

NEW TRENDS IN COUPLED SIMULATIONS FEATURING DOMAIN DECOMPOSITION AND METACOMPUTING

PHILIPPE D'ANFRAY¹, LAURENCE HALPERN² AND JULIETTE RYAN³

Abstract. In this paper we test the feasibility of coupling two heterogeneous mathematical modeling integrated within two different codes residing on distant sites. A prototype is developed using Schwarz type domain decomposition as the mathematical tool for coupling. The computing technology for coupling uses a CORBA environment to implement a distributed client-server programming model. Domain decomposition methods are well suited to reducing complex physical phenomena into a sequence of parallel subproblems in time and space. The whole process is easily tuned to underlying hardware requirements.

Mathematics Subject Classification. 65M55, 65Y05.

Received: December 7, 2001. Revised: May 23, 2002.

1. INTRODUCTION

Within the realm of partial differential equations, multifields models can quite significantly reduce the problem complexity while increasing the numerical accuracy. The complexity may be due to heterogeneous physics, mathematics, discretisations, or heterogeneous computing environments.

A solution to the physical and mathematical aspect is domain decomposition methods, well suited to reducing complex physical phenomena into a sequence of subproblems in time and space easier to model and which to some extent can be solved simultaneously.

Many tools have been developed recently to achieve such cooperative computing that often require thorough modifications of codes to integrate new functionalities for coupling and much thinking as to where and who is in charge of the coupling. The goal of the CORBA standards developed by the OMG (Object Management Group) is to simplify these procedures.

Recent extensions of Schwarz type domain decomposition to time decomposition allow greater flexibility in terms of granularity of the computation. Time-space domains can be adjusted to the underlying computing hardware and networks.

The aim of this paper is to test the feasibility of this process, applied to heterogeneous mathematical modeling of the convection diffusion problem, running on different computing sites with the use of, and monitored by, the CORBA environment.

Keywords and phrases. Domain decomposition, evolution equations, coupling of applications, heterogeneous computations, distributed computing, meta-computing, CORBA.

¹ CEA DTI-CISC, 91191 Gif-sur-Yvette Cedex France (on leave from ONERA DTIM-CHP) and LAGA, Université Paris XIII, 93430 Villeteuse, France. e-mail: Philippe.Anfray@cea.fr

² LAGA, Université Paris XIII, 93430 Villeteuse, France. e-mail: halpern@math.univ-paris13.fr

³ ONERA DTIM-CHP and LAGA, Université Paris XIII, 93430 Villeteuse, France. e-mail: ryan@math.univ-paris13.fr

The first part of this paper describes the time and space domain decomposition technique involved, the second part is a description of the CORBA technology. Finally an application is presented with the case of a convection diffusion problem around an airfoil. The domain is split into 2 subdomains, finite elements round the airfoil and finite differences in the far field, each domain being computed with two different codes coupled via CORBA.

2. TOOLS IN DOMAIN DECOMPOSITION METHODS

Let $(0, T)$ be a bounded time interval, and Ω a domain in \mathbb{R}^2 with boundary $\Gamma = \partial\Omega$. Let us consider the convection diffusion equation

$$\mathcal{P}U \equiv \frac{\partial U}{\partial t} + \mathbf{a} \cdot \nabla U - \nu \Delta U = 0 \text{ in } \Omega \times (0, T), \tag{1}$$

associated with the initial data $U(\cdot, 0) = U^0$ in Ω at time 0 and the Dirichlet boundary condition $U = 0$ on $\Gamma = \partial\Omega$.

The viscosity ν is a strictly positive coefficient. The advection velocity \mathbf{a} is a smooth function of x , in our case always non-zero in Ω .

A general technique to solve the latter is to first discretise in time with an implicit scheme such as the second order Crank-Nicolson scheme :

$$\frac{u^{n+1} - u^n}{\Delta t} + \frac{1}{2} [\mathbf{a} \cdot \nabla u^{n+1} - \nu \Delta u^{n+1}] + \frac{1}{2} [\mathbf{a} \cdot \nabla u^n - \nu \Delta u^n] = 0 \text{ in } \Omega, \tag{2}$$

where u^n is an approximation of $U(\cdot, t_n)$. At each step n , u^{n+1} is solution of a steady equation such as $\mathcal{L}u = F$ where

$$\mathcal{L} \equiv \frac{2}{\Delta t} + \mathbf{a} \cdot \nabla - \nu \Delta$$

and $F = (\frac{2}{\Delta t} - \mathbf{a} \cdot \nabla + \nu \Delta) u^n$.

In the rest of this paper, the quantity $\frac{2}{\Delta t}$ will be denoted by c .

Fast solvers for steady problems can now be applied such as non-overlapping domain decomposition methods using Schwarz type algorithms designed in [4] which are presented in Section 2.1. Algorithms global in time, namely Schwarz waveform relaxation algorithms as designed in [2] will be presented in Section 2.2.

2.1. Steady problems

As stated previously, we are led to the problem

$$\begin{cases} \mathcal{L}u \equiv (c + \mathbf{a} \cdot \nabla - \nu \Delta)u = F & \text{in } \Omega \\ u = 0 & \text{on } \Gamma = \partial\Omega. \end{cases} \tag{3}$$

Let us split Ω in two subdomains: $\bar{\Omega} = \bar{\Omega}_1 \cup \bar{\Omega}_2$ and denote $\Gamma_{12} = \bar{\Omega}_1 \cap \bar{\Omega}_2$. On Γ_{12} , there are two unit normals at each point: \mathbf{n}_1 is the outgoing normal to Ω_1 , \mathbf{n}_2 the outgoing normal to Ω_2 and $\mathbf{n}_1 + \mathbf{n}_2 = \mathbf{0}$. Γ_1 will denote the part of Γ bounding Ω_1 and Γ_2 that of Γ bounding Ω_2 . We shall always suppose in the theoretical part, that the advection velocity \mathbf{a} is constant and, to set matters, entering Ω_2 , i.e. that $\mathbf{a} \cdot \mathbf{n}_1 > \mathbf{0}$ along Γ_{12} , although the methodology carries over to more general situations (see [4]).

Let us now introduce the non-overlapping Schwarz algorithm.

$$\begin{cases} \mathcal{L}v^{k+1} = F, & \text{in } \Omega_1 \\ v^{k+1} = 0 & \text{on } \Gamma_1 \\ \mathcal{B}_1 v^{k+1} = \mathcal{B}_1 w^k & \text{on } \Gamma_{12} \end{cases} \quad \begin{cases} \mathcal{L}w^{k+1} = F, & \text{in } \Omega_2 \\ w^{k+1} = 0 & \text{on } \Gamma_2 \\ \mathcal{B}_2 w^{k+1} = \mathcal{B}_2 v^k & \text{on } \Gamma_{12} \end{cases} \tag{4}$$




FIGURE 1. non-overlapping 2 domains decomposition.

where \mathcal{B}_1 et \mathcal{B}_2 are given by

$$\begin{cases} \mathcal{B}_1 = \frac{\partial}{\partial n_1} - \frac{a}{\nu} + p \\ \mathcal{B}_2 = \frac{\partial}{\partial n_2} + p \end{cases} \tag{5}$$

and p is a positive real number.

Suppose $\Omega = \mathbb{R}^2$, Ω_1 and Ω_2 two half spaces delimited by $\Gamma_{12} = \{(x, y), \alpha x + \beta y = 0\}$. For simplicity we shall write $\mathbf{n} = \mathbf{n}_1 = (\alpha, \beta)$, so that $\mathbf{n}_2 = -\mathbf{n}$. We shall perform the change of variables $(X = \alpha x + \beta y, Y = -\beta x + \alpha y)$, and the two half spaces become $\Omega_1 = \{(X, Y), X < 0\}$ and $\Omega_2 = \{(X, Y), X > 0\}$. By a partial Fourier transformation along the Y variable (the dual variable is denoted by ξ), we can calculate the iterate errors $V^k = v^k - u, W^k = w^k - u$:

$$\begin{cases} \hat{V}^k = \alpha^k e^{\lambda_+ x}, & \hat{W}^k = \beta^k e^{\lambda_- x} \\ \left(\lambda_+ - \frac{a}{\nu} + p\right) \alpha^{k+1} = \left(\lambda_- - \frac{a}{\nu} + p\right) \beta^k \\ (-\lambda_- + p) \beta^{k+1} = (-\lambda_+ + p) \alpha^k \end{cases}$$

where λ_{\pm} are the roots of the characteristic polynomial

$$P(i\xi; \lambda) = -\nu\lambda^2 + \mathbf{a} \cdot \mathbf{n}\lambda + (c + i\mathbf{a} \cdot \boldsymbol{\tau}\xi + \nu\xi^2)$$

such that $\Re\lambda_+(i\xi) > 0$ and $\Re\lambda_-(i\xi) < 0$, where $\boldsymbol{\tau}$ denotes the unit vector tangent to the interface such that $(\mathbf{n}, \boldsymbol{\tau}) = +\pi/2$.

The convergence rate is thus given by:

$$\delta(i\xi) = \frac{\lambda_+(i\xi) - p}{\lambda_-(i\xi) - p}. \tag{6}$$

In [3], optimization of the convergence rate was proved to be successful, for more general transmission conditions. We choose an optimal value of p defined by

$$\inf_{p \in \mathbb{R}^+} \sup_{\xi \in [\xi_0, \xi_1]} |\delta(i\xi)| \tag{7}$$

where ξ_0 is related to the length of the Y interval, $\xi_0 = \frac{\pi}{L}$, and ξ_1 to the Y discretisation, $\xi_1 = \frac{\pi}{\Delta Y}$. There is a unique p_{opt} minimizing the convergence rate, and it is characterized by $\delta(i\xi_0) = \delta(i\xi_1)$ [5].

2.2. Evolution equations

An opposite approach has been proposed in [1]: the evolution equation itself is solved with a domain decomposition technique allowing for different time steps in each subdomain.

Let us resume the evolution equation

$$\mathcal{P}U \equiv \frac{\partial U}{\partial t} + \mathbf{a} \cdot \nabla U - \nu \Delta U = 0 \text{ in } \Omega \times (0, T),$$

associated with the initial data $U(\cdot, 0) = U^0$ in Ω at time 0 and the Dirichlet boundary condition $U = 0$ on $\Gamma = \partial\Omega$. A global Schwarz algorithm in time can be developed in the two subdomains $\Omega_1 \times (0, T)$, and $\Omega_2 \times (0, T)$,

$$\begin{cases} \mathcal{P}v^{k+1} = 0, & \text{in } \Omega_1 \times (0, T) \\ v^{k+1}(\cdot, 0) = U^0 & \text{in } \Omega_1 \\ v^{k+1} = 0 & \text{on } \Gamma_1 \times (0, T) \\ \mathcal{B}_1 v^{k+1} = \mathcal{B}_1 w^k & \text{on } \Gamma_{12} \times (0, T) \end{cases} \tag{8}$$

$$\begin{cases} \mathcal{P}w^{k+1} = 0, & \text{in } \Omega_2 \times (0, T) \\ w^{k+1}(\cdot, 0) = U^0 & \text{in } \Omega_2 \\ w^{k+1} = 0 & \text{on } \Gamma_2 \times (0, T) \\ \mathcal{B}_2 w^{k+1} = \mathcal{B}_2 v^k & \text{on } \Gamma_{12} \times (0, T) \end{cases} \tag{9}$$

with the transmission operators given in (5). Again, in the case of two nonoverlapping half spaces, the error can be identified through a Fourier transform (parameter ξ) in the tangential direction and a Laplace transform in time (parameter s with $\eta = \Re(s) \geq \eta_0 > 0$).

