

A frontal solver tuned for fully-coupled non-linear hygro-thermo-mechanical problems

Mauro Bianco¹, Gianfranco Bilardi¹, Francesco Pesavento²,
Geppino Pucci¹, and Bernhard A. Schrefler³

¹*Dipartimento di Elettronica e Informatica, Università di Padova, Padova, Italy.*
{bianco1,bilardi,geppo}@dei.unipd.it

²*Dipartimento di Costruzioni e Trasporti, Università di Padova, Padova, Italy.*
pesa@dic.unipd.it

³**Corresponding author:** *Dipartimento di Costruzioni e Trasporti, Università di Padova*
via F. Marzolo, 9 35131 - Padova, Italy. email: bas@dic.unipd.it
Tel. +39 049 827 5611, Fax. +39 049 827 5604,

SUMMARY

Solving fully-coupled non-linear hygro-thermo-mechanical problems relative to the behavior of concrete at high temperatures using monolithic models is nowadays a very interesting and challenging computational problem. These models require an extensive use of computational resources, such as main memory and computational time, due to the great number of variables and the numerical characteristics of the coefficients of the linear systems involved.

In this paper a number of different variants of a frontal solver used within HITECOSP, an application developed within the BRITE Euram III “HITECO” EU project, to solve multiphase porous media problems, are presented and evaluated with respect to their numerical accuracy and performance. When developing the variants, several optimization techniques have been adopted, such as data structure, cache and branches optimizations. Specifically numerical accuracy has been evaluated using a modified component-wise backward error analysis.

The main result of this work is a new solver which is both much faster and more accurate than the original one. Specifically, the code runs over 5 times faster and numerical errors are reduced of up to three order of magnitude. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS: frontal method, performance optimization, numerical errors analysis

1. INTRODUCTION

Many successful methods exist for the solution of algebraic equations arising from the discretization of uncoupled problems. For coupled problems, especially if they involve several fields, the problem is still open. As far as coupled thermo-hygro-mechanical problems are

Contract/grant sponsor: This work was supported, in part, by MURST of Italy within the framework of the Centre for Science and Application of Advanced Computation Paradigms of the University of Padova.

concerned, several methods have been investigated or are under investigation, but a fully satisfactory answer has not yet been found. The most reliable but also the most expensive approach is the monolithic approach, where all field equations are solved simultaneously [1]. Operator split techniques like the staggered method [2, 3, 4] allow to solve smaller size problems at each instant, but usually require iterations between the fields which reduce their advantage. Recent investigations on two field problems have shown that the LATIN method behaves rather well [5]. As far as parallel computing is concerned, domain decomposition [6, 7, 8], multifrontal techniques (single and multilevel) [6, 7, 8, 9] and asynchronous methods [10] have been investigated. Except for the asynchronous methods, which, at least up to now have only been made to work for linear two field problems, the overhead due to parallelization introduced by these methods make them less attractive than the monolithic method. Algebraic multigrid methods [11] which are currently under investigation seem however promising. Taking into account this situation, the monolithic method is still much in use and we concentrate here on this method trying to improve its performance.

First we explain why the coupled thermo-hygro-mechanical problems on which we work, are so complicated to deserve special attention.

Two different approaches are usual in the simulation of hygro-thermal phenomena and related mechanical effects in porous media. In the phenomenological approach, moisture and heat transport are described by diffusive type differential equations with temperature- and moisture content-dependent coefficients. The model coefficients are determined by inverse problem solution, i.e., using known results of experimental tests to obtain the best agreement between theoretical prediction and experimental evidence (e.g., in the sense of least square method). Thus, they are very accurate for interpolation and rather poor for extrapolation of the known experimental results. Moreover, various physical phenomena are lumped together, there is no distinction between different phases of water, hence phase changes cannot be taken into account and finally model parameters often have no clear physical interpretation.

In mechanistic models which we use [12], governing equations are usually more complicated formally, but their coefficients have clear physical meaning and often are related to classical material parameters, like for example porosity, intrinsic permeability, diffusivity of vapour in air, etc. When some relations between structure parameters and transport properties are found, usually they are valid for a class of similar materials. Mass and energy fluxes are expressed by means of gradients of thermodynamic potentials. Phase changes and mass- and energy sources (sinks) related to them are taken into account. Moreover, some additional couplings, e.g. effect of material damaging on intrinsic permeability or capillary and vapour pressures (moisture content) on skeleton stresses, can be considered.

For these reasons the second approach is considered more complete from a physical point of view, but results in a particularly complex mathematical system, in terms of both number of equations and their coefficients, which are strongly nonlinear.

We concentrate here on a particular coupled multi-physics problem which deals with concrete under high temperature conditions. Such a model allows for instance to make residual lifetime prediction in concrete vessels of nuclear reactors or to predict the behaviour of concrete walls in tunnel fires etc., [1, 13, 14]. The model has been implemented in the computer code HITECOSP in the framework of the BRITE Euram III "HITECO" [15] research project. This software uses a frontal technique to solve the final system resulting from the FE implementation of the model. The aim of our work to improve the efficiency of its frontal solver in terms of performance as well as numerical accuracy, in order to exploit the various characteristics imposed by the

mechanistic model.

Improvements in term of performance have been obtained implementing a number of optimizations (see Section 4.2), such as improving the data structures, enhancing the spatial and temporal locality of the memory accesses, and branch optimizations. With these optimizations execution times improves of a factor 4 with respect to original HITECOSP program.

In order to improve numerical accuracy in solving the linear systems arising from the Newton-Raphson like procedure used for the mechanistic model, several pivoting strategies have been implemented and evaluated based on modified component-wise backward error analysis (see Section 5) which allows to plot the magnitude of the minimal perturbation each equation sustains in order to obtain the approximated solution of the system. Actually, these perturbations can be compared with the roundoff unit error. From this analysis it follows that the best strategy in terms of accuracy is the best also in terms of performance. In particular we have noticed errors of the same order of magnitude of the roundoff unit error and a further speed-up with respect to the optimized version of the original solver.

The rest of the paper is organized as follows. In Section 2 the mathematical model is described, while Section 4 deals with the various optimizations we have introduced and the pivoting strategy we have implemented. The metrics adopted to evaluate accuracy and performance are described in Section 5. Finally, the test cases used to evaluate our solvers are described in Section 6 while our results are shown in Section 7.

2. MATHEMATICAL MODEL

In the model presented here, concrete is considered as a partially saturated porous material [1, 13] consisting of a solid phase, two gas phases and three water phases. The theoretical framework is based on the works of Whitaker [16], Bear [17], Bear and Bachmat [18, 19], Hassanizadeh and Gray [20, 21, 22] and Lewis and Schrefler [23].

