

# DOCUMENTOS DE TRABAJO

# BILTOKI

D.T. 2012.01

MPI parallel programming of mixed integer optimization problem  
using CPLEX with COIN-OR

Unai Aldasoro, M.Araceli Garín, María Merino y Gloria Pérez

eman ta zabal zazu



Universidad Euskal Herriko  
del País Vasco Unibertsitatea

Facultad de Ciencias Económicas.  
Avda. Lehendakari Aguirre, 83  
48015 BILBAