



HAL
open science

A parallel solver for incompressible fluid flows

Yushan Wang, Marc Baboulin, Jack Dongarra, Joel Falcou, Yann Fraigneau,
Olivier Le Maitre

► **To cite this version:**

Yushan Wang, Marc Baboulin, Jack Dongarra, Joel Falcou, Yann Fraigneau, et al.. A parallel solver for incompressible fluid flows. [Research Report] RR-8247, INRIA. 2013, pp.16. hal-00790144

HAL Id: hal-00790144

<https://inria.hal.science/hal-00790144>

Submitted on 19 Feb 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A parallel solver for incompressible fluid flows

Yushan Wang, Marc Baboulin, Jack Dongarra, Joël Falcou, Yann Fraigneau, Olivier Le Maître

**RESEARCH
REPORT**

N° 8247

Février 2013

Project-Team Grand-Large



A parallel solver for incompressible fluid flows

Yushan Wang^{*}, Marc Baboulin[†], Jack Dongarra[‡], Joël Falcou[§],
Yann Fraigneau[¶], Olivier Le Maître^{||}

Project-Team Grand-Large

Research Report n° 8247 — Février 2013 — 16 pages

Abstract: The Navier-Stokes equations describe a large class of fluid flows but are difficult to solve analytically because of their nonlinearity. We present in this paper a parallel solver for the 3-D Navier-Stokes equations of incompressible unsteady flows with constant coefficients, discretized by the finite difference method. We apply a prediction-projection method that transforms the Navier-Stokes equations into three Helmholtz equations and one Poisson equation. For each Helmholtz system, we apply the Alternating Direction Implicit (ADI) method resulting in three tridiagonal systems. The Poisson equation is solved using partial diagonalization which transforms the Laplacian operator into a tridiagonal one. We present an implementation based on MPI where the computations are performed on each subdomain and information is exchanged at the interfaces between subdomains. We describe in particular how the solution of tridiagonal systems can be accelerated using vectorization techniques.

Key-words: Navier-Stokes equations, prediction-projection method, ADI method, partial diagonalization, parallel computing, tridiagonal systems, SIMD extension

^{*} LRI, Université Paris-Sud, France (yushan.wang@lri.fr).

[†] Inria Saclay and Université Paris-Sud, France (marc.baboulin@inria.fr).

[‡] University of Tennessee and Oak Ridge National Laboratory, USA, University of Manchester, United Kingdom (dongarra@eecs.utk.edu).

[§] LRI, Université Paris-Sud, France (joel.falcou@lri.fr).

[¶] LIMSI-CNRS, France (yann.fraigneau@limsi.fr).

^{||} LIMSI-CNRS, France (olm@limsi.fr).

**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

1 rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing
Campus de l'École Polytechnique
91120 Palaiseau

Un solveur parallèle pour les écoulements de fluides incompressibles

Résumé : Les équations de Navier-Stokes décrivent une large classe d'écoulements de fluides mais sont difficiles à résoudre analytiquement à cause de leur non-linéarité. Nous présentons dans cet article un solveur parallèle pour les équations de Navier-Stokes 3-D pour des fluides incompressibles turbulents avec coefficients constants, discrétisés par la méthode des différences finies. Nous appliquons une méthode de prédiction-projection qui transforme les équations de Navier-Stokes en trois équations d'Helmholtz et une équation de Poisson. Pour chaque système d'Helmholtz, nous appliquons la méthode ADI (Alternating Direction Implicit) pour obtenir trois systèmes tridiagonaux. L'équation de Poisson est résolue à partir d'une diagonalisation partielle qui transforme l'opérateur Laplacien en un système triadiagonal. Nous présentons une implémentation basée sur MPI où les calculs sont effectués sur chaque sous-domaine et les informations sont échangées aux interfaces. Nous décrivons en particulier comment la résolution des systèmes tridiagonaux peut être accélérée en utilisant des techniques de vectorisation.

Mots-clés : Equations de Navier-Stokes, méthode de prédiction-projection, méthode ADI, diagonalisation partielle, calcul parallèle, systèmes tridiagonaux, extension SIMD

1 Introduction

The incompressible Navier-Stokes (NS) equations express the Newton's second law applied to fluid motion, with the assumptions that the fluid density is constant and the fluid stress is the sum of a viscous term (proportional to the gradient of the velocity) and an isotropic pressure term. These equations are widely used to model fluid flows in numerous scientific and industrial applications (*e.g.* in aeronautics, aerospace, meteorology, thermo-hydraulics, and climatology). Efficient numerical methods for the NS equations have then received considerable attention over the last decades, as their simulation remains a challenging problem especially for flows exhibiting complex transient, turbulent dynamics and features at many different scales. Alternative numerical methods for the NS equations include *finite difference* (FD) [4, 15, 19, 5], *finite element* [29, 31], *finite volume* [4, 15], *spectral methods* [11, 33] and mesh-free methods such as the *vortex method* [13] or *smoothed particle hydrodynamics* [27, 25, 26]. These methods have pros and cons, depending on the problem considered. In addition, the evolution of computer architectures and the emergence of massively parallel computers call for the development of new algorithms to fully exploit hardware advances.