Refinements such as non-linearities due to temperature and pressures, hydration-dehydration, evaporation-condensation, adsorption-desorption, phenomena are considered, (see [1, 13]). Different physical mechanisms governing the liquid and gas transport in the pores of partially saturated concrete are clearly distinguished, i.e. capillary water and gas flows driven by their pressure gradients, adsorbed water surface diffusion caused by saturation gradients, as well as air and vapour diffusion driven by vapour density gradients, [1, 13]. Concrete damaging effects arising from coupled hygro-thermal and mechanical interaction are considered by use of the isotropic non-local damage theory and a further coupling between intrinsic permeability and mechanical damage has been introduced to take into account the changes of material microstructure. More-over improvements, regarding the possibility to simulate the behaviour of the material at temperatures which largely exceed the critical point of water and the real behaviour of gases present in the pores of concrete, i.e. the gases are treated as real gas, have been recently introduced in the model, [14].

The final mathematical model consists of four balance equations: mass conservation of dry air, mass conservation of the water species (both in liquid and gaseous state, taking phase changes, i.e. evaporation-condensation, adsorption-desorption and hydration-dehydration process, into account), enthalpy conservation of the whole medium (latent heat of phase changes and heat effects of hydration or dehydration processes are considered) and

linear momentum of the multiphase system. The system is completed with an appropriate set of constitutive and state equations, as well as some thermodynamic relationships, [23]. The governing equations of the model are expressed in terms of the chosen state variables: gas pressure p^g , capillary pressure p^c , temperature T and displacement vector of the solid matrix u and are directly derived from [23], where the constitutive law for the solid skeleton density and water density are experimentally determined. In particular the density of solid skeleton has to respect the solid mass conservation equation which is not a basic equation of the model. The governing equations of the model proposed, considering negligible both the inertial forces and the convective heat flux related to solid phase and taking into account the Bishops stresses [24], are the following:

Dry air mass conservation equation:

$$\begin{aligned} & -n \frac{\partial S_w}{\partial t} - \beta_s (1-n) S_g \frac{\partial T}{\partial t} + S_g \nabla \cdot \mathbf{v}^s + \frac{S_g n}{\rho^a} \frac{\partial \rho^a}{\partial t} + \frac{1}{\rho^a} \nabla \cdot \mathbf{J}_g^a + \\ & + \frac{1}{\rho^a} \nabla \cdot (n S_g \rho^a \mathbf{v}^{gs}) + \frac{(1-n) S_g}{\rho^s} \frac{\partial \rho^s}{\partial \Gamma_{hydr}} \frac{\partial \Gamma_{hydr}}{\partial t} = \frac{\dot{m}_{dehydr}}{\rho^s} S_g \end{aligned}$$

Water species (liquid + vapour) mass conservation equation:

$$\begin{aligned} & n(\rho^w - \rho^v) \frac{\partial S_w}{\partial t} - \beta_{swg} \frac{\partial T}{\partial t} + (\rho^v S_g + \rho^w S_w) \nabla \cdot \mathbf{v}^s + S_g n \frac{\partial \rho^v}{\partial t} + \\ & + \nabla \cdot \mathbf{J}_g^v + \nabla \cdot (n S_g \rho^v \mathbf{v}^{gs}) + \nabla \cdot (n S_w \rho^w \mathbf{v}^{ws}) + \\ & + \frac{(1-n)(S_g \rho^v + \rho^w S_w)}{\rho^s} \frac{\partial \rho^s}{\partial \Gamma_{hydr}} \frac{\partial \Gamma_{hydr}}{\partial t} = \frac{\rho^v S_g + \rho^w S_w - \rho^s}{\rho^s} \dot{m}_{dehydr} \end{aligned}$$

Energy conservation equation (enthalpy balance):

$$\begin{aligned} & (\rho C_p)_{ef} \frac{\partial T}{\partial t} + (\rho_w C_p^w \mathbf{v}^{ws} + \rho_g C_p^g \mathbf{v}^{gs}) \cdot \nabla T - \\ & - \nabla \cdot (\lambda_{ef} \nabla T) = -\dot{m}_{phase} \Delta H_{phase} + \dot{m}_{dehydr} \Delta H_{dehydr} \end{aligned}$$

Linear momentum equation (equilibrium equation):

$$\operatorname{div} \{ \boldsymbol{\sigma}' - \mathbf{I} (p^g - S_w p^c) \} + \rho \mathbf{g} = 0$$

For the closure of the model a set of thermodynamic and constitutive relationships are needed. In the present work only few of them will be shown in detail, in particular as far as damage mechanics and absolute permeability-damage parameter coupling are concerned. (For further information see [23, 1, 13, 14, 24]).

3. NUMERICAL SOLUTION

The system of governing equations of the model, formed by two mass conservation equations, one linear momentum conservation equation and an enthalpy conservation equation, after application of the finite element method for discretization in space becomes:

$$\begin{aligned}
\mathbf{C}_{gg} \frac{\partial \bar{\mathbf{p}}^g}{\partial t} + \mathbf{C}_{gc} \frac{\partial \bar{\mathbf{p}}^c}{\partial t} + \mathbf{C}_{gt} \frac{\partial \bar{\mathbf{T}}}{\partial t} + \mathbf{C}_{gu} \frac{\partial \bar{\mathbf{u}}}{\partial t} + \mathbf{K}_{gg} \bar{\mathbf{p}}^g + \mathbf{K}_{gc} \bar{\mathbf{p}}^c + \mathbf{K}_{gt} \bar{\mathbf{T}} &= \mathbf{f}_g \\
\mathbf{C}_{cc} \frac{\partial \bar{\mathbf{p}}^c}{\partial t} + \mathbf{C}_{ct} \frac{\partial \bar{\mathbf{T}}}{\partial t} + \mathbf{C}_{cu} \frac{\partial \bar{\mathbf{u}}}{\partial t} + \mathbf{K}_{cg} \bar{\mathbf{p}}^g + \mathbf{K}_{cc} \bar{\mathbf{p}}^c + \mathbf{K}_{ct} \bar{\mathbf{T}} &= \mathbf{f}_c \\
\mathbf{C}_{tc} \frac{\partial \bar{\mathbf{p}}^c}{\partial t} + \mathbf{C}_{tt} \frac{\partial \bar{\mathbf{T}}}{\partial t} + \mathbf{C}_{tu} \frac{\partial \bar{\mathbf{u}}}{\partial t} + \mathbf{K}_{tg} \bar{\mathbf{p}}^g + \mathbf{K}_{tc} \bar{\mathbf{p}}^c + \mathbf{K}_{tt} \bar{\mathbf{T}} &= \mathbf{f}_t \\
\mathbf{K}_{ug} \bar{\mathbf{p}}^g + \mathbf{K}_{uc} \bar{\mathbf{p}}^c + \mathbf{K}_{ut} \bar{\mathbf{T}} + \mathbf{K}_{uu} \bar{\mathbf{u}} &= \mathbf{f}_u
\end{aligned} \tag{1}$$