More precisely, the partial Fourier transform in relation to Y and the Laplace transform in relation to t is defined as:

$$\hat{u}(x, \xi, s) = \frac{1}{2\pi} \int_{\mathbb{R}} \int_{\mathbb{R}_+} u(x, Y, t) e^{-i\xi Y - st} dY dt, \quad s = \eta + i\omega, \eta \geq \eta_0.$$

The new convergence rate is $\delta(i\xi, s)$ given formally by (6), where $\lambda_{\pm}(i\xi, s)$ are now the roots of the characteristic polynomial

$$P(i\xi, s; \lambda) = -\nu\lambda^2 + \mathbf{a} \cdot \mathbf{n}\lambda + (c + i\mathbf{a} \cdot \boldsymbol{\tau}\xi + \nu\xi^2).$$

The optimal value of p is now chosen so as to minimize

$$\sup_{\xi \in [\xi_0, \xi_1], \eta \geq \eta_0, \omega \in [\omega_0, \omega_1]} |\delta(i\xi, s)|$$

where ξ_0 is related to the length of the Y interval, $\xi_0 = \frac{\pi}{L}$, and ξ_1 to the Y discretisation, $\xi_1 = \frac{\pi}{\Delta Y}$.

In the same way, $\omega_0 = \frac{\pi}{T}$ and $\omega_1 = \frac{\pi}{\Delta T}$ while η_0 is chosen arbitrarily small.

For analyticity reasons, it is sufficient to make sure of an upper bound for $\eta = \eta_0$ and p is computed to realize

$$\inf_{p \in \mathbb{R}^+} \sup_{\xi \in [\xi_0, \xi_1], \eta = \eta_0, \omega \in [\omega_0, \omega_1]} |\delta(i\xi, s)|. \tag{10}$$

This ends the presentation of the mathematical coupling. We now proceed to introduce the coupling technology.

3. METACOMPUTING

Development of software and hardware technologies has lead to the new paradigm of **distributed** or **meta computing**. Object Oriented programming, client server model and software bus paradigm will be used to set up a new class of applications which might involve several simulation codes running on different sites and multiple interacting users. One key point here is to be able to deal with heterogeneity of computing languages, operating systems and hardware platforms.

In this framework applications are seen as “object-servers” that interact transparently through the network. We introduce briefly some useful programming and run time models for distributed applications.

3.1. Programming and run time models

The first step is to encapsulate our applications in “objects”. An object has attributes –mainly the encapsulated data– which define its state and a behavior described by its methods –or procedures–. In a general object oriented programming model one sends “messages” to objects to request execution of a local method. Hence, as described in the last sentence, typical objects are “servers”. They are able to perform a given task, *i.e.* “provide a service” upon reception of a message, *i.e.* “upon request from a client”.

It is rather straightforward to design a client-server architecture from existing software. Just think of the usual “main program” of the application as a client code requesting services from a “library part”. The functions called in this main program and the data they are working on belong to the server side, hence parts of the server can be viewed as a library. In a standard stand-alone application everything is on the same platform. In a distributed context, the caller –client side– and the library –server side– may reside on separate platforms.




FIGURE 2. A client/server application (1).

Just like hardware components are connected by a bus in a computer, a software bus is a model where applications are seen as components that can be “plugged” on the bus. All plugged components may interact without modifications of the existing ones. Figure 3 shows a “software bus” vision of a distributed system involving several servers.




FIGURE 3. A software bus.

When Object Oriented Programming is used to design a client-Server architecture, we will use the term “distributed objects” and will refer to “objects living on the server side” and “objects living on the client side”.




FIGURE 4. Corba, a user's view.

In this framework, the run time model is defined as objects interacting by exchanging messages. Those objects may live anywhere in the distributed system. This is transparent to the user. Here we will use C++ for object oriented programming but an interesting alternative is to use the Java programming language [6, 7] for large numerical simulations [8, 9]. The Remote Method Invocation (RMI) mechanism introduced in Java [10] allows one to design distributed objects applications.

Most recent software environments now address the “**Computing Grid**” where project developments integrate geographically distributed resources shared by groups of users [11, 12]. The challenge for many research projects such as **Globus** [13, 14] or **Legion** [15, 16] is to build powerful and “user friendly” hardware and software environments that provide uniform and network transparent access to those resources (computing power, information, applications, data bases, ...).

Our project is obviously more limited but we aim at designing applications as a set of software components that can be reused within any general framework.

3.2. Corba

Corba [17–19] stands for “*Common Object Request Broker Architecture*”. It is an attempt to define a standard for programming distributed applications with the three models discussed above: object oriented; client-server and software bus. As stated before this is “Distributed Objects” but objects will now communicate through the **Corba** software bus.

Corba is defined by the **OMG** [20] which is a consortium of users and hardware or software vendors. It is an open standard in the sense that all members of the **OMG** consortium agree to it and may examine new proposals. The **OMG** does not provide any implementation of the standard.

Let us just cite some main goals of **Corba** which are particularly relevant to the context of distributed numerical simulation:

- link with existing programming languages (C, C++, Java, ...);
- automatic and transparent activation of objects; transparency of method invocation: “local” or “remote”;
- self descriptive system. Requests can be made at compile time –statically– or built at run-time –dynamically–;
- interoperability of **Corba** buses –*i.e.* of various implementations– through well defined protocols.