In this paper, we focus on FD discretizations (of second order in time and space) of the three-dimensional NS equations which are solved using the classical prediction-projection method [12, 23, 18, 10]. This method relies on an explicit treatment of the non-linear terms and an implicit treatment of the viscous and pressure terms to ensure stability. As detailed in Section 2, it leads to the solution of 3-D diffusion (velocity) and Poisson (pressure) equations that constitute the main computational workload of the simulation. Further, we restrict ourselves to spatial discretizations on structured cartesian grids resulting from the tensorization of one-dimensional grid. Although not general, the structured character of these grids greatly facilitates the design and efficient implementation of the algorithms on modern computer architectures, and encompasses many situations of interest. Even for these simple meshes, 3-D parallel solvers are needed since complex flow simulations and targeted applications commonly require grid resolution corresponding to $10^7 - 10^9$ unknowns to properly capture the flow dynamics. These large simulations require the use of parallel computers. Another important advantage brought by the prediction-projection approach on tensored grids is that it leads to Helmholtz-like or Poisson equations that can be consistently reduced to one-dimensional operators using respectively operator splitting and partial diagonalization [2]. The second order discretization results in tridiagonal systems that can be solved efficiently by taking advantage of current parallel architectures. The parallel strategy is based on the idea of dividing the computational work between several processors which perform their calculations concurrently. We use a domain decomposition method in which the computational domain is divided into subdomains, each associated to one processor. The Poisson and Helmholtz-like equations are then solved over each subdomain in parallel using the Schur complement method to ensure the continuity of the solution at the subdomains interfaces. This involves data communication between neighboring subdomains for the interface solution. Communications are handled by the *Message-Passing-Interface* (MPI) [16]. We also rely on the BLAS [1], LAPACK [3] and ScaLAPACK [9] kernels from the MKL [21] library to perform the computations. In addition to the combination of the numerical methods mentioned above, an important contribution of this work is the implementation of a vectorized version of Thomas algorithm for a fast solution of the tridiagonal systems.

This paper is organized as follows. Section 2 summarizes the solution of the Navier-Stokes equations for incompressible fluid flows and the prediction-projection method. Then, the FD discretization in space and time is introduced together with the directional splitting and partial diagonalization to obtain the tridiagonal systems to be solved. Section 3 discusses the parallel implementation and emphasizes on vectorization techniques to improve the Thomas algorithm

for the solution of the tridiagonal systems. In Section 4, several performance measurements of our Navier-Stokes solver are provided using a current parallel system. Concluding remarks are presented in Section 5.

2 Navier-Stokes Problem

We consider the unsteady flow of an incompressible Newtonian fluid governed by the Navier–Stokes (NS) equations that can be expressed as

$$\frac{\partial \mathbf{V}}{\partial t} + \nabla \cdot (\mathbf{V}^T \otimes \mathbf{V}) = -\nabla P + \frac{1}{\text{Re}} \Delta \mathbf{V}, \quad (1)$$

$$\nabla \cdot \mathbf{V} = 0, \quad (2)$$

where $\mathbf{V} = (u(x, y, z, t), v(x, y, z, t), w(x, y, z, t))^T$ is the velocity vector and $P = P(x, y, z, t)$ is the pressure. Eqs. (1) and (2) state the conservation of momentum and mass respectively. We have denoted ∇ the nabla and Δ the Laplacian operators, while \otimes is the Kronecker product [17]. In the following, we use the shorthand notation $\mathcal{CT} := \nabla \cdot (\mathbf{V}^T \otimes \mathbf{V})$ for the convective term. In Eq. (1), $\text{Re} := \rho UL/\mu$ is the Reynolds number based on a reference length L , a reference velocity U , the fluid density ρ , and the dynamic viscosity μ . The Reynolds number is a good indicator of the fluid state (laminar, transient, turbulent). We see from Eq. (1) that the larger Re is, the more dominant the convection is, with respect to viscous effects. This leads to the onset of flow instabilities and a turbulent state of the flow for high Reynolds numbers. In turbulent flows, the order of magnitude for the ratio between the largest and the smallest length scale of eddy structures is $\text{Re}^{\frac{3}{4}}$. So, a finer discretization is required when the Reynolds number increases and the resulting order of magnitude for the 3D mesh size will be $\text{Re}^{\frac{3}{4}}$. Solving the NS equations also involves defining the initial and boundary conditions. The details about the implementations of the boundary condition will not be discussed in this paper. We mention that our code can deal with mixed Dirichlet and Neumann (homogeneous or not), and periodic/symmetric boundary conditions. The initial conditions for velocity and pressure will depend on the physical problem.