As usual, the primary variables are expressed by their nodal values and the shape functions as follows:

$$\begin{aligned}
\mathbf{p}^g &= \mathbf{N}_p \bar{\mathbf{p}}^g \\
\mathbf{p}^c &= \mathbf{N}_p \bar{\mathbf{p}}^c \\
\mathbf{T} &= \mathbf{N}_t \bar{\mathbf{T}} \\
\mathbf{u} &= \mathbf{N}_u \bar{\mathbf{u}}
\end{aligned}$$

Finite differences in time are used for the solution of the initial value problem. The method used in this text produces an identical system of equations obtained with the weighted residual method in time with point collocation. We can write the system (1) in a more concise form as

$$\mathbf{B} \frac{\partial \mathbf{X}}{\partial t} + \mathbf{C} \mathbf{X} = \mathbf{F} \tag{2}$$

where the matrices B, C, X and F are the following:

$$\mathbf{C} = \begin{bmatrix} \mathbf{K}_{gg} & \mathbf{K}_{gc} & \mathbf{K}_{gt} & \mathbf{0} \\ \mathbf{K}_{cg} & \mathbf{K}_{cc} & \mathbf{K}_{ct} & \mathbf{0} \\ \mathbf{K}_{tg} & \mathbf{K}_{tc} & \mathbf{K}_{tt} & \mathbf{0} \\ \mathbf{K}_{ug} & \mathbf{K}_{uc} & \mathbf{K}_{ut} & \mathbf{K}_{uu} \end{bmatrix} ; \quad \mathbf{B} = \begin{bmatrix} \mathbf{C}_{gg} & \mathbf{C}_{gc} & \mathbf{C}_{gt} & \mathbf{C}_{gu} \\ \mathbf{0} & \mathbf{C}_{cc} & \mathbf{C}_{ct} & \mathbf{C}_{cu} \\ \mathbf{0} & \mathbf{C}_{tc} & \mathbf{C}_{tt} & \mathbf{C}_{tu} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}$$

$$\mathbf{F} = \begin{Bmatrix} \mathbf{f}_g \\ \mathbf{f}_c \\ \mathbf{f}_t \\ \mathbf{f}_u \end{Bmatrix} ; \quad \mathbf{X} = \begin{Bmatrix} \bar{\mathbf{p}}^g \\ \bar{\mathbf{p}}^c \\ \bar{\mathbf{T}} \\ \bar{\mathbf{u}} \end{Bmatrix}$$

The matrices B and C are non-symmetric and depend on X. The discretisation in the time domain is carried out by the generalised trapezoidal method, also known as the generalised midpoint rule:

$$\left(\frac{\partial \mathbf{X}}{\partial t} \right)_{n+\theta} = \frac{\mathbf{X}_{n+1} - \mathbf{X}_n}{\Delta t}$$

$$0 < \theta \leq 1$$

$$\mathbf{X}_{n+\theta} = (1 - \theta) \mathbf{X}_n + \theta \mathbf{X}_{n+1}$$

where Δt is the time step length, \mathbf{X}_n and \mathbf{X}_{n+1} are the state vectors at times t_n and t_{n+1} , and where, depending of the value of θ , it is possible to obtain different schemes for integration of the system in time:

$\theta = 1$ fully implicit scheme (Euler backward);
 $\theta = 0.5$ Crank-Nicholson scheme;
 $\theta = 0$ fully explicit scheme (Euler forward); then (2) at time $t_{n+\theta}$ becomes:

$$\begin{aligned} \Psi(\mathbf{X}_{n+1}) &= [\mathbf{B} + \theta \Delta t \mathbf{C}]_{n+\theta} \mathbf{X}_{n+1} + \\ &- [\mathbf{B} - (1 - \theta) \Delta t \mathbf{C}]_{n+\theta} \mathbf{X}_n - \Delta t \mathbf{F}_{n+\theta} = 0 \end{aligned} \quad (3)$$

Where

$$\Psi(\mathbf{X}_{n+1}) = [\Psi^g(\mathbf{X}_{n+1}), \Psi^c(\mathbf{X}_{n+1}), \Psi^t(\mathbf{X}_{n+1}), \Psi^u(\mathbf{X}_{n+1})]^T$$

the matrices in (3) are evaluated at time $t_{n+\theta}$.

An examination of the set of governing equations (3) reveals that the coupled system (2) is non-linear; hence it requires linearising by an iterative Newton-Raphson procedure of the form:

$$\Psi^\kappa(\mathbf{X}_{n+1}^i) = - \left. \frac{\partial \Psi^\kappa}{\partial \mathbf{X}} \right|_{\mathbf{X}_{n+1}^i} \Delta \mathbf{X}_{n+1}^i, \quad \kappa = g, c, t, u$$

where the Jacobian matrix is defined by:

$$\left. \frac{\partial \Psi}{\partial \mathbf{X}} \right|_{\mathbf{X}_{n+1}^i} = \left[\begin{array}{cccc} \frac{\partial \Psi^g}{\partial \bar{\mathbf{p}}^g} & \frac{\partial \Psi^g}{\partial \bar{\mathbf{p}}^c} & \frac{\partial \Psi^g}{\partial \bar{\mathbf{T}}} & \frac{\partial \Psi^g}{\partial \bar{\mathbf{u}}} \\ \frac{\partial \Psi^c}{\partial \bar{\mathbf{p}}^g} & \frac{\partial \Psi^c}{\partial \bar{\mathbf{p}}^c} & \frac{\partial \Psi^c}{\partial \bar{\mathbf{T}}} & \frac{\partial \Psi^c}{\partial \bar{\mathbf{u}}} \\ \frac{\partial \Psi^t}{\partial \bar{\mathbf{p}}^g} & \frac{\partial \Psi^t}{\partial \bar{\mathbf{p}}^c} & \frac{\partial \Psi^t}{\partial \bar{\mathbf{T}}} & \frac{\partial \Psi^t}{\partial \bar{\mathbf{u}}} \\ \frac{\partial \Psi^u}{\partial \bar{\mathbf{p}}^g} & \frac{\partial \Psi^u}{\partial \bar{\mathbf{p}}^c} & \frac{\partial \Psi^u}{\partial \bar{\mathbf{T}}} & \frac{\partial \Psi^u}{\partial \bar{\mathbf{u}}} \end{array} \right]_{\mathbf{X}=\mathbf{X}_{n+1}^i}$$

and the increment vector of the primary variables is:

$$\Delta \mathbf{X}_{n+1}^i = [(\Delta \bar{\mathbf{p}}^g)^i, (\Delta \bar{\mathbf{p}}^c)^i, \Delta \bar{\mathbf{T}}^i, \Delta \bar{\mathbf{u}}^i]^T$$

where i is the iteration count.

