval $[Re/1000, 1]$ and λ^* chosen as 10 and 100. The smoothed inlet conditions were employed to prevent reverse flow and overshoot conditions at the inflow associated with interpolation of uniform inflow onto the quadratic basis.

These cases were designed to correspond well with dilute solutions of long-chain polymers dissolved at varying concentrations in a solvent (e.g. nitro-cellulose in butyl acetate, see Reference [34]). The model error in approximating the flow as Newtonian is of interest. The global relative error in L^2 norm, $\|u_h^{\text{Newt}} - u_h\|_0/\|u_h\|_0$, is graphed versus viscosity ratio in Figure 5 for each of the cases specified above. The error incurred ranges between 3 and 20% in the cases simulated here and is exponential (or worse) in μ_∞/μ_0. Mid-channel velocity profiles for the various cases are plotted in Figure 6. Each group of profiles represents an

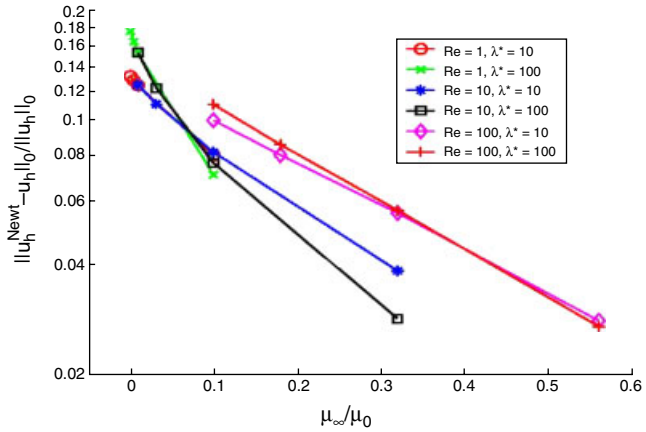

Figure 5. Modelling error when comparing the solution of the non-Newtonian model to the solution of a Newtonian approximation.

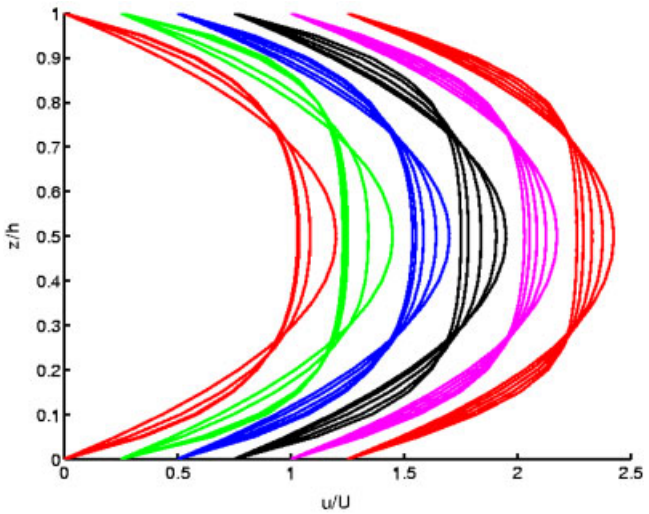

Figure 6. Solution profiles: from left to right, groups of $(Re, \lambda^*) = (1, 10)$, $(1, 100)$, $(10, 10)$, $(10, 100)$, $(100, 10)$, and $(100, 100)$.

(Re, λ^*) pair as in Figure 5 and increasing bluntness of the profile corresponds to decreasing μ_∞/μ_0. It is evident that decreasing Re and increasing λ^* also tend to flatten the profile.

From (26) and (29), we recall that the behaviour of this model is determined by the apparent viscosity in a manner analogous to that discussed in Section 2 concerning the effect of the non-linear material property k and associated 'model discrepancy'. (Recall (3)–(5).) Following

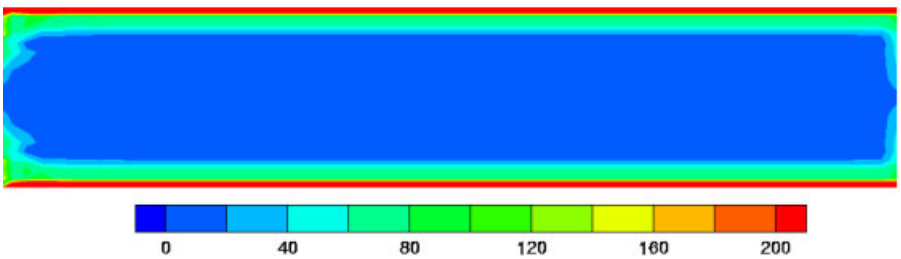

Figure 7. Local model discrepancy, $\mu_\infty/\mu_0 = 0.01$ and $\lambda^* = 10$.