2.1 Prediction-Projection method

In this section, we present the prediction-projection method applied on the NS equations. The spatial discretization uses second-order finite differences approximations of the differential operator. The domain is discretized using $N_x \times N_y \times N_z$ cells, in the x , y and z directions. The discrete velocity and pressure unknowns are defined along a staggered arrangement, in which the pressure is defined on the center of each cell, the velocity components are defined on the center of the corresponding edge, see Fig. 1, to prevent spurious oscillations [24].

We denote the time level with upper-indices and Δt the time-step size. We rely on a second order approximation of the time derivative (BDF2) [6, 20], an implicit treatment of the viscous and pressure term, and a second order explicit extrapolation of the convective term.

Then the discrete NS equations can be expressed as

$$\frac{3\mathbf{V}^{n+1} - 4\mathbf{V}^n + \mathbf{V}^{n-1}}{2\Delta t} + \widetilde{\mathcal{CT}}^{n+1} = -\nabla P^{n+1} + \frac{1}{\text{Re}} \Delta \mathbf{V}^{n+1}, \quad (3)$$

$$\nabla \cdot \mathbf{V}^{n+1} = 0, \quad (4)$$




Figure 1: Staggered mesh structure.

where $\widetilde{\mathcal{C}\mathcal{T}}^{n+1} = \mathcal{C}\mathcal{T}^n - \mathcal{C}\mathcal{T}^{n-1}$.

The integration for one time-step consists of two steps: a prediction followed by a projection. The latter step is also called the correction step.

2.1.1 Prediction step

In the prediction step, an approximate solution for the momentum equation is sought without accounting for the incompressibility condition in Eq. (4). More specifically, we compute an auxiliary velocity \mathbf{V}^*

$$\frac{3\mathbf{V}^* - 4\mathbf{V}^n + \mathbf{V}^{n-1}}{2\Delta t} + \widetilde{\mathcal{C}\mathcal{T}}^{n+1} = -\nabla P^n + \frac{1}{\text{Re}}\Delta\mathbf{V}^*. \quad (5)$$

Note that in Eq. (5), the pressure P^{n+1} has been substituted by P^n . Defining $\sigma := 2\Delta t/3$, the Eq. (5) for \mathbf{V}^* can be recast as

$$\left(I - \frac{\sigma}{\text{Re}}\Delta\right)\mathbf{V}^* = S, \quad S := \frac{4\mathbf{V}^n - \mathbf{V}^{n-1}}{3} - \sigma(\widetilde{\mathcal{C}\mathcal{T}}^{n+1} + \nabla P^n). \quad (6)$$

To solve Eq. (6), we apply the Alternating Direction Implicit (ADI) method which approximates the 3-D operator $(I - \frac{\sigma}{\text{Re}}\Delta)$ as a product of three 1-D operators:

$$\left(I - \frac{\sigma}{\text{Re}}\Delta\right) = \left(I - \frac{\sigma}{\text{Re}}\Delta_x\right)\left(I - \frac{\sigma}{\text{Re}}\Delta_y\right)\left(I - \frac{\sigma}{\text{Re}}\Delta_z\right) + O(\Delta t^2),$$

where Δ_i is the second order derivative along the i -direction. With this product, we can transform the 3-D Helmholtz-like equation (Eq. (6)) into a system of three 1-D Helmholtz-like equations

$$\left(I - \frac{\sigma}{\text{Re}}\Delta_x\right)T_1 = S, \quad \left(I - \frac{\sigma}{\text{Re}}\Delta_y\right)T_2 = T_1, \quad \left(I - \frac{\sigma}{\text{Re}}\Delta_z\right)\mathbf{V}^* = T_2. \quad (7)$$

In the remainder of this paper, Eq. (7) will be simply referred to as Helmholtz equation.

2.1.2 Projection step

In general, \mathbf{V}^* is not divergence-free so it is corrected in the second step. The correction is sought as a potential velocity field as follows. Taking the divergence of the difference of Eqs. (3) and (5), it comes:

$$\nabla \cdot (\mathbf{V}^{n+1} - \mathbf{V}^*) = -\sigma \Delta(P^{n+1} - P^n) + \frac{\sigma}{\text{Re}} \nabla \cdot \Delta(\mathbf{V}^{n+1} - \mathbf{V}^*). \quad (8)$$

Since $\nabla \cdot \mathbf{V}^{n+1} = 0$ and since, for any vector \mathbf{U} , we have $\nabla \cdot \Delta \mathbf{U} = \Delta \nabla \cdot \mathbf{U}$, then Eq. (8) becomes

$$\nabla \cdot \mathbf{V}^* = \sigma \Delta \left(P^{n+1} - P^n - \frac{1}{\text{Re}} \nabla \cdot \mathbf{V}^* \right). \quad (9)$$

Denoting $\phi := P^{n+1} - P^n + \frac{1}{\text{Re}} \nabla \cdot \mathbf{V}^*$, then ϕ satisfies the Poisson equation

$$\Delta \phi = \frac{1}{\sigma} \nabla \cdot \mathbf{V}^*, \quad (10)$$

which is complemented with appropriate boundary conditions. After computing ϕ , we update the velocity and pressure fields using

$$\mathbf{V}^{n+1} = \mathbf{V}^* - \sigma \phi, \quad P^{n+1} = P^n + \phi - \frac{1}{\text{Re}} \nabla \cdot \mathbf{V}^*.$$

There exist several methods to solve the Poisson equation given in Eq. (10). For example, the fast Fourier transformation (FFT), fast sine/cosine transformations, and multigrid methods [28]. In this paper, we consider a direct method based on partial diagonalization.