During the computation, for each time step, the primary variable vector \mathbf{X}_{n+1} is updated after each iteration, i.e.:

$$\mathbf{X}_{n+1}^{i+1} = \mathbf{X}_{n+1}^i + \Delta \mathbf{X}_{n+1}^i$$

In the Newton-Raphson method, the elements of the Jacobian matrix are updated after each iteration, which is computationally expensive. Thus, in practice the modified Newton-Raphson method is often preferred, where the Jacobian matrix is calculated only once at the beginning of each time step or after a fixed number of iterations.

To take into account damage of concrete, a two stage solution strategy is applied at every time step. First an intermediate problem, keeping the constant damage value obtained at the previous time step, is solved. Then, starting from this intermediate state, the final solution is obtained, for all state variables and damage parameter, by the modified Newton-Raphson method, using the tangential or Jacobian matrix from the last iteration of the first stage. Such an approach allowed us to avoid differentiation with respect to the damage and to obtain a

converging solution. Sometimes this derivative increases very rapidly until the damage reaches its maximum value and then remains equal to zero, causing divergence of the solution.

An unoptimized version of the above model has been implemented in the computer code HITECOSP, in the framework of the BRITE Euram III ‘‘HITECO’’ research project [15]. This code provides the starting base upon which we develop our optimization strategies aiming at improving both performance and numerical accuracy of the code.

4. FRONTAL METHOD: OVERVIEW AND IMPLEMENTATION

The large linear systems arising from the Newton-Raphson method described in Section (3) are solved using the *frontal method* (see [25, 9] for a full description). The frontal method solves a linear system by working, at each step, only on a portion of the matrix (called *frontal matrix*), hence it is useful in those situations where core memory becomes the critical resource. The method works by sequentially executing two phases on each element of the finite element grid: an *assembly* phase and an *elimination* phase. During the assembly phase, the frontal matrix is augmented by the appropriate number of columns and rows relative, respectively, to the variables associated to the element and the equations containing those variables, and the matrix entries are updated to account for the new element. An entry becomes *fully-summed* if it will not receive further updates in any subsequent assembly phase. A column (resp., row) becomes fully-summed when all its entries become fully summed. A variable corresponding to a fully-summed column is also said fully-summed.

During the elimination phase, Gaussian elimination is applied to the frontal matrix, choosing the pivot in the block at the intersection of fully-summed rows and fully-summed columns. At each Gaussian elimination step, the pivot row is *eliminated*, i.e., it is stored somewhere into memory (typically onto a disk, if the problem is too large to fit in main memory). After the last elimination phase, back substitution on the reduced linear system is executed.

4.1. Overview of pivoting strategies

Recall that in the frontal method only a part of the matrix of the system is available at any given time, hence any pivoting strategy must be adopted to cope with this scenario. In particular, the pivot must always be chosen among those entries of the frontal matrix which reside in fully-summed rows and fully-summed columns.

Many strategies have been developed either to speed up the frontal solver or to improve its numerical stability. In this section we describe those strategies which we have implemented in order to find a solution that achieves the best tradeoff between stability and performance for our particular physical problems.

Let A be the frontal matrix in a given elimination step. *Numerical pivoting* [9] entails choosing the pivot among the entries a_{ij} residing in fully-summed columns such that

$$|a_{ij}| \geq \alpha \max_k |a_{ik}|, \quad (4)$$

where $0 < \alpha \leq 1$ is a numerical constant and i is the index of a fully-summed row. Numerical pivoting was adopted to reduce the approximation error introduced by Gaussian elimination for the frontal method.

If an eligible pivot is not found, then the next element is assembled and a new search for a pivot is performed. This strategy is called *postordering for stability* [9]. The algorithm clearly terminates since eventually all the rows and all the columns of the frontal matrix will become fully-summed.

Often it is claimed in the literature that postordering for stability affects performance only slightly, while numerical stability is substantially increased. However, our experiments show that this is not the case for our physical problems. The failure of postordering in our scenario seems to be primarily due to the fact that the matrix entries in our mixed physical problems can be up to 40 orders of magnitude apart. Moreover, when a pivot is chosen, during the elimination step, fill-ins are produced in other rows, thus creating bonds among variables that do not belong to neighboring elements in the finite element mesh. Often these bonds prevent a row to be eliminated since the entries in its fully summed columns do not satisfy condition (4). In the next elimination steps this row continues to be filled with additional nonzero entries, hence the likelihood that it will not be chosen for elimination keeps on increasing, in a sort of “positive feedback” effect. Indeed we have observed extreme cases where rows entered in the frontal matrix at the very beginning of the solver’s activity remains in the matrix until the very end. This phenomenon introduces two problems: not only does it cause the frontal matrix to grow inordinately, slowing down the program, but also worsens the numerical stability of the method, since a row that is present for a long time in the frontal matrix will sustain many operations on it, which is likely to amplify accumulation errors.

Another popular pivoting strategy is known as *minimum degree* [26, 9]. This strategy was proposed as a greedy way to reduce fill-ins when performing Gaussian elimination and was proved to be suited for symmetric, positive-definite matrices. Under minimum degree, the pivot is chosen as the diagonal element of a row with the minimum number of entries. Under the frontal method, the minimum degree strategy may be applied to the frontal matrix, choosing the pivot on the diagonal entry of the row with minimum number of entries in the block formed by the intersection between fully summed rows and fully summed columns. Since the full matrix of our systems has a symmetric structure, choosing pivots on the diagonal also preserves this symmetry inside the frontal matrix, allowing the data structures to be simplified.

It has to be remarked that the minimum degree strategy does not make any numerical consideration on the chosen pivot and it was originally developed for matrices that do not need such numerical precautions, e.g., positive definite symmetric matrices. Although our matrices, featuring great differences between numerical values of their entries, appear to be unsuitable for an application of minimum degree pivoting, our experiments have shown that the strategy is an effective way of reducing the accumulation error caused by postordering, perhaps due to the fact that it substantially reduces the amount of floating point operations. Indeed, a careful implementation of the minimum degree strategy has proven to feature both excellent performance and numerical accuracy for our problems.

The original application, HITECOSP, from which this work started, uses the following hybrid strategy. Before the pivot is chosen, if the absolute value of the previous pivot is less than a fixed numerical threshold value (10^{-4} in our cases), then the fully-summed rows are normalized so they will contain only values included between -1 and 1 . After that, the pivot is chosen as the entry with the maximum absolute value among those in the intersection between fully-summed rows and columns. No postordering is performed. This strategy seems to work well for our physical problems. Namely, it exhibits good numerical accuracy and lends itself to an efficient implementation, which however requires a complete redesign of the relevant data

structures.