this idea, we investigate the modelling error associated with the apparent viscosity by plotting in Figure 7 the contour map of the model discrepancy (residual for the horizontal velocity component) for the case $\mu_\infty/\mu_0 = 0.01$ and $\lambda^* = 10$ obtained by using the Newtonian velocity in the Powell–Eyring viscosity model. For this channel flow problem there are noteworthy differences near the two no-slip walls, where the Newtonian viscous stress differs significantly from the stress for the Powell–Eyring fluid. This figure also suggests that one may introduce the non-Newtonian model locally in the boundary layer regions identified by this simple model error. There are also differences in the inflow and outflow regions where the boundary conditions impose different stresses on the different fluid types.

3.3. *Glacier flow*

The movement of a glacier can be approximated as a slow flowing non-Newtonian fluid. Several variants of the (implicit) non-linear stress–strain relationship, or flow law $D_{ij}(\mathbf{u}) = f(\sigma)\sigma_{ij}$, have been used, where $D_{ij}(\boldsymbol{u})$ and σ_{ij} are again the strain rate and stress tensors, respectively, and $\sigma = \sqrt{\sigma_{ij}\sigma_{ij}}$. Examples that have been described under the classification of a 'Glen's flow law class' include:

$$D_{ij}(\boldsymbol{u}) = f_1(\sigma)\sigma_{ij} = A\sigma^{n-1}\sigma_{ij} \tag{37}$$

$$D_{ij}(\boldsymbol{u}) = f_2(\sigma)\sigma_{ij} = \frac{A}{\lambda^{n-1}}\frac{\sinh(\lambda\sigma)^n}{\lambda\sigma}\sigma_{ij} \tag{38}$$

$$D_{ij}(\boldsymbol{u}) = f_3(\sigma)\sigma_{ij} = A(T_0^2 + \sigma^2)^{(n-1)/2}\sigma_{ij} \tag{39}$$

$$D_{ij}(\boldsymbol{u}) = f_4(\sigma)\sigma_{ij} = A(T_0^{n-1} + \sigma^{n-1})\sigma_{ij} \tag{40}$$

Relations (38), (39) are presented in Reference [35], while (40) is used in Reference [15] and (41) in Reference [36]. These functions f_i above are shown for a typical choice of their parameters ($A = 1.0$, $n = 3.2$, $\lambda = 1.0$, $T_0 = 1.0$) in Figure 8. For $T_0 = 0$ all four laws are similar for small σ, with (37), (39), (40) being identical. For $T_0 \neq 0$, (39) and (40) are similar for small σ. For large σ, (37), (39) and (40) have the same rate of growth. Also for the special case $n = 3$, (39) and (40) are the same.

As these equations are used to simulate real glaciers, it is important to understand the differences that the choice of flow law induces. Clearly, if the choice of flow law is in question and the associated model error may be significant, then controlling the discretization error to

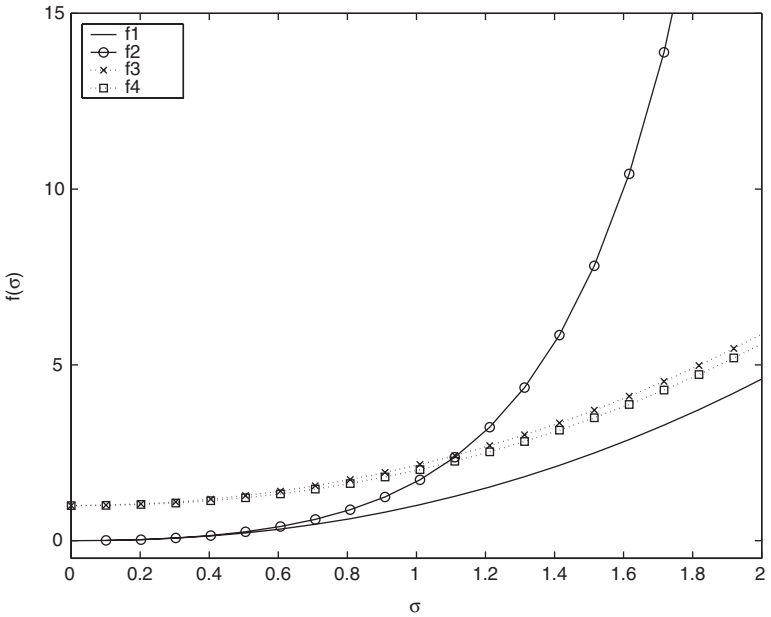

Figure 8. Glacier flow laws for $n = 3.2$, $T_0 = 1.0$, $A = 1.0$, $\lambda = 1.0$.