Let us denote $\Delta \phi = \frac{1}{\sigma} \nabla \cdot \mathbf{V}^* = S$, and rewrite $\Delta = \Delta_x + \Delta_y + \Delta_z$. Then the second order central schemes for these 1-D operators Δ_x , Δ_y and Δ_z give three tridiagonal matrices for the considered direction. To further simplify the computation, we transform tridiagonal matrices into diagonal ones by using their eigendecomposition [32].

Given a real square ($N \times N$) matrix A with N linearly independent eigenvectors q_i ($i = 1, \dots, N$), then A can be factorized as $A = Q \Lambda Q^{-1}$ where Λ is the diagonal matrix whose diagonal elements are the corresponding eigenvalues, and Q is the square ($N \times N$) matrix whose i th column is the eigenvector q_i of A .

For our problem, we decompose only Δ_x , Δ_y and keep Δ_z , then we have

$$(Q_x \Lambda_x Q_x^{-1} + Q_y \Lambda_y Q_y^{-1} + \Delta_z) \phi = S. \quad (11)$$

Then we multiply Eq. (11) with $Q_x^{-1} Q_y^{-1}$, and observe that

$$Q_x^{-1} Q_y^{-1} = Q_y^{-1} Q_x^{-1}, \quad Q_x^{-1} \Lambda_y^{-1} = \Lambda_y^{-1} Q_x^{-1}, \quad Q_y^{-1} \Lambda_x^{-1} = \Lambda_x^{-1} Q_y^{-1},$$

so we have

$$(\Lambda_x + \Lambda_y + \Delta_z) \tilde{\phi} = \tilde{S}, \quad \text{where } \tilde{\phi} = Q_y^{-1} Q_x^{-1} \phi \text{ and } \tilde{S} = Q_y^{-1} Q_x^{-1} S. \quad (12)$$

So the original Poisson equation (Eq. 10) is now transformed into the form of Eq. (12) which is a 1-D tridiagonal equation for the z -direction.

3 Parallel implementation for solving Navier-Stokes equations

3.1 Domain decomposition and Schur complement method

In our parallelization strategy we divide the domain equally into subdomains and we assign each subdomain to one MPI process in such a way that each process has the same amount of computational workload. Once the subdomains and the boundary conditions have been defined, the computations on each subdomain can be performed in parallel. The boundaries here include the real domain boundaries and the interfaces between the subdomains. The classical Schur complement method [32] is used to ensure continuity of the solution at the interface between subdomains, and the information at the interfaces is exchanged via MPI routines.

Fig. 2 shows an example of a 2-D domain decomposition with boundary cells. For instance in subdomain IV, the top and right boundaries are real domain boundaries while the left and lower boundaries are actually interfaces with neighboring subdomains.

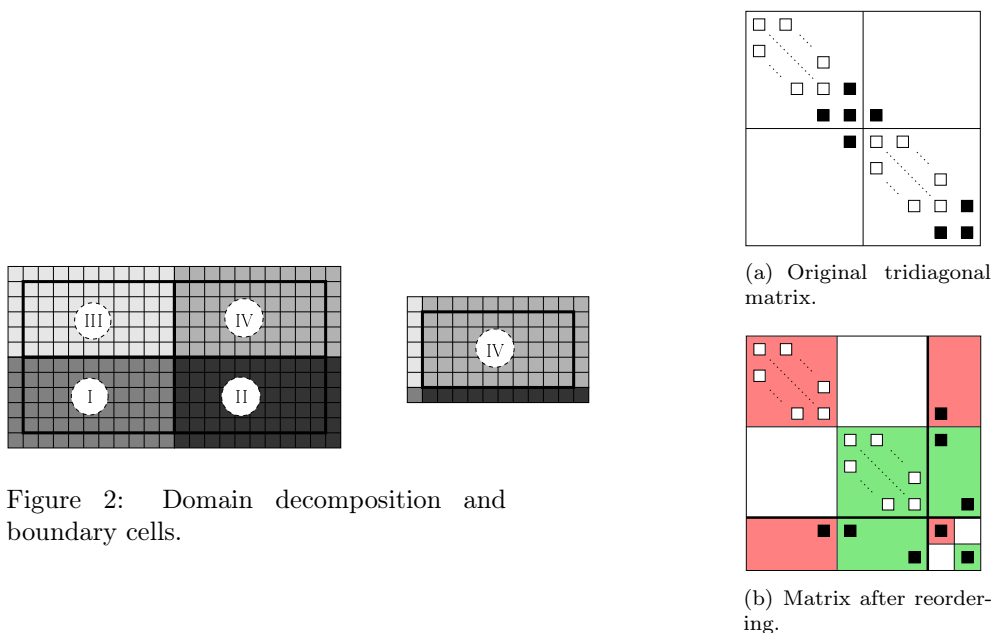


Figure 2: Domain decomposition and boundary cells.