4.2. Our solutions

In this section we describe the frontal solvers which we have implemented. Each variant is characterized by a short name (in parentheses) which suggests the specific pivoting strategy adopted by the variant. Our first intervention has aimed at improving performance of the HITECOSP software (HIT) by providing an optimized implementation of its solver. In particular, the greatest improvement in performance has been obtained by the redesign of the main data structures in order to reduce at the minimum the number of linear searches inside arrays.

Another important issue that has been considered in redesigning HITECOSP's solver is the enhancement of the temporal and spatial locality of the memory accesses performed. To this purpose, extensive cache optimization has been applied, such as performing operations (e.g., row elimination, pivot search, etc.) within the solver by column rather than by row, in order to exploit the column major allocation of matrices featured by the FORTRAN compiler.

Another important source of performance enhancement has come from conditional branch optimization. As an example, consider that when computing on a sparse matrix, many operations are performed to no effect on zero entries (e.g., divisions and multiplications). However in most modern microprocessor architectures (and, in particular, on the ALPHA platform where our experiments run), keeping these operations improves performance since they take less cycles than those necessary by the processor to recover from a mispredicted branch. Indeed, conditional branch elimination in HITECOSP has improved the performance of the resulting code of up to 20%. This version of the frontal solver, implementing the same pivoting strategy as HITECOSP, named BASE, exhibits rather good performance.

The above code optimizations have also been employed to speed up the execution of the other solvers developed within our study. However, the main justification for designing new solvers mainly stems from our desire to compare the efficiency and numerical stability of the pivoting approach of HITECOSP with the other more established strategies described in the previous section. As a first step, basic numerical pivoting (as illustrated before) was implemented. Specifically we have developed a version without postordering that chooses as pivot the element that maximizes the value of α in (4) (NUMPIV), and another one implementing postordering (NUMPP0) which set $\alpha = 10^{-6}$ in (4).

Next, we have implemented the minimum degree strategy (MINDEG). This latter solver chooses the pivot on the diagonal of the frontal matrix and is endowed with recovery features when an entry in the diagonal is zero (however, this has never occurred in our experiments). A further optimization stems from the fact that, since the structure of the frontal matrix depends only on that of the finite element mesh, and, under the minimum degree strategy, the pivots depend only on the structure of the matrix, the pivotal sequence remains the same for all the executions of the frontal solver over the different iterations of the Newton-Raphson method (unless the chosen pivot is zero, which requires special care). Hence, it is possible to store the pivotal sequence during the first execution and to use it for the next ones, choosing the pivot, at each stage, with a single memory access. Our version of the minimum degree algorithm stores the pivotal sequence after the first execution of the frontal solver and uses it in the subsequent calls.

Finally, for the purpose of comparison, we have produced a further implementation (HSL)

which uses the free version of the HSL (Harvell Subroutine Library) library [27]. Specifically, we have used the MA32 routine [28] that implements a frontal solver featuring a sophisticated numerical pivoting strategy. Version HSL gives us insight to compare our strategies against standard solutions available using third-party software.

5. COMPARISON METRICS

All the solver versions introduced in the previous section are evaluated in terms of their numerical stability (limited to the solution of the linear system) and performance. The next two sections discuss the metrics used to measure these two characteristics.

Evaluating numerical stability.

Let \tilde{x} be the computed (approximated) solution to one of the linear systems $Ax = b$ arising during the solution of a FEM problem. Approximation errors have been evaluated using two different metrics. The first is the two-norm of the residual r , that is,

$$\|r\| = \|b - A\tilde{x}\| \quad (5)$$

This quantity gives an indirect indication of the distance between the real solution and the approximated one.

The second metric adopted is the *component-wise backward error* [29]

$$w = \min\{\varepsilon : (A + \Delta A)\tilde{x} = b + \Delta b, |\Delta A| \leq \varepsilon|A|, |\Delta b| \leq \varepsilon|b|\} \quad (6)$$

(The absolute values and the comparisons are intended to be component-wise). It can be proved (see [29]) that, setting $0/0 = 0$ and $\xi/0 = \infty$ if $\xi \neq 0$, w can be computed as

$$w = \max_i \frac{|r_i|}{(|A| |\tilde{x}| + |b|)_i} \quad (7)$$

There is evidence in the literature that the component-wise backward error is more sensitive to instability than the residual-based error. Intuitively, w measures the minimum variation that the matrix of the system and the right hand side vector should sustain to obtain the approximated \tilde{x} solution.

In this work, in order to have a more detailed description of the approximation errors introduced by our solvers, we have refined the component-wise backward error metric as follows. Let a_i denote the i -th row of A . Define the *ith equation error* to be

$$w_i = \min\{\varepsilon : (a_i + \Delta a_i)\tilde{x} = b_i + \Delta b_i, |\Delta a_i| \leq \varepsilon|a_i|, |\Delta b_i| \leq \varepsilon|b_i|\}. \quad (8)$$

Value w_i gives a measure of “how well” vector \tilde{x} satisfies the i th equation of the system. We have:

Theorem 1. *Let v be a vector with $v_i = \frac{r_i}{(|A| |\tilde{x}| + |b|)_i}$. Then $w_i = v_i$*

Proof: Let $\Delta \hat{a}_i$ and $\Delta \hat{b}_i$ be a pair of minimal perturbations associated to w_i . Since $|\Delta \hat{a}_i| \leq w_i|a_i|$ and $|\Delta \hat{b}_i| \leq w_i|b_i|$, we have that

$$|r_i| = |b_i - a_i\tilde{x}| = |\Delta \hat{a}_i\tilde{x} - \Delta \hat{b}_i| \leq |\Delta \hat{a}_i||\tilde{x}| + |\Delta \hat{b}_i| \leq w_i(|a_i||\tilde{x}| + |b_i|),$$

whence $w_i \geq |v_i|$. Also, from the definition of v it follows that $r_i = v_i(|a_i|\tilde{x} + |b_i|)$. Define now $\Delta a'_i = v_i|a_i|\text{diag}(\text{sign}(\tilde{x}))$ and $\Delta b'_i = -v_i|b_i|$. It is easy to see that $(a_i + \Delta a'_i)\tilde{x} - (b_i + \Delta b'_i) = 0$. Therefore, since $|\Delta a'_i| = |v_i||a_i|$ and $|\Delta b'_i| = |v_i||b_i|$, it follows that $w_i \leq |v_i|$, and the theorem follows. \square

It is easy to see that $\Delta \hat{A} = \text{diag}(v)|A|\text{diag}(\text{sign}(\tilde{x}))$ and $\Delta \hat{b} = -\text{diag}(v)|b|$ are such that $|\Delta \hat{A}| \leq w|A|$, $|\Delta \hat{b}| \leq w|b|$ and $(A + \Delta \hat{A})\tilde{x} - (b + \Delta \hat{b}) = 0$. Moreover, the above theorem proves that all the perturbations are the minimum possible, in the sense indicated by (8). Hence, vector v provides a readily obtainable indication about the minimum perturbation that each equation should sustain to obtain the approximated solution. In particular any element of v can be compared against the roundoff unit error to gain immediate appreciation of the significance of the corresponding perturbations. Finally, observe that the standard component-wise error metric w can be obtained as $w = \|v\|_\infty$.