The user's view of a **Corba** system is quite simple, see Figure 4. The server may be located anywhere but is accessed through an interface. The client issues requests *i.e.* method invocations, on a “proxy”. **Corba** (the request broker, the bus, ...) takes care of the requests and potential results. An implementation of a standard like **Corba** is called a **middleware** as it makes the link between the applications (software level) and computers and their operating systems (hardware level).

The **Corba** standard includes the definition of an Interface Description Language called **IDL**. Using **IDL** allows us to be independent from the programming language used *in fine*. **IDL** is not a programming language since it is just used to describe interfaces.

The C++ language [21] is now widely used for large scale simulations [22]. For C++ users, the **IDL** specification almost matches the specifications in the **public** part of a **class** construct. The **IDL** compiler makes “projections” of the specifications onto the desired programming languages for the client and for the server. The projection

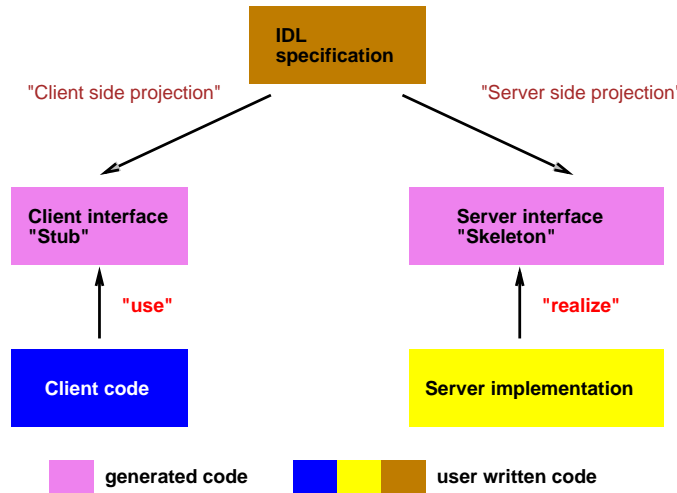


FIGURE 5. Structure of a “Corba software”.




FIGURE 6. A client/server application (2).

on the client side is called a **stub**: it is the interface used by the server. The projection on the server side is called a **skeleton**: here the operations have to be implemented.

If an object oriented language such as C++ is used for both projections the whole process is easy to handle through inheritance. Here, the IDL compiler generates an interface (*i.e.* specification part of a C++ class) from the IDL description. Stub and Skeleton are generated as classes derived from this interface by inheritance and finally the server implementation inherits from the generated skeleton as in Figure 5.

Going back to Figure 2 we see in Figure 6 that we have filled the dashed boxes: the stub and the skeleton are the components which allow us to distribute the application. Local and remote calls now have the same syntax and the same semantic.

3.3. Conclusions

We have shown how an application can be developed as an “object server” and then can be run through remote control. Obviously this approach stands for new applications, but legacy software (typically Fortran codes) can easily be “encapsulated” within an object framework (*e.g.* a C++ class) and turned into such a server without major restructuring.

As a consequence software developed inside this framework does not need to be ported locally to be accessible. It might even be a commercial application made available by an “Application Service Provider” (ASP).

Once a server is developed, it is ready to be used as a component within a distributed application. Of courses new services will be needed for the components to communicate. It remains now to test the feasibility of this technology on a domain decomposition test case.

4. APPLICATION

An important aspect of domain decomposition is that it allows for geometric simplification with different discretisations along with various modelisations. Complex geometrical elements are time consuming and should be as few as possible. A typical problem is meshing with smoothness the connection of a wing to the fuselage, or any complex body within a large regular domain. Our test case concerns the coupling of finite elements with finite differences discretisations of a same equation, relevant to this problem.

4.1. The test case

The equation is the two-dimensional convection diffusion of a scalar (temperature), see equation (1) given in Section 2 with $\nu = 0.01$ and $\mathbf{a} = (1, 0)$, describing a heat bubble hitting an up-going Naca airfoil. The computing domain consists of two subdomains, the first around the airfoil using triangles for a better definition of the profile, the second domain discretising the far field with regular rectangles. External boundary conditions are imposed: Dirichlet values on the upstream ($u = v_1$), top ($u = 0$), bottom sides ($u = 0$) and an homogeneous Neumann boundary condition on the downstream side and the airfoil, see Figure 7.




FIGURE 7. Non-overlapping 2 domains decomposition.

4.2. Numerical algorithms

Two independent different codes (EF, DF) were coupled, see Figures 8 and 9. Both solve the two-dimensional convection diffusion of the scalar u in equation (1) with the same time scheme (Crank-Nicolson).

EF space discretisation uses P1 finite elements, on the triangular mesh close to the airfoil, the implicit part in the time scheme is solved using a Block Lower Upper decomposition.

DF space discretisation is based on second order finite differences on a cartesian grid, the implicit part in the time scheme is solved iteratively with a GMRES solver. The Finite Difference grid meshes in one block the whole domain including the airfoil, points situated on the grid within the inner Finite Element domain are set to zero and boundary conditions are imposed on the interface with the finite element domain.




FIGURE 8. Finite element (EF) and finite difference (DF) grids.




FIGURE 9. Closeup of the airfoil.

These two codes are coupled using the Schwarz iterative procedure presented above, with Robin type transmission operators $\frac{\partial u}{\partial n} + \alpha_i u$ on the interface.

4.2.1. *Computation of the Robin coefficients*

As shown in Section 2, the $(\alpha_i)_{i=1,2}$ are chosen so as to minimise the convergence rate of the Schwarz algorithm and are given by $\alpha_2 = \mathbf{a} \cdot \mathbf{n} / \nu - \alpha_1$ and $\alpha_1 = p$ such that $\inf_{p \in \mathbb{R}^+} \sup_{\xi \in [\xi_0, \xi_1]} \left| \frac{\lambda_+ - p}{\lambda_- - p} \right|$ where ξ_0 is related to the length of the y interval, $\xi_0 = \frac{\pi}{L}$, and ξ_1 to the y discretisation, $\xi_1 = \frac{\pi}{\Delta y}$.