Figure 3: Illustration of variable reordering for Schur Complement method for the case of 2 subdomains. (\square : interior variables. \blacksquare : interface variables.)

We suppose now that we have obtained the tridiagonal system for one direction (corresponding for instance to the x-direction, as depicted in Fig. 2). The matrix structure is shown by Fig. 3(a). This tridiagonal system is solved by the Schur complement method. By reordering the variables so that the interface variables are listed last, we obtain a new system where the matrix has a structure as depicted in Fig. 3(b).

We recall below how the solution is computed via the Schur complement method. Given a

block system

$$\begin{pmatrix} A_k & A_{ki} \\ A_{ik} & A_i \end{pmatrix} \begin{pmatrix} X_k \\ X_i \end{pmatrix} = \begin{pmatrix} S_k \\ S_i \end{pmatrix}, \quad (13)$$

where indices k and i refer respectively to interior and interface variables.

From the first block equation in Eq. (13) the unknown vector X_k can be expressed as

$$X_k = A_k^{-1}(S_k - A_{ki}X_i), \quad (14)$$

and we substitute this into the second equation to obtain

$$(A_i - A_{ik}A_k^{-1}A_{ki})X_i = S_i - A_{ik}A_k^{-1}S_k. \quad (15)$$

If the Schur complement matrix $A_i - A_{ik}A_k^{-1}A_{ki}$ can be formed and the linear system (15) can be solved, all the interface variables X_i are then available and the remaining unknowns can be computed via Eq. (14). Due to the particular block structure of A_k and A_i , we have r separate systems that can be solved in parallel, where r is the number of subdomains along the direction in which we solve the original tridiagonal system.

3.2 Solving tridiagonal systems

We describe in this section how the block-tridiagonal systems resulting from the discretizations of Eq. (7) (using ADI method) and Eq. (10) (using partial diagonalization) can be solved efficiently. There have been several efforts in recent years to develop efficient tridiagonal solvers [8, 22, 30]. A well-known procedure is based on the Thomas algorithm [7] which basically corresponds to a Gaussian elimination without pivoting.

Given a tridiagonal system $Ax = s$, where $A \in \mathbb{R}^{n \times n}$ is diagonally dominant (nonsingular), and $x, s \in \mathbb{R}^n$, the Thomas algorithm consists of two steps. The first step is a forward elimination where we eliminate the lower diagonal coefficients and transform the original matrix into an upper triangular one. The second step is a backward substitution that solves the upper triangular system. This requires $\mathcal{O}(n)$ operations.

To better understand the structure of the linear system, let us take a simple 2-D example as depicted in Fig. 4 which uses a 4×3 2D mesh and let us consider the operator Δ_x in Eq. (7). We order the unknowns first by row and then by column as shown in Fig. 4. Then the matrix corresponding to the operator Δ_x can be represented as the matrix given in Fig. 5 (with $m = 3$ tridiagonal blocks). Similarly, if we consider the operator Δ_y in Eq. (7), the unknowns will be ordered first by column then by row and the resulting matrix will contain 4 tridiagonal blocks.

We apply the Thomas algorithm within each tridiagonal block. If the tridiagonal blocks are identical, then we consider a smaller system with multiple right-hand sides (RHS) that can be expressed as $AX = S$ with $A \in \mathbb{R}^{n \times n}$ and $X, S \in \mathbb{R}^{n \times m}$. In the following we explain how the solution of this system can be accelerated using vectorization techniques.

Since the late 90's, processor manufacturers provide specialized processing units called multimedia extensions or Single Instruction Multiple Data (SIMD) extensions. This feature allows processors to exploit the latent data parallelism available in applications by executing a given instruction simultaneously on multiple data stored in a single special register. However, taking advantage of SIMD extensions in applications remains a complex task. We can find for instance in [14] a description of Boost.SIMD, a high-level C++ library to program SIMD architectures. Boost.SIMD provides both expressiveness and performance by using generic programming to handle vectorization in a portable way.

9	10	11	12
5	6	7	8
1	2	3	4

Figure 4: 2D $n \times m$ mesh with $n = 4, m = 3$.




Figure 5: An example of block tridiagonal matrix for operator Δ_x .