A plot of vector v (using the equation indices as the abscissae) can be used to ascertain whether numerical errors tend to affect some groups of equations more than others. We feel that this is particularly useful in multi-physics applications as the ones treated in this paper.

All the metrics described above have been collected over several iterations of each solver. No significant variation of each metric has been observed over the different iterations. However, in what follows, we report the maximum errors encountered on each test case.

Measuring performances.

Performance is measured both in terms of computational time and rate of floating point operations (Mflops) relatively to the frontal solver only. While computational time is a direct and well understood measure of performance, the rate of floating point operations is a more technical one. A low rate implies that the solver spends most of the time in managing data structures, in idle time due to data dependencies and outstanding accesses to main memory, etc., but not in producing effective data for the output. Hence, the former measure provides a direct indication of the efficiency of the chosen pivoting strategy, while the latter gives an insight on how efficiently the code is exploiting the underlying architecture.

6. TEST CASES

The various solver versions have been tested on a number of test cases arising in several practical scenarios and characterized by an increasing complexity of the underlying physical system. As for the solver version, each test case is indicated by a short name (in parentheses):

1. *small column (smcol)*: a regular 10×10 mesh of 100 elements in which all the degrees of freedom, except for the ones related to displacements, are set to zero.
2. *wall (wall)*: 69 elements lined up in a row where the fifth degree of freedom (y -displacement) is fixed to zero;
3. *container (cont)*: 288 elements outlining a container;
4. *column (col)*: a square section of a column made of a 20×20 mesh of 400 elements;
5. *big column (bigcol)*: like *column* but made of a 25×25 mesh of 625 elements;

Such a variety of test cases allows us to evaluate the behavior of the solver variants when the complexity of the physics behind the problem to be solved varies, from simpler (*smcol*) to

harder (*bigcol*).

7. RESULTS

The solver versions shown above have been tested on an Alpha workstation which uses a 21264 processor clocked at 666Mhz with two floating point units that make it capable of a 1354Mflops peak performance

The next two subsections report the results obtained by evaluating the various versions of the solvers described above. In section 7.1 we analyze the numerical stability properties exhibited by the solvers with respect to the metrics discussed in section 5. In section 7.2 we examine the performance achieved by the solvers.

		HIT	BASE	NEWPIV	NEWPPO	MINDEG	HSL
<i>smcol</i>	$\ r\ $	$1 \cdot 10^{-12}$	$1 \cdot 10^{-12}$	$1 \cdot 10^{-12}$	$1 \cdot 10^{-12}$	$1 \cdot 10^{-12}$	$1 \cdot 10^{-12}$
	w	$3 \cdot 10^{-16}$	$3 \cdot 10^{-16}$	$3 \cdot 10^{-16}$	$3 \cdot 10^{-16}$	$3 \cdot 10^{-16}$	$4 \cdot 10^{-16}$
<i>wall</i>	$\ r\ $	$5 \cdot 10^{-14}$	$1 \cdot 10^{-13}$	$7 \cdot 10^{-12}$	$4 \cdot 10^{-8}$	$6 \cdot 10^{-14}$	$7 \cdot 10^{-14}$
	w	$4 \cdot 10^{-14}$	$1 \cdot 10^{-13}$	$1 \cdot 10^{-11}$	$3 \cdot 10^{-6}$	$4 \cdot 10^{-16}$	$1 \cdot 10^{-6}$
<i>cont</i>	$\ r\ $	$6 \cdot 10^{-14}$	$8 \cdot 10^{-14}$	$3 \cdot 10^{-12}$	$1 \cdot 10^{-3}$	$6 \cdot 10^{-14}$	$1 \cdot 10^{-12}$
	w	$6 \cdot 10^{-13}$	$5 \cdot 10^{-13}$	$6 \cdot 10^{-11}$	$2 \cdot 10^{-4}$	$6 \cdot 10^{-16}$	$5 \cdot 10^{-3}$
<i>col</i>	$\ r\ $	$1 \cdot 10^{-14}$	$2 \cdot 10^{-14}$	$2 \cdot 10^{-11}$	$3 \cdot 10^{-5}$	$1 \cdot 10^{-14}$	$1 \cdot 10^{-13}$
	w	$4 \cdot 10^{-12}$	$2 \cdot 10^{-12}$	$2 \cdot 10^{-5}$	$2 \cdot 10^{-3}$	$9 \cdot 10^{-16}$	$4 \cdot 10^{-4}$
<i>bigcol</i>	$\ r\ $	$1 \cdot 10^{-13}$	$1 \cdot 10^{-13}$	$1 \cdot 10^{-6}$	$4 \cdot 10^{-5}$	$7 \cdot 10^{-14}$	$6 \cdot 10^{-13}$
	w	$2 \cdot 10^{-12}$	$6 \cdot 10^{-13}$	$7 \cdot 10^{-5}$	$5 \cdot 10^{-3}$	$1 \cdot 10^{-15}$	$5 \cdot 10^{-5}$

Table I. Residual errors and component-wise backward errors exhibited by the various solvers for each test cases.

Table I reports error metrics (5) and (6) for all the solver versions and test cases described before.

7.1. Numerical Quality

We note that MINDEG exhibits the least errors (in order of magnitude), scaling extremely well as the physical problems become more complex. Also, the table shows that HIT and BASE do not exhibit exactly the same errors, with BASE featuring slightly larger errors. This is explained by the fact that HIT uses some extra heuristic precautions to reduce fill-ins. We have chosen not to implement these expedients in BASE, since they complicate the code while not providing significant improvements in term of either accuracy or performance.

If we compare NEWPIV with NEWPPO, implementing, respectively, numerical pivoting without and with postordering, we see that the errors exhibited by the latter (in all cases except *smcol*) are orders of magnitude worse than those exhibited by the former. Indeed, for the two largest test cases, the errors exhibited by NEWPPO in the solution of the linear systems became so large to prevent the Newton-Raphson method from converge. This provides numerical evidence that postordering does not achieve its intended purpose in our physical scenario.