the level of the model error would be appropriate. Secondly, the parameters in these flow laws are not known accurately and are known to depend on other complicating factors such as temperature and the nature of the local ice crystal structure, which may vary throughout the glacier [36]. Thus the degree to which the problem is robust with respect to these parameters will determine the practical utility of any numerical simulations, indicating whether it is sensible to draw quantitative or qualitative conclusions from the models.

A related issue is the sensitivity of the problem to modelling boundary conditions (e.g. see Reference [37] for more on this issue). A slip region is introduced by Colinge and Rappaz [15] to model base melting at part of a contact region. This transition from 'slip' to 'stick' can be modelled in various ways. In Reference [15] the 'point' model of the adjacent stick–slip zones introduces a discontinuity in the gradient of the solution. This implies that large local gradients may be poorly approximated in the flow model. The presence of the singularity also degrades the rate of convergence [14]. Alternatively, a regularized transition zone model for the boundary condition may be a more realistic approximation to the physics and also yield a more tractable numerical problem. The model differences are clearly of interest.

A common model for glacier flow treats the glacier as an idealized ice slab sliding down an inclined plane. The axes are chosen so that x is in the direction of flow, z is vertical, and y is across the glacier. The ice slab is considered an incompressible viscous non-Newtonian fluid. To simplify the problem, several assumptions are typically used: (i) steady flow; (ii) plane strain approximation: no y dependency; and (iii) perturbation expansion using a scaling parameter: $\varepsilon = H/L \ll 1$, with L, H the characteristic horizontal/vertical extent of the ice sheet. Introducing these assumptions into the momentum balance, mass conservation and stress–strain

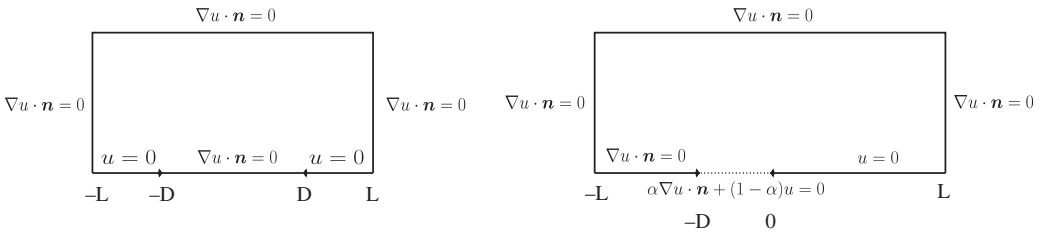

Figure 9. Domains used for testing behaviour of Glacier flow laws.

relations the model reduces to the simplified 2-D glacier model of Colinge and Blatter [38]:

$$-\nabla \cdot (k(|\nabla u|)\nabla u) = \tfrac{1}{2} \tag{41}$$

over $\Omega = (-L, L) \times (0, 2)$, with $k(|\nabla u|) = 1/f(\sigma) = 1/f(\sigma_{xx}, \sigma_{xz})$ and

$$f(\sigma)\sigma_{xx} = \frac{\partial u}{\partial x}, \quad f(\sigma)\sigma_{xz} = \frac{\partial u}{\partial z} \tag{42}$$

where u is the glacier velocity along the slab and f is the stress–strain rate relationship.