These coefficients were numerically computed using an iterative zooming technique. p is first determined (as p_0) in a initial interval $[0, p_{\max}]$ meshed with N points. Then another $p = p_1$ is searched in $[p_0 - \epsilon, p_0 + \epsilon]$ meshed with N points. Search is ended when in the case of steady coupling $\frac{\lambda_+(k \min) - p}{\lambda_-(k \min) - p} = \frac{\lambda_+(k \max) - p}{\lambda_-(k \max) - p}$ [5], and in the unsteady case when $p_i - p_{i+1} < 10^{-12}$ ($k \min$ and $k \max$ are minimum and maximum frequencies). Three or four iterations were sufficient for $N = 1000$.

In Figures 10 and 11, convergence rate for the Robin value minimising (7) and (10) is shown *versus* space frequencies in the steady case, and *versus* space and time frequencies in the unsteady case.

4.2.2. *Steady coupling*

In the steady case the algorithm is the following: at each time step with index n , and for each Schwarz iteration with index k , the following problems are solved:

$$\begin{cases} \mathcal{L}_1 v_k^{n+1} = F_1, & \text{in } \Omega_1 \\ v_k^{n+1} = v_1 & \text{on } \Gamma_1 \\ \frac{\partial v_k^{n+1}}{\partial n} = 0 & \text{on } \Gamma'_1 \\ \mathcal{B}_1 v_k^{n+1} = \mathcal{B}_1 w_k^n & \text{on } \Gamma_{12} \end{cases} \tag{11}$$




FIGURE 10. Steady case.




FIGURE 11. Unsteady case.

where \mathcal{L}_1 is the second order finite difference discretisation of $[Id + dt/2(\mathbf{a}\nabla - \nu\Delta)]$ and $F_1 = [Id - dt/2(\mathbf{a}\nabla - \nu\Delta)]v_k^n$ and the boundary condition is imposed explicitly

$$\begin{cases} \mathcal{L}_2 w_{k+1}^{n+1} = F_2, & \text{in } \Omega_2 \\ \frac{\partial w_k^{n+1}}{\partial n} = 0 & \text{on } \Gamma_2 \\ \mathcal{B}_2 w_k^{n+1} = \mathcal{B}_2 v_{k+1}^n & \text{on } \Gamma_{12}. \end{cases} \quad (12)$$

These second set of equations are discretised using finite element integration in space and finite differences in time. They become

$$\int_{\Gamma_{12}} \alpha_2 w_{k+1}^{n+1} v \, d\gamma + \int_{\Omega_2} (w_{k+1}^{n+1} v + dt/2(\mathbf{a}\nabla)w_{k+1}^{n+1} v + \nu dt/2 \nabla w_{k+1}^{n+1} \nabla v) \, dx = \int_{\Omega_2} \{w_k^n v - [(dt/2\mathbf{a}\nabla)w_k^n v + \nu dt/2 \nabla w_k^n \nabla v]\} \, dx + dt/2 \int_{\Gamma_{12}} \left(\mathcal{B}_2 v_{k+1}^n + \frac{\partial w_k^n}{\partial n} \right) v \, d\gamma \quad (13)$$

where v is any element of the P1 basis and $\mathcal{B}_2 u = \alpha_2 u + \frac{\partial u}{\partial n}$.

The Schwarz iterative procedure can be seen as a 4-block Gauss-Seidel method where two of the blocks contain the discretisation of the subdomains $(A_{ii})_{i=1,2}$ and the other two the coupling along the artificial boundaries $(A_{ij})_{i,j=1,2}$

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}. \quad (14)$$

At this point, two remarks can be made. 1) It is well known that any Krylov iterative technique is an improvement in terms of number of iterations over the Gauss-Seidel solver. 2) Krylov techniques imply storing residual vectors of size the number of points in all the domains, and can be quite storage consuming if applied directly to (14). Therefore the problem is transformed into a linear problem on the interface Γ_{12} .

For simplicity, we shall note $v_{k+1}^{n+1} = v$ and $w_{k+1}^{n+1} = w$.

$v = v_0 + \bar{v}$ and $w = w_0 + \bar{w}$ where v_0 and w_0 are solutions of the problems

$$\begin{cases} \mathcal{L}_1 v_0 = F_1, & \text{in } \Omega_1 \\ v_0 = v_1 & \text{on } \Gamma_1 \\ \frac{\partial v_0}{\partial n} = 0 & \text{on } \Gamma'_1 \\ \mathcal{B}_1 v_0 = 0 & \text{on } \Gamma_{12} \end{cases} \tag{15}$$

$$\begin{cases} \mathcal{L}_2 w_0 = F_2, & \text{in } \Omega_2 \\ \frac{\partial w_0}{\partial n} = 0 & \text{on } \Gamma_2 \\ \mathcal{B}_2 w_0 = 0 & \text{on } \Gamma_{12} \end{cases} \tag{16}$$

\bar{v} and \bar{w} are now solutions of

$$\begin{cases} \mathcal{L}_1 \bar{v} = 0, & \text{in } \Omega_1 \\ \bar{v} = 0 & \text{on } \Gamma_1 \\ \frac{\partial \bar{v}}{\partial n} = 0 & \text{on } \Gamma'_1 \\ \mathcal{B}_1 \bar{v} = \mathcal{B}_1 \bar{w} & \text{on } \Gamma_{12} \end{cases} \tag{17}$$

$$\begin{cases} \mathcal{L}_2 \bar{w} = 0, & \text{in } \Omega_2 \\ \frac{\partial \bar{w}}{\partial n} = 0 & \text{on } \Gamma_2 \\ \mathcal{B}_2 \bar{w} = \mathcal{B}_2 \bar{v} & \text{on } \Gamma_{12}. \end{cases} \tag{18}$$

Both (17) and (18) are linear problems with respect to the boundary conditions γ_1, γ_2 on Γ_{12} .