A first requirement to achieve performance using SIMD is to ensure memory alignment. In our case, this means that the Fortran arrays have to be allocated by a dedicated C++ function that uses the *ISO_C_BINDING* interface.

The SIMD opportunity in our algorithm stands from the fact that solving the tridiagonal systems with multiple RHS can be performed in parallel. This means that once we have an SIMD vector containing the tridiagonal values of a set of rows, the backward substitution of these rows can be issued as a single SIMD instruction. The triangularization followed by the backward substitution will be accelerated by the fact that data handled at each step of the algorithm will be contiguous in memory using a shuffling operation. Shuffling is a typical idiom of SIMD programming that replaces some class of complex memory access patterns by computation on simpler patterns. Depending on the SIMD extension, these shuffling operations can either be limited (like in SSE2 where shuffling can only occur piecewise inside a given register) or be arbitrary (*e.g.* in AltiVec where shuffling corresponds to permutations of bytes). Shuffling also enables deinterleaving of scattered data that can be fetched from main memory and bring back into a single, contiguous SIMD register. In our case, we use shuffling to aggregate values from the sparse representation of the system into a set of contiguous values located at the proper place so that the elimination and backward substitution phases can take place inside an SIMD vector. Expected benefit should be a faster code due to less memory access in different part of the memory (thus limiting cache misses) and augmented SIMD speedup due to the layout of the data after the shuffling and better use of pipelining.

We illustrate this method in Fig. 6 by considering 2 RHS (case of SSE where a double precision vector holds 2 values) where we would perform the following operations:

- load 3 registers units with coefficients $(a_i, a_{i+1}), (b_i, b_{i+1})$ and (c_i, c_{i+1}) coming from respectively the lower, main and upper diagonals,
- shuffle the data to have a 2×3 local matrix $\begin{pmatrix} a_i & b_i & c_i \\ a_{i+1} & b_{i+1} & \end{pmatrix}$,
- perform two Gaussian eliminations to eliminate the coefficients (a_i, a_{i+1}) in the lower diagonal,
- after elimination of the a_i 's, we perform a backward substitution to solve the upper triangular

system with the RHS coefficients loaded in one register (*e.g.* 2 RHS for SSE or 4 RHS for AVX SIMD extension).




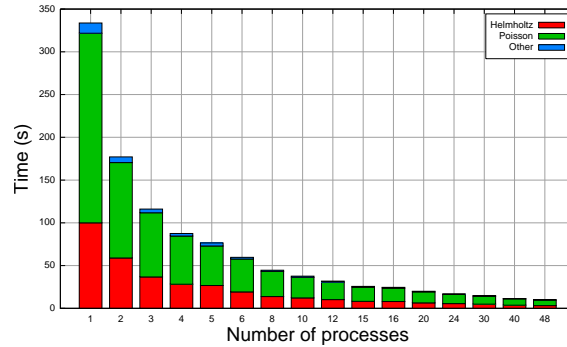
Figure 6: Thomas algorithm with vectorization.

Our implementation uses this method to handle multiple RHS. Note that the memory alignment mentioned previously enables here the “load” process to be performed efficiently. We point out that this vectorization approach can be also used in other parts of our solver, for instance to compute the source term faster.


4 Performance results

The following experiments were carried out using a MagnyCours-48 system from University of Tennessee. This machine has a NUMA architecture and is composed of four AMD Opteron 6172 (with a SSE-4a instruction set) running at 2.1GHz with twelve cores each (48 cores total) and 128GB of memory. Our solver is linked with the LAPACK and ScaLAPACK routines from the 10.3.6 version of the Intel MKL [21] library and communications are performed using OpenMPI 1.4.3. In our experiments, we use one MPI process per core and each MPI process is a single thread.

In Fig. 7(a) we consider a 3-D vortex problem with mesh size 240^3 , *i.e.* about 1.4×10^7 unknowns. We represent the performance (in seconds) for one iteration of the NS solver including the Helmholtz and Poisson equations, and miscellaneous tasks (mainly I/O and velocity/pressure updates). The number of threads varies from 1 to 48. We observe in Fig. 7(a) that the CPU time decreases significantly with the number of threads, showing then a good scalability of the solver. We observe that the Helmholtz equation represents about 30% of the global computational time and this percentage remains the same when the number of threads increases.



(a) Time breakdown for one iteration of the NS solver.



(b) Scalability of NS solver.

Figure 7: Performance of the Navier-Stokes solver.

Fig. 7(b) shows the parallel speedup of the solver. We obtain a speedup up of 33 using 48 cores. We also observe in Fig. 7(b) that the solution for the Poisson equation scales slightly better than for the Helmholtz equation. This is because in the Poisson equation we have only one tridiagonal system to solve (vs three systems in Helmholtz equation) and thus there is less information exchanged (as mentioned in Section 3, when we solve a tridiagonal system via the Schur complement method, we have a reordering of the variables that results in additional communication).