In the present study two different rectangular domains were used—one to investigate the effect of the flow law and flow law parameters on solution behaviour, and the second to examine the effect of smoothing a stick–slip zone to locally regularize the problem. Both domains are of the form $[-L, L] \times [0, 2]$ with $\nabla u \cdot \mathbf{n} = 0$ on the lines $z = 2$ and $x = \pm L$. On $z = 0$ the first domain has a slip zone located between two non-slip zones. In particular, $\nabla u \cdot \mathbf{n} = 0$ on $z = 0$ for $|x| < D$ and $u = 0$ outside this region. The second domain uses the boundary condition $\alpha u + (1 - \alpha)\nabla u \cdot \mathbf{n} = 0$ on $z = 0$. Here $\alpha = \alpha(x) = 0$ for $x < -D$, $\alpha(x) = 1$ for $x > 0$ and $\alpha(x)$ has smooth transition between 0 and 1 over the interval $(-D, 0)$. The function used here is the cubic $\alpha(x) = 3/D^2(x + D)^2 - 2/D^3(x + D)^3$, which has $\alpha(-D) = 0$, $\alpha(0) = 1$, $\alpha'(-D) = 0$ and $\alpha'(0) = 0$. The two domains are shown in Figure 9. Unless stated otherwise calculations were performed on uniformly refined grids. The finest level of the mesh usually had about 100 000 degrees of freedom.

Solution procedure: Linear convergence of successive approximation iteration has been established theoretically under certain assumptions in Reference [14]. This scheme is demonstrated to be robust in the present simulations whereas Newton iteration fails in some cases because the initial iterate is outside the contraction region. Therefore, we first apply several iterations of successive approximation and then revert to Newton iteration which converges quadratically near the solution. These algorithms were implemented using the deal.II finite element library [39].

Dependence on flow law: Comparisons between the solutions for the stated flow laws were carried out with the parameters $A = 1.0$, $n = 3.2$, $\lambda = 1$ and $T_0^2 = 0.1$ in a domain of length 60 ($L = 30$) having a slip zone of length 4 ($D = 2$). The contour map of u for a typical solution is shown in Figure 10, where the solution ranges from $u = 0$ on non-slip regions to $u = 3.09$ in the centre of the slip region. Although the general behaviour of solutions for the four flow laws was similar, the values for the flow law (38) were inconsistent with the others, differing by approximately a factor of 3. One possible source of this discrepancy is the large difference

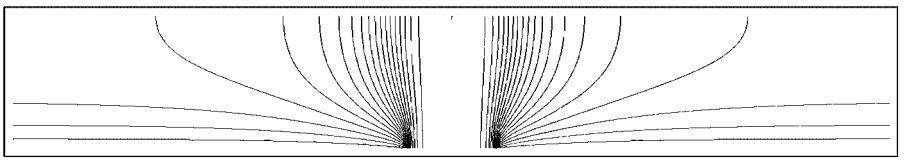

Figure 10. Typical solution behaviour for the first test problem. This example was for flow law (39) with $A = 1.0$, $n = 3.2$, $T_0^2 = 0.1$.

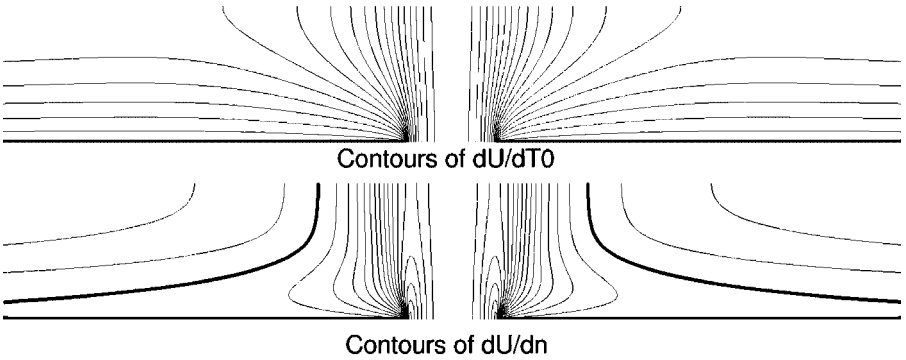

Figure 11. The sensitivity of u to the parameters T_0 and n shown as contours of $\partial u/\partial T_0$ and $\partial u/\partial n$ for Equation (39). Zero contours are shown darker. (Contour spacing for $\partial u/\partial T_0$ is approximately 0.12 and for $\partial u/\partial n$ it is approximately 0.03.)

in the flow laws in regions of high velocity gradient, such as near the interface between the stick–slip boundary conditions.

For the other three laws the solutions were only slightly different. The relative differences between solutions using laws (37) and (39) were around 15% (9% near the non-slip boundary, 11% in the slip region and 16% on the surface away from the slip region). The differences between (39) and (40) were around 4% (2% near the non-slip boundary, 2.7% in the slip region and 4% on the surface away from the slip region). Note that absolute errors were largest in both cases at the surface of the slip region and zero on the non-slip boundary. In both cases the modelling error due the choice of constitutive equation is considerable.