Determining the coupled values (v, w) can now be seen as solving the following 2-field linear operator $G(\gamma_1, \gamma_2) = -(\mathcal{B}_1(v_0 - w_0), \mathcal{B}_2(v_0 - w_0))$ where $G(\gamma_1, \gamma_2) = (\mathcal{B}_1(\bar{v} - \bar{w}), \mathcal{B}_2(\bar{v} - \bar{w}))$.

A Krylov technique applied to this problem will have residuals the size of the interface which is small compared to the whole domain.

This iterative solver has to be applied at each time step in the case of the steady coupling. A natural preconditioner is to project the initial solution v_0^{n+1} on former search vectors defined when computing v^n . This is easily implemented within the FOM solver (Full Orthogonalisation Method).

Remark. $\mathcal{B}_1(\bar{w})$ is computed with the EF code in the following way: imposing $\gamma_2 = \alpha_2 \bar{w} + \frac{\partial \bar{w}}{\partial n}$ provides after resolution \bar{w} , thus $\frac{\partial \bar{w}}{\partial n} = \gamma_2 - \alpha_2 \bar{w}$ and $\mathcal{B}_1(\bar{w}) = (\alpha_1 - \alpha_2) \bar{w} + \gamma_2$.

$\mathcal{B}_2(\bar{v})$ is computed using finite differences.

4.2.3. Unsteady coupling

Having reduced the problem to one on the interface, the unsteady coupling can be seen as $G(\gamma_1(t), \gamma_2(t)) = (\mathcal{B}_1(\bar{v}(t) - \bar{w}(t)), \mathcal{B}_2(\bar{v}(t) - \bar{w}(t)))$ where t is a vector (t_1, \dots, t_N) representing time discretisation of $[T_1, T_2]$ and residuals have size the interface dimension times N . Thus the unsteady coupling algorithm is identical to the steady one. The only difference is that instead of computing at each time step $G(\gamma_1(t_i), \gamma_2(t_i))$ implying small communications, the unsteady algorithm computes directly the block $G(\gamma_1(t), \gamma_2(t))$ every N time steps, thus a greater amount of data exchanged but less often.




FIGURE 12. Effect of the preconditioner.

Schematically:

the steady algorithm is:

```
do time iterations 1:N
do Krylov iterations
  residual vectors =
  size of interface
```

the unsteady algorithm is:

```
do Krylov iterations
do time iterations 1:N
  residual vectors =
  size of interface x N
```

4.3. Convergence results on the test case

Steady case

In Figure 12 convergence with and without preconditioning is presented for two time steps, situated when the heat blob is entering the finite element domain.

Unsteady case

As can be seen in Figure 13, the unsteady domain decomposition technique over a time window of 10 time steps converges in terms of number of iterations similarly to 10 preconditioned steady time steps. The unsteady coupling acts as a global preconditioner.

5. IMPLEMENTATION WITH CORBA

In this section, we describe how we implemented the coupling of two **Fortran** numerical applications. Legacy softwares are encapsulated within **C++** classes, “turned into servers” and immersed in a **Corba** object oriented client-server environment. Finally we can implement the coupling without porting and major restructuring of the original applications.

5.1. From Fortran codes to object servers

For our project two applications are involved, they *a priori* need to be ported locally and then coupled. Porting an application always leads to many annoying problems and one never knows what to do with the computations which are specific to the coupling.




FIGURE 13. Unsteady case.

apli.h	apli.c	main.c
<pre> class apli { private: // // application data ... public: // // prototypes // of methods ... compute (...); ... }; </pre>	<pre> #include <... .h> #include "apli.h" // // implementation // of methods apli::compute (...) { // call existing // Fortran routine ... } </pre>	<pre> #include <... .h> #include "apli.h" int main() { // // create an "apli object" // apli *A=new apli (...); // // solve a problem ! // ...A->compute(...); ... } </pre>

FIGURE 14. C++ encapsulation.

To overcome these problems we encapsulate the “user’s relevant” components of an application so that they become methods of an “object server”. This is done by designing a C++ class as described in Figure 14.

Following the methodology presented in Section 3.1, the methods are the services to be provided by the `apli` object server. Those methods match the subroutines called in the main Fortran program and a few simple requests for setting initial data, etc. The variables declared in the main Fortran program will have their counterpart encapsulated in the private part of a `apli` object. A client code will look like the body of the main Fortran program and invoke services on a reference to a `apli` object.

Then we derive a IDL specification from the C++ description of the class and the IDL compiler is used to generate C++ stub and skeleton used in Figure 6. Now, the implementation will inherit from the generated skeleton and will implement the functionalities of the object server for the given application as in Figure 5: it will be simply derived from the previous C++ implementation of `apli`.




FIGURE 15. Software components for coupling.

The client code will use the stub interface as the specification for the server and will also simply be derived from the C++ main program. It will invoke services on a reference to a remote `appli` object through the “proxy” as shown in Figure 4.

If we compare with the initial “monolithic” application, we are now able to solve the same problem by running the code through remote control, *i.e.* by sending requests to an object server. In this distributed context any user may write a client code looking like his usual “main” program and run it **locally**. The application will connect to a remote server which will carry over computations and solve the problem. No porting is to be done. Furthermore, we did not really have to “enter” the code: we just analysed the main program, identified application data and exhibited the relevant functionalities. An important point is that we used codes which were not originally developed for distributed simulations. By wrapping the existing subroutines in our servers we salvaged all the scientific knowledge imbedded in these applications.