In a second experiment, we consider a 3-D vortex problem where the size increases with the number of threads and with a fixed mesh size per subdomain (weak scalability). The local size is set to $240 \times 240 \times 10$ and 10 time iterations are performed. We observe in Fig. 8 that the time spent in solving Helmholtz and Poisson equations does not vary (about 60 and 110 milliseconds respectively) because the number of unknowns computed is always $240^2 \times 10 = 5.76 \times 10^5$. However the global CPU time increases linearly with the number of processes due to a larger amount of I/O after each iteration. Indeed, the numerical solution is stored after each time iteration in order to generate a detailed animation that visualizes the fluid movement. In this numerical test, we only divide the domain along one direction (z -direction here) to keep the shape of the interfaces. Thus, the amount of information to be exchanged (the number of interface elements) is $240^2 \times (\text{number of MPI processes} - 1)$.

Let us now study more specifically the performance of the tridiagonal solver described in Section 3.2. We compare the performance of a vectorized Thomas algorithm (as illustrated by




Figure 8: Scalability of problems with variable size.

Fig. 6) with the LAPACK 3.2 routine DGTSV from Netlib, linked with the MKL BLAS library, which solves a general tridiagonal system using Gaussian elimination with partial pivoting (note that, since our matrices are diagonally dominant, the routine DGTSV does not pivot and only the search for pivot is performed).

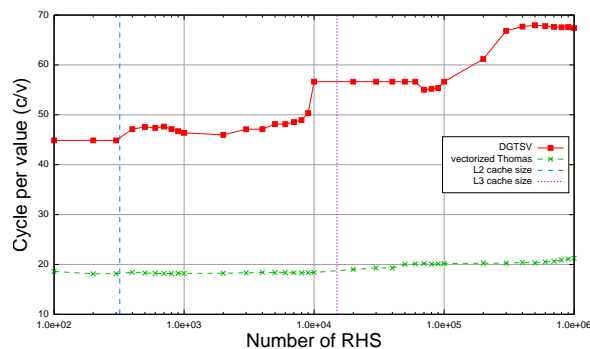
In Fig. 9(a), we plot the number of cycles per value, *i.e.* the amount of CPU cycles required to compute one element of the result (we operate here on tridiagonal blocks of size 100 in order to exhibit the cache effects). This metric allows us to do a fine grain analysis of both the impact of vectorization and the impact of memory access. For the curve related to DGTSV, we observe a cycle/value amount that jumps when we hit the cache size (L2 and L3). This is due to the fact that every computation in the elimination and backward substitution processes requires more memory access to be completed. However, when we consider the vectorized Thomas algorithm, we have a constant amount of CPU cycle/value up to the largest sizes, due to the shuffling that replaces memory access by computation. Fig. 9(b) shows the time ratio between DGTSV and the vectorized Thomas algorithm. For large problem sizes (*e.g.* more than 10^6 unknowns), we observe that the vectorized version of Thomas algorithm outperforms DGTSV with a factor 3.

5 Conclusion


In this paper, we presented a numerical method and a parallel implementation for the solution of the 3-D Navier-Stokes equations. The resulting solver is scalable on problems of realistic size and the solution of the tridiagonal systems can be significantly enhanced by using vectorization. We point out that this solver can also take into account other characteristics of the fluid or of the domain that are not discussed in details in the paper. For example, several boundary conditions are implemented and the domain can be treated either as continuous or discontinuous. Note that other physical factors like temperature and density are also considered in the solver.

Acknowledgments

This work was supported by Région Île-de-France and Digitéo (<http://www.digiteo.fr>), CAL-IFHA project, contract No 2011-038D. We also thank Pierre Esterie (University Paris-Sud) for his support in using Boost SIMD.



(a) Cycles per value for Thomas algorithm.



(b) Time ratio between DGTSV and vectorized Thomas.

Figure 9: Performance of vectorized Thomas algorithm.

References

- [1] Basic Linear Algebra Subprograms Technical Forum Standard. *Int. J. of High Performance Computing Applications*, 16(1), 2002.
- [2] Hervé Abdi. The eigen-decomposition: Eigenvalues and eigenvectors. In Neil J. Salkind, editor, *Encyclopedia of Measurement and Statistics*, pages 304–308. Sage Publications, 2007.
- [3] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, and *et al.* *LAPACK Users' Guide*. SIAM, Philadelphia, 1999. Third edition.
- [4] J. Anderson. *Computational Fluid Dynamics: The Basics with Applications*. McGraw-Hill, 1995.
- [5] Joao P. D. Angeli, Andrea M. P. Valli, Neyval C. Reis Jr., and Alberto F. De Souza. Finite difference simulations of the Navier-Stokes equations using parallel distributed computing. 15th Symposium on Computer Architecture and High Performance Computing, 2003.
- [6] Uri. M. Ascher and Linda R. Petzold. *Computational Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, 1998.
- [7] K. Atkinson and W. Han. *Elementary Numerical Analysis*. John Wiley & Sons, 2003.