Dependence on model parameters: Let us consider flow law (39). The problem then has two parameters, T_0 and n for which there are a wide range of acceptable values. For example, values reported by Paterson [36] for n ranged between 2.8 and 5.2 while values for A ranged from 0.57×10^{-15} to $5.6 \times 10^{-15}\,\mathrm{s}^{-1}\,\mathrm{kPa}^{-3}$. No information about the variability of T_0 was available. As these parameters are also temperature dependent, they vary throughout the glacier. Thus it is important to determine how sensitive the solution is to variations in these parameters. This is achieved here by comparing a variety of simulations.

Although there was no obvious qualitative sensitivity to the parameters, significant quantitative changes occurred throughout the domain. These can be seen in Figure 11. In particular

small changes δT_0 in T_0 result in changes of order δT_0 near the surface of the glacier, with a peak of $\partial u/\partial T_0 \approx 2.4$ above the slip zone (where $u \approx 3$). The difference decreases smoothly to zero on the non-slip boundary but is non-zero for the slip boundary. Small changes in n introduce a small change throughout the domain of order δn, with larger changes in the slip zone. Values for $\partial u/\partial n$ range from -0.1 on the upper surface to 0.5 near the stick–slip interface. Due to the high level of uncertainty in the model parameters, error due to parameter selection will be comparable to the error due to choice of constitutive equation.

Boundary condition effects: The stick–slip boundary causes a discontinuity in the gradient of the solution, which decreases the global convergence rate of the solution. Due to the large changes in ∇u locally, this model of the boundary condition has significant implications on the behaviour of the respective flow models (37)–(40). Errors in the computed local gradient will exacerbate the constitutive model discrepancies. Adaptive mesh refinement toward the singular point can alleviate this problem [14]. At a sufficient distance away from the singular point the interior estimates for the approximate solution will achieve the optimal rate even in the non-adaptive case. However, for practical applications adaptive refinement will clearly reduce only the discretization error, but not the modelling error associated with either the choice of constitutive model parameters or the model for the base slip boundary condition. The use of a pure stick–slip model and the associated singularity can lead to other difficulties concerning the validity of the non-linear model for rough data, convergence of the successive approximation and Newton algorithms, and possibly oscillations due to the inadequacy of the local mesh.

Locally smoothing the boundary data may introduce additional local modelling error but also regularizes the local solution behaviour and improves the global approximation properties. In particular, here we approximate the boundary conditions using a non-constant Robin boundary condition of the form $\alpha \nabla u \cdot \mathbf{n} + (1-\alpha)u = 0$ where $\alpha = \alpha(x)$ varies from 0 to 1 in the smoothing region. There is a conversion of local numerical error due to the singularity to local modelling error due to the changed boundary conditions.

In a numerical test, the smoothing was applied over a distance of $D = 0.1$ in a domain of length 12 ($L = 6$). The smooth problem has superior accuracy and rates of convergence. The magnitude of the difference between the two solutions was observed to be $O(D)$ in a region of size $O(D)$ around the stick–slip point, where D was the length of the smoothed region. The change in the behaviour of the solution near the transition can be seen in Figure 12. In particular, note that there is no longer an abrupt change in the angles of the contours near the transition. The change in the solution away from this region was significantly smaller.

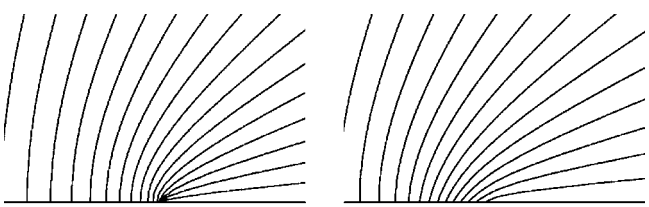

Figure 12. Close-up views of contours near a sharp (left) or smoothed (right) stick–slip transition.

Even near the stick–slip region the error associated with the regularization was significantly smaller than that associated with the uncertainty in the parameters.

We conclude that if the interest is in the global behaviour, it is acceptable to smooth the boundary conditions to provide a better rate of convergence. If the local behaviour near a 'true' singularity is of interest then smoothing the boundary will introduce a modelling error which should be avoided if it is larger than the other sources of error. In this example these effects were significantly smaller than the errors associated with uncertainty in the choice of flow law and its parameters. Consequently smoothing the interface did not degrade the accuracy of the solution further.