Corba allows us to achieve interoperability. Our software construct is independent from the programming languages we use and from the operating systems and hardware the client and the server are running on.

5.2. Design of the coupling

As we are able to run the codes through remote control, it is now possible to undertake the coupling of applications provided we make available the basic functionalities needed to express interactions between them **in a client code**.

From a technical point of view, many software architectures can be foreseen for this system, where all additional computations are taken over by the client or by another secondary module. Two typical software architectures are shown in Figure 15. In the first diagram the client controls the two applications and does all the additional work. In the second diagram, the client just controls the processes while additional work is delegated to a “coupling module”.

If we were able to build a “real software bus” (*cf.* Fig. 3), *i.e.* work with a software component model, any software architecture for the system would be transparent to the user. Unfortunately, this is not really the case with current implementations of Corba which do not provide enough abstraction for that: clearly here codes for servers would be different **but** encapsulations would not be affected. For this test case we will keep to the first diagram. The new Corba component model (CCM) recently released by the OMG will provide more flexibility in the design of distributed applications [23].

To undertake this coupling, *i.e.* to implement the algorithms developed in Section 4 we need to add new functionalities –services– to our servers:

- identify boundaries;
- extract the values of a given unknown on a given boundary;
- extract the values of time or space derivatives of a given unknown on a given boundary;
- modify the values of a given unknown on a given boundary;
- access various scalar values (coefficients for optimal convergence, ...).

Our servers need to be upgraded as they must include the new “coupling functionalities” but these features are rather simple, “generic” and easy to implement in **Fortran** in the initial code. Thanks to those functionalities, data can be transferred transparently from one application to any other hence interpolations and additional computations can be performed by any module.

The coupling will be entirely expressed in the new client code in which we will monitor both applications and perform the additional computations.

Just to give a flavour of the corresponding IDL **specification**, additional functionalities for both servers will look like:

```

// ...
// here, "fields" are arrays of double
// borders are identified by "border_id"
// ...
// access values on a boundary
//
void get_u      (in long border_id, out field Uo );
void get_du_dn (in long border_id, out field Uo );
//
// get coefficients for optimal convergence
//
void get_alpha (out double alpha                );
void get_beta  (out double beta                 );
//
// update values on a boundary
//
void update_u  (in long border_id, in  field Ui );
// ...

```

FIGURE 16. IDL upgrades for servers.

5.3. Handling parallelism

Obviously, in this new client code, many things can be done in parallel. For example, each server can be monitored inside a **thread** (with the suitable synchronizations. . .) to implement the inter-application parallelism level in the client code. Multi threaded implementations of **Corba** generally offer a “java-like” abstraction of threads through a **C++** thread class. Also in **C++**, multiple inheritance allows to take care of thread parallelism in the object server itself.

On the other hand, intra application parallelism must be handled too. Here the original codes are parallel applications based on the **MPI** library [24, 25]. At the **Corba** level –from the client point of view– we do not know that they are parallel applications and cannot take benefit from it. The consequence is that, at some step, exchanges are not performed among processes of the parallel applications but sequentialization has to occur since only one process of each application can dialog through **Corba** as shown in Figure 17. Another problem

in Figure 17 is that MPI applications are no longer SPMD codes but should be based on a master/slave model, the node 0 obviously playing a special part.

Some research projects try to take care of this flaw as in COBRA [26] or PARDIS [27] where the IDL language is extended to handle “distributed” data types.

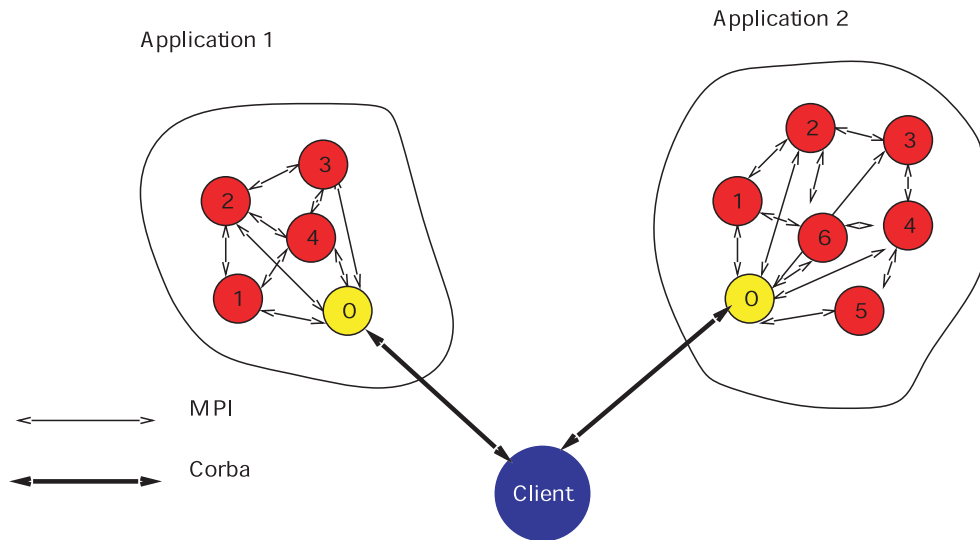


FIGURE 17. A Corba client monitoring Mpi-based applications.

Anyway a good practice in distributed and parallel programming is to try to reduce communications by improving the “locality” of the algorithms. Clearly the overall performance is tightly connected to the characteristics of the network –bandwidth of course but also, in this context, “startup” time– and the availability of an optimized environment specific protocol for Corba.