-
- [8] Dario A. Bini and Beatrice Meini. The cyclic reduction algorithm: From Poisson equation to stochastic processes and beyond. *Journal of Numerical Algorithms*, 51(1):23–60, 2009.
- [9] L. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I. Dhillon, and et al. *ScaLAPACK Users’ Guide*. SIAM, Philadelphia, 1997.
- [10] David L. Brown, Richardo Cortez, and Michael L. Minion. Accurate projection methods for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, pages 464–499, 2001.
- [11] A. Q. C. Canuto, M. Meneguzzi, and T. Zang. *Spectral Methods: Fundamentals in Single Domains*. Springer, 2006.
- [12] Alexandre J. Chorin. Numerical solution of the Navier-Stokes equations. *Mathematics of Computation*, 22(104):745–762, 1968.
- [13] G. Cottet and P. Koumoutsakos. *Vortex Methods Theory and Practice*. Cambridge University Press, 2000.
- [14] Pierre Esterie, Mathias Gaunard, Joel Falcou, Jean-Thierry Lapresté, and Brigitte Rozoy. Boost.simd: generic programming for portable simdization. In *Proceedings of the 21st international conference on Parallel architectures and compilation techniques*, pages 431–432, New York, NY, USA, 2012.
- [15] J. H. Ferziger and M. Perić. *Computational Methods for Fluid Dynamics*. Springer, 3 edition, 1996.
- [16] Message Passing Interface Forum. *MPI : A Message-Passing Interface Standard*. Int. J. Supercomputer Applications and High Performance Computing, 1994.
- [17] A. Graham. *Kronecker Products and Matrix Calculus with Application*. Wiley, New York, 1981.
- [18] J. L. Guermond, P. D. Mineev, and J. Shen. An overview of projection methods for incompressible flows. *CMAME*, 195:6011–6045, 2006.
- [19] F. Harlow and J. E. Welch. Numerical calculation of time-deendent viscous incompressible flow of fluid with free surface. *The Physics of Fluids*, 8(12):2182–2189, 1965.
- [20] Willem Hundsdorfer. Partially implicit BDF2 blends for convection dominated flows. *Journal of Numerical Analysis*, 38(6):1763–1783, 2001.
- [21] Intel. *Math Kernel Library (MKL)*. <http://www.intel.com/software/products/mkl/>.
- [22] George Em Karniadakis and Robert M. Kirby II. *Parallel Scientific Computing in C++ and MPI*. Cambridge University Press, 2003.
- [23] J. Kim and P. Moin. Application of a fractional-step method to incompressible Naver-Stokes equations. *Journal of Computational Physics*, 59:308–323, 1985.
- [24] Randall J. Leveque. *Numerical Methods for Conservation Laws*. Lectures in Mathematics. Eth Zurich, 2 edition, 1999.
- [25] Shaofan Li and Wing Kam Liu. Meshfree and particle methods and their applications. *Applied Mechanics Reviews*, 55:1–34, 2002.

-
- [26] Dan Martin, Phil Colella, Noel Keen, A. Deane, G. Brenner, A. Ecer, and *et al.* An incompressible Navier-Stokes with particles algorithm and parallel implementation. In *Proceedings of the 2005 International Conference on Parallel Computational Fluid Dynamics*, pages 461–468, College Park, MD, USA, 2006.
- [27] J. J. Monaghan. Smoothed particle hydrodynamics. *Annual Review of Astronomy and Astrophysics*, 30:543–574, 1992.
- [28] S. Noury, S. Boivin, and Olivier Le Maître. *A fast Poisson solver for OpenCL using multigrid methods*. GPU Pro2: Advanced Rendering Techniques, 2011.
- [29] Olivier Pironneau. *The Finite Element Method for Fluids*. John Wiley & Sons, 1 edition, 1990.
- [30] Eric Polizzi and Ahmed H. Sameh. A parallel hybrid banded system solver: the SPIKE algorithm. *Journal of Parallel Computing*, 32:177–194, 2006.
- [31] Rolf Rannacher. Finite element methods for the incompressible Navier-Stokes equations. *Fundamental Directions in Mathematical Fluid Mechanics*, pages 191–293, 2000.
- [32] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, 2 edition, 2003.
- [33] A. Vincent and M. Meneguzzi. The spatial structure and statistical properties of homogeneous turbulence. *Journal of Fluid Mechanics*, 225:1–20, 1991.

Contents

1	Introduction	3
2	Navier-Stokes Problem	4
2.1	Prediction-Projection method	4
2.1.1	Prediction step	5
2.1.2	Projection step	6
3	Parallel implementation for solving Navier-Stokes equations	7
3.1	Domain decomposition and Schur complement method	7
3.2	Solving tridiagonal systems	8
4	Performance results	10
5	Conclusion	12



**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

1 rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing
Campus de l'École Polytechnique
91120 Palaiseau

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399