4. CONCLUSION

The main theme of the investigation is the assessment of modelling error, particularly with reference to choice of constitutive models. We also consider the question of sensitivity of the solution to the parameters that enter these constitutive models. The ideas we explore are, of course, very general but we have elected to study them here within the context of flow and transport problems. The three related applications are the investigation of constitutive relations for density driven flow and transport in porous medium contaminant plumes, simulation of Powell–Eyring-type generalized Newtonian flow relations in viscous flow, and Glen's flow law in simulation of stresses in glaciers. Adaptive mesh refinement is also applied in some of the studies to control discretization error, particularly in the vicinity of a singularity. The sensitivity to imperfections in material properties or even the mesh as well as those resulting from the nature of the underlying physical problem are considered here but not in depth. This is a subject of growing research interest. Likewise, in our sensitivity studies with respect to parameters and material properties we have considered the effect of perturbations and especially local effects. The general issue of model error implications concerning uncertainty and risk assessment are important. Finally, we remark that, in research and practice, attention has focused on reliability of computations and the use of adaptive refinement for error control of a specified model rather than other sources of error. However, it is well recognized that models involve assumptions and approximations. This implies that model errors may in some cases be quite significant. Here we first discuss this issue and describe some approaches for simple model error indicators analogous to those used in adaptive mesh refinement. We illustrate some of these issues and approaches in the applications and we emphasize that control of discretization error in situations where model error is in question needs careful attention.

ACKNOWLEDGEMENTS

This work has been supported in part by the DoD HPCMP PET Programme, Sandia Grant #56522 and the ICES fellowship programme.

REFERENCES

1. Bangerth W, Rannacher R. *Adaptive Finite Element Methods for Differential Equations.* Birkhäuser Verlag: Basel, 2003.
2. Carey GF, Anderson ML, Carnes B, Kirk B. Some aspects of adaptive grid technology related to boundary and interior layers. *Journal of Computational and Applied Mathematics* 2004; **16**:55–86.