To achieve a more profound understanding on the numerical behavior of the implemented variants, in Figure 1, we show all the components of the v vector (see sec. 5) for different solvers running the *bigcol* test case, plotted (in logarithmic scale) against their respective indices. For

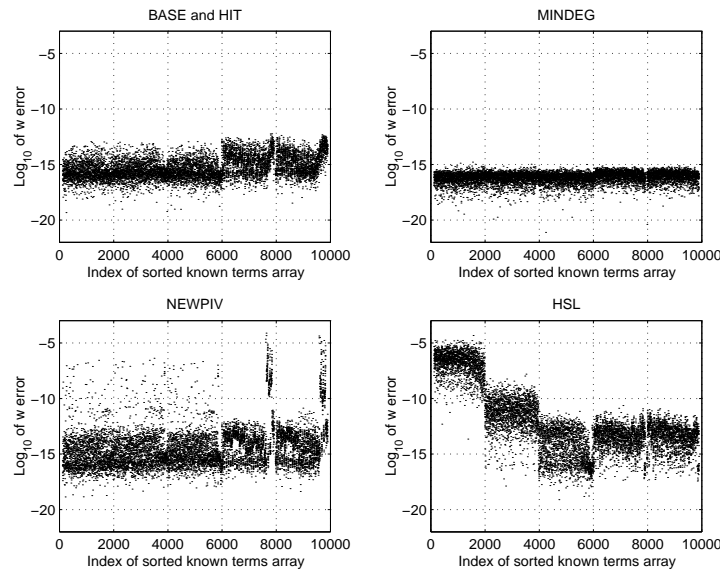


Figure 1. Plot of the component-wise errors for the *bigcol* test case, exhibited by the various solvers. The abscissae are the equations indices and the ordinates are the base 10 logarithms of the errors.

our specific physical problems, we have that indices between 1 and 2000 are relative gas pressure equations, indices between 2001 and 4000 are relative to capillary pressure, indices between 4001 and 6000 are relative to temperatures, and, finally, the remaining indices are relative to displacement equations. It is interesting to observe how errors tend to cluster according to the type of equation.

It is clear from the figure that MINDEG exhibits extremely homogeneous errors that are all close to the roundoff unit error (which is about 10^{-16} for our machine), while BASE, although still behaving quite well, tends to show a more varied range of errors, which implies that different equations are solved with different degrees of accuracy. As for NEWPIV, we can see a rather good behavior on average, but the plot highlights a nonnegligible set of outliers with high values of the error (which, in turn, yields a high value of the metric w). Finally, looking at the plot for HSL, we may note how it behaves very poorly for two entire groups of equations (especially those related to pressures) while it is only slightly worse than BASE for the other two groups.

7.2. Performances

Table II shows the execution times and the floating point operation rates (Mflops) exhibited by the frontal solvers. Each test case involves several time steps, with each time step in turn requiring the solution of a number of linear systems, one per Newton-Raphson iteration. More specifically, the number of linear systems solved in each test case is about 120 for *wall* and *smcol*, 45 for *cont* and *col*, and 30 for *bigcol*. For the last three test cases, the most computationally intensive ones, we have that the solver accounts for about 75% of the total time for MINDEG, and goes up to 95% for HIT. This fact justifies our focusing on the solver only, rather than

		HIT	BASE	NEWPIV	NEWPPO	MINDEG	HSL
<i>wall</i>	Time (s)	4.6	2.23	2.48	1.96	1.07	4.48
	Mflops	59.54	109.18	134.15	163.08	161.34	231.73
<i>smcol</i>	Time (s)	6.05	1.98	2.12	1.83	1.64	6.97
	Mflops	50.28	164.3	168.94	179.05	180.45	201.91
<i>cont</i>	Time (s)	77.5	27.5	30.4	790.0	22.9	207.1
	Mflops	79.63	237.63	232.06	98.4	262.02	206.47
<i>col</i>	Time (s)	314.1	116.8	124.5	3212	96.5	1230.5
	Mflops	54.08	224.09	222.62	7.89	256.93	154.76
<i>bigcol</i>	Time (s)	1276.4	260.1	274.0	$+\infty$	225.9	1609.5
	Mflops	30.58	232.53	226.6	–	261.77	137.47

Table II. Times and floating point operations rates (frontal solver only). Note the $+\infty$ entries in the NEWPPO columns are related to that test cases for which the method does not converge as the errors become too large.




Figure 2. Time taken by the various solvers on a single linear system as a function of the size of the system itself.

other parts of the program, such as the numerical integration routines.

In figure 2 we plot the time taken by the various solvers on a single linear system as a function the size of the system itself. We want to remark that the size of the system is not, however, the only parameter that affects performance, since, depending on the solver used, other peculiarities may influence execution times, such as, for example, the shape of the mesh.

Looking at Table II, we note that MINDEG exhibits the by far best performance. Together with the numerical stability data presented in the previous section, this implies that MINDEG achieves both performance and accuracy at the same time, which is somewhat surprising since an increase in accuracy often comes at the expense of a deterioration of performance.




Figure 3. Floating point operations rates exhibited by the solvers varying the size of the systems resolved.

Unlike HIT and HSL, MINDEG, BASE, and NEWPIV seem to be able to sustain high flops rate when the size of the problem increases (see fig. 3). Such scalability rewards our redesign of the data structures, which affords access times to data which are independent of the amount of the data itself. Comparing BASE against HIT, we see that redesigning the data structures, optimizing data accesses and carefully eliminating conditional branches alone made the solver more than 4 times faster. Changing the pivoting strategy yielded an extra time saving: indeed, MINDEG is more than 5.6 times faster than HIT for the *bigcol* test case.

Looking at the numerical pivoting strategies, we note that HSL exhibits a rather unsatisfactory performance. This seems to be mainly due to postordering. Comparing NEWPPO to NEWPIV, for the test cases for which both strategies converge, we note that when postordering is extensively used (as for *cont* and *col* test cases), the execution time explodes. This can be explained by the feedback effect described in section 4.1 and due to postordering. It has to be remarked that HSL behaves quite well for the *wall* and *smcol* test cases, for which it exhibits the highest flops rate. In fact, when the physical problem is simple, HSL becomes competitive. However, as the problem becomes more complex, the time performance of HSL degrades, even though its floating point operations rate remains quite high. As a bottomline we can say that HSL features a very good implementation (high flops rate), but its pivoting strategy, however, turns out to be a poor choice for our physical problems.

Going back to NEWPIV and NEWPPO, we note that for the *wall* and *smcol* test cases, the latter exhibits better execution times than the former. This is due to the fact that, while NEWPIV maximizes α in equation (4), NEWPPO simply picks the first element which satisfies (4) for a fixed $\alpha = 10^{-6}$. This strategy proves to be beneficial for performance since fewer entries of the frontal matrix need to be scanned. The gain in performance is however limited to those simple cases where postordering is rarely applied. We have chosen not to pick the best possible pivot when implementing postordering since we have observed that otherwise some rows would

remain longer in the frontal matrix, which has detrimental effects on both time performance and accuracy.

8. CONCLUSION

When solving very non linear and strictly coupled physical problems where there may be many orders of magnitude among the numerical values involved, our experiments suggest that the best strategy is to strive for simplicity. Indeed, the MINDEG version of the solver does not do any numerical consideration when choosing the pivot in Gaussian elimination, but only structural ones. Yet, this suffices to get excellent performance and good accuracy.