5.4. Conclusions

Once a server is developed, only a few “generic” additional features are necessary to undertake coupling with another application. Implementation of these new functionalities is in most cases straightforward. Servers can be sequential or parallel applications. If legacy software is encapsulated its original performances are not affected.

The implementation of the coupling itself remains external to the original applications. It is totally contained in the client code (or another secondary module also controlled by the client as in Fig. 15). Inter application parallelism is also handled in the client code.

Availability of implementations of new paradigms: component model, “parallel objects”, will allow more ease and flexibility in designing and implementing distributed parallel applications.

6. NUMERICAL RESULTS AND OPENING REMARKS

In Figure 18, the evolution of the heat blob around the airfoil in the unsteady case (the results are similar to the steady case) with a window of 10 time steps ($\delta t = .05$) is shown. Coupling is done through Corba. As can be seen, going through the finite element border does not interfere with the heat blob convection, thus validating in this case the numerical algorithm.

For the first 5 time windows, convergence takes under a 100 iterations, while when the blob leaves the finite element domain, number of iterations is reduced by half. The whole computation takes under 5 minutes, validating the coupling prototype.




FIGURE 18. Time evolution of a heat blob around a Naca airfoil.

To conclude, several remarks can be made:

First, one of the key points of the domain decomposition method presented here is its great pliability in terms of border exchanges and local computation, thus allowing better adaptability to the hardware. By adding the time dimension, communications can be adjusted to latency and bandwidth, independently of the space size of the domains. It is particularly well-suited to distributed computing. Connections with high latency and large bandwidth will be most effective with large time windows, and data transfers can be overlapped by local computations. Low latency and small bandwidth systems can work on a the steady coupling (a one time step window) without modifying the precision of the results.

Secondly, this work has shown the feasibility of coupling heterogeneous methods on distant sites. The software technology used is simple and provides immediate access to the codes. In this context the coupling is non-intrusive and is described in an external module.

Finally, such techniques allow fast prototyping of new algorithms and methods and is a step towards total flexibility in CFD and other fields of numerical computations.

REFERENCES

- [1] M. Gander and L. Halpern, Méthodes de relaxation d'ondes (SWR) pour l'équation de la chaleur en dimension 1 (submitted to CRAS).
- [2] M.J. Gander, L. Halpern and F. Nataf, Optimal convergence for overlapping and non-overlapping Schwarz waveform relaxation, in *Eleventh international Conference of Domain Decomposition Methods*, C.-H. Lai, P. Bjørstad, M. Cross and O. Widlund Eds. (1999).
- [3] C. Japhet, *Conditions aux limites artificielles et décomposition de domaine : méthode 002 (optimisée d'ordre 2). Application à la résolution de problèmes en mécanique des fluides*. Thèse, École Polytechnique (1997).
- [4] C. Japhet, F. Nataf and F. Rogier, The Optimized Order 2 method. Application to convection-diffusion problems. *Future Generation Computer Systems FUTURE* **18** (2001).
- [5] F. Nataf, communication personnelle.
- [6] M. Grand, *Java Language Reference*. 2nd Edition, O'Reilly (1997), ISBN 1-56592-326-X.
- [7] *The Java Programming Language*, <http://java.sun.com/>.
- [8] *Java Grande Forum*, information at <http://www.javagrande.org>
- [9] *Java Numerics*, information at <http://math.nist.gov/javanumerics>
- [10] J. Farley, *Java Distributed Computing*. O'Reilly (1998), ISBN 1-565-92206-9.
- [11] GRID, *The GRID Forum*, <http://www.gridforum.org/>.
- [12] EGRID, *The European Grid Forum*, <http://www.egrid.org>
- [13] A. Chervenak, I. Foster, C. Kesselman, C. Salisbury and S. Tuecke, The Data Grid: Towards an Architecture for the Distributed Management and Analysis of Large Scientific Datasets. Available on line at [14]. *J. Network Comput. Appl.* **23** (2001) 187–200.
- [14] *The GLOBUS project*, information at <http://www.globus.org/>.
- [15] S. Chapin, J. Karpovich and A. Grimshaw, The Legion Resource Management System, in *Proc. of the 5th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP'99)*. San Juan, Porto Rico (1999).
- [16] *The LEGION project at the University of Virginia, USA*, <http://legion.virginia.edu/>.
- [17] J. Siegel et al., *CORBA Fundamentals and Programming*. J. Wiley & Sons (1996), ISBN 0-471-12148-7.
- [18] M. Henning and S. Vinoski, *Advanced CORBA Programming with C++*. Addison-Wesley (1999), ISBN 0201379279.
- [19] *Corba: Common Object Request Broker Architecture*, information at <http://www.corba.org>
- [20] *OMG the Object Management Group*, <http://www.omg.org>
- [21] B. Stroustrup, *The C++ programming language*. 3rd Edition, Addison-Wesley (1998), ISBN 0-201-88954-4.
- [22] J. Barton and L. Nackman, *Scientific and Engineering C++*. Addison-Wesley (1994).
- [23] *OMG CCM Implementers Group, CORBA Component Model Tutorial*, Document-ccm/02-04-01, available at <http://www.omg.org>
- [24] W. Gropp, E. Lusk and A. Skjellum, *Using MPI*. 2nd Edition, MIT Press (1999), ISBN 0-262-57132-3.
- [25] *MPI: Message Passing Interface*, all documents can be retrieved from <http://www.mpi-forum.org/>. For information and implementations of the standard see <http://www-unix.mcs.anl.gov/mpi/>.
- [26] T. Priol, C. René and G. Alléon, *Programming SCI Clusters Using Parallel CORBA Objects*. INRIA-IRISA Report 3649 (1999).
- [27] K. Keahey and D. Gannon, PARDIS: CORBA-based Architecture for Application-level Parallel Distributed Computation, in *Proc. of Supercomputing'97* (1997).