3. Runnels SR, Carey GF. Finite element simulation of phase change using capacitance methods. *Numerical Heat Transfer Part B* 1991; **19**:13–30.
4. Prandtl L. On fluid motions with very small friction. *Verhdlg. 3. Intern. Math. Kongr*., 1904, also *NACA Technical Memorandum No. 452*, 1928.
5. Carey GF, Oden JT. *Finite Elements*: *Computational Aspects.* Prentice-Hall: Englewood Cliffs, NJ, 1984.
6. Carey GF, Richardson WB, Reed C, Mulvaney B. *Circuit*, *Device and Process Simulation*: *Mathematical and Numerical Aspects*. Wiley: New York, 1996.
7. Freitas CJ. Editorial policy statement on the control of numerical accuracy. *Journal of Fluids Engineering* (ASME) 1993; **115**(2):339–340 (originally presented in [8] "... it will not accept for publication any paper reporting the numerical solution of a fluids engineering problem that fails to address the task of systematic truncation error testing and accuracy estimation.").
8. Roache PJ, Ghia K, White F. Editorial policy statement on the control of numerical accuracy. *Journal of Fluids Engineering* (ASME) 1986; **108**(1):2
9. Verfürth R. *A Review of A Posteriori Error Estimation and Adaptive Mesh Refinement Techniques.* Wiley/Teubner: New York, Stuttgart, 1996.
10. Biswas D, Carey GF. Least-squares finite-element method to predict areal composition variations in large hydrocarbon reservoirs. *SPE Journal* 1998; 307–315.
11. Tarantola A. *Inverse Problem Theory.* Elsevier: Amsterdam, 1987.
12. Babuška I, Jerina K, Li Y, Smith P. Quantitative assessment of the accuracy of constitutive laws for plasticity with an emphasis on cyclic deformation. *Technical Note BN-1146*, University of Maryland, College Park, USA, 1993.
13. Anderson M, Bangerth W, Carey G. Analysis of parameter sensitivity and experimental design for a class of non-linear partial differential equations. *International Journal for Numerical Methods in Fluids* 2004 (accepted).
14. Chow S, Carey GF, Anderson ML. Finite element approximations of a glaciology problem. *M2AN Mathematical Modelling and Numerical Analysis* 2004, in press.
15. Colinge J, Rappaz J. A strongly nonlinear problem arising in glaciology. *M2AN Mathematical Modelling and Numerical Analysis* 1999; **33**(2):395–406.
16. Picasso M, Rappaz J, Reist A, Funk M, Blatter H. Numerical simulation of the motion of a two dimensional glacier. *International Journal for Numerical Methods in Engineering* 2004; **60**:995–1009.
17. Voss CI, Souza WR. Variable density flow and solute transport simulation of regional aquifers containing a narrow freshwater-saltwater transition zone. *Water Resources Research* 1987; **23**:1851–1866.
18. Elder J. Transient convection in a porous medium. *Journal of Fluid Mechanics* 1967; **27**:609–623.
19. Kolditz O, Ratke R, Diersch H-J, Zielke W. Coupled groundwater flow and transport: 1. Verification of variable density flow and transport models. *Advances in Water Resources* 1997; **21**:27–46.
20. Holzbecher E. *Modeling Density-driven Flow in Porous Media*: *Principles*, *Numerics*, *Software.* Springer: Berlin, 1998.
21. Ségol G. *Classic Groundwater Simulations—Proving and Improving Numerical Models.* Prentice-Hall: Englewood Cliffs, NJ, 1994.
22. Cussler EL. *Diffusion: Mass Transfer in Fluid Systems.* Cambridge University Press: Cambridge, 1984.
23. Intraval, Phase 1, case 13, *Final Report. Technical Report*, OECD 1991.
24. Stuyfzand PJ. An accurate relatively simple calculation of the saturation index of calcite for fresh to salt water. *Journal of Hydrology* 1989; **105**:95–107.
25. Lever DA, Jackson CP. On the equations for the flow of concentrated salt solution through a porous medium. *DOE/RW/85.100*, Department of the Environment, U.K.
26. Woods JA, Barth W. Modelling brine density as a function of sodium chloride concentration. Unpublished CFDLab Presentation, University of Texas, 2003.
27. Woods JA, Carey GF. Upwelling and downwelling behaviour in the Elder–Voss–Souza benchmark. *Advances in Water Resources* 2004, in press.
28. Simmons CT, Narayan KA, Woods JA, Herczeg AL. Groundwater flow and solute transport at the Mourquong saline-water disposal basin, Murray Basin, southeastern Australia. *Hydrogeology Journal* 2002; **10**:278–295.
29. Kirk BS, Peterson JW. `libMesh`. A C++ Finite Element Library. CFDLab, 2003. URL `http://libmesh.sourceforge.net`
30. Kelly DW, Gago JPdSR, Zienkiewicz OC, Babuška I. A posteriori error analysis and adaptive processes in the finite element method. I. Error analysis. *International Journal for Numerical Methods in Engineering* 1983; **19**(11):1593–1619.
31. Cramer SD, Marchello JM. Numerical evaluation of models describing non-Newtonian behavior. *AIChE Journal* 1968; **14**(6):980–983.
32. Eyring H. Viscosity, plasticity, and diffusion as examples of absolute reaction rates. *Journal of Chemical Physics* 1936; **4**:283–287.
33. Baranger J, El Amri H. Estimateurs a posteriori d'erreurs pour le calcul adaptatif découlements quasi-Newtoniens. *RAIRO Modélisation Mathématique et Analyse Numérique* 1991; **25**:31–48.

34. Ree FH, Ree T, Eyring H. Relaxation theory of transport problems in condensed systems. *Industrial Engineering and Chemistry* 1958; **50**(7):1036–1040.
35. Oden JT, Vemaganti KS. Estimation of local modeling error and goal-oriented adaptive modeling of heterogeneous materials. I. Error estimates and adaptive algorithms. *Journal of Computational Physics* 2000; **164**(1):22–47.
36. Paterson WSB. *The Physics of Glaciers* (2nd edn). Pergamon: Oxford, 1981.
37. Seager MK, Carey GF. Adaptive domain extension and adaptive grids for unbounded spherical elliptic PDEs. *SIAM Journal on Scientific and Statistical Computing* 1990; **11**(1):92–111.
38. Colinge J, Blatter H. Stress and velocity fields in glaciers: part I. Finite difference schemes for higher-order glacier models. *Journal of Glaciology* 1998; **44**(148):448–456.
39. Bangerth W, Hartmann R, Kanschat G. `deal.II`. *Technical Reference*, Differential Equations Analysis Library, IWR, 2004. URL `http://www.dealii.org`