Redesigning the data structures and performing code optimizations has proved to be the most effective way to speed-up the program considering that BASE achieves an improvement of a factor 4 with respect to HIT. A further improvement is then obtained by simplifying the pivoting strategy.

A possible further improvement in performance, that we mean to investigate, is to find the right tradeoff between avoiding linear searches inside the arrays and limiting indirect addressing. Specifically, observe that data structures designed to avoid linear searches make large use of indirect addressing, which, however, may disrupt temporal locality and slow down the algorithm by forcing the processor to wait for the data to become available from main memory.

A further way to improve performance is parallelizing the code. This implies choosing between a number of architecture-dependent alternatives, such as adopting a full *multi-frontal* [9, 30] versus a *multiple fronts* strategy [31], and a well suited partitioning algorithm depending on the of problems that have to be resolved.

REFERENCES

1. Gawin D, Majorana CE, Pesavento F, Schrefler BA. A fully coupled multiphase FE model of hygro-thermo- mechanical behaviour of concrete at high temperature. In Computational Mechanics., Onate, E. & Idelsohn, S.R. (eds.), New Trends and Applications, Proc. of the 4th World Congress on Computational Mechanics, Buenos Aires 1998; 1–19. Barcelona: CIMNE, 1998.
2. Park KC, Felippa CA. Partitioned analysis of coupled system, in Computational Methods for Transient Analysis, T. Belytschko and T.R.J. Hughes (eds), Elsevier, Amsterdam, 1983; 157-219
3. Turska E, Wisniewski K, Schrefler BA. Error propagation of staggered solution procedures for transient problems, Computer Methods in Appl. Mech. and Engng., 1994; 144:177–188.
4. Turska E, Schrefler BA. On convergence conditions of partitioned solution procedures for consolidations problems, Computer Methods in Applied Mech. and Engng., 1993; 106:51–63.
5. Dureisseix D, Ladeveze P, Schrefler BA. LATIN strategy for coupled fluid-solid problems in the domain, from Computational Fluid and Solid Mechanics, Vol.2 (K.J. Bathe ed.) Elsevier, 2001; 1143–1146.
6. Wang X, Gawin D, Schrefler BA. A parallel algorithm for thermo-hydro-mechanical analysis of deforming porous media, Computational Mechanics, 1996; 19:94–104.
7. Wang X, Schrefler BA. A multifrontal parallel algorithm for coupled thermo-hydro-mechanical analysis of deforming porous media, Int. J. Num. Meth. Engng., 1998; 43:1069–1083
8. Wang X, Baggio P, Schrefler BA. A multi-level frontal algorithm for finite element analysis and its implementation on parallel computation, Engineering Computations, 1999; **16**(4):406–427.
9. Duff IS, Erisman AM, Reid JK. *Direct Methods for Sparse Matrices*, Clarendon Press, 1986.
10. Wang X, Matteazzi R, Schrefler BA, Zienkiewicz OC. An asynchronous parallel algorithm for consolidation problems, chapter 3 from ‘Modeling in Geodynamics’, M. Zaman, J.R. Booker, G. Gioda, eds., Wiley, Chichester, 2000; 51–64.

11. Wang X, Schrefler BA. Fully coupled thermo-hydro-mechanics analysis by an algebraic multigrid method, submitted for publication
12. Schrefler BA. Mechanics and thermodynamics of saturated-unsaturated porous materials and quantitative solutions, submitted for publication
13. Gawin D, Majorana CE, Schrefler BA. Numerical analysis of hygro-thermic behaviour and damage of concrete at high temperature. In *Mech. Cohes.-Frict. Mater.* 1999; 4:37-74.
14. Gawin D, Pesavento F, Schrefler BA. Modelling of hygro-thermal behaviour and damage of concrete at temperature above the critical point of water. Submitted for publication.
15. BRITE Euram III BRPR-CT95-0065 1999 "HITECO". Understanding and industrial applications of High Performance Concrete in High Temperature Environment, Final Report, 1999.
16. Whitaker S. Simultaneous heat mass and momentum transfer in porous media: a theory of drying. In *Advances in heat transfer* 13. New York: Accademic Press, 1977.
17. Bear J. *Dynamics of Fluids in Porous Media*. New York: Dover, 1988.
18. Bear J, Bachmat Y. Macroscopic modelling of transport phenomena in porous media, 2: applications to mass momentum and energy transfer. In *Transp. in Porous Media*, 1986; 1:241-269.
19. Bear J, Bachmat Y. *Introduction to Modelling of Transport Phenomena in Porous Media*. Dordrecht: Kluwer, 1990.
20. Hassanizadeh M, Gray WG. General conservation equations for multiphase systems: 1. Averaging technique. In *Adv. Water Res.* 1979; 2:131-144.
21. Hassanizadeh M, Gray WG. General conservation equations for multiphase systems: 2 Mass, momenta, energy and entropy equations. In *Adv. Water Res.* 1979; 2:191-203.
22. Hassanizadeh M, Gray WG. General conservation equations for multiphase systems: 3 Constitutive Theory for Porous Media. In *Adv. Water Res.* 1980; 3:25-40.
23. Lewis RW, Schrefler BA. *The Finite Element Method in the Static and Dynamic Deformation and Consolidation of Porous Media*. Chichester: Wiley & Sons, 1998.
24. Schrefler BA, Gawin D. The effective stress principle: incremental or finite form? *Int. J. for Num. and Anal. Meth. in Geomechanics*, 1996; **20**(11):785-815.
25. Irons BM. A frontal solution program for finite element analysis *Int. J. Numer. Meth. Engng*, 1970; 2:5-32.
26. Tinney WF, Walker JW. Direct solutions of sparse network equations by optimally ordered triangular factorization, *Proc. IEEE* 55, 1967; 1801-1809,
27. HSL (Formerly the Harwell Subroutine Library). <http://www.cse.clrc.ac.uk/Activity/HSL>
28. Duff IS, Reid JK. MA32 - a package for solving sparse unsymmetric systems using the frontal method, *Report R.10079*, HMSO, London, 1981.
29. Oettli W, Prager W. "Compatibility of Approximate Solution of Linear Equations with Given Error Bounds for Coefficients and Right-Hand Sides", *Numerische Mathematik*, 1964; 6:405-409.
30. Duff IS, Reid JK. The multifrontal solution of unsymmetric sets of linear equations *SIAM J. Sci. Stat. Comput.*, September 1984; **5**(3):633-641,
31. Duff IS, Scott JA. The use of multiple fronts in Gaussian Elimination, *Technical Report RAL-TR-94-040*, Department for Computation and Information, Rutherford Appleton Laboratory, September 1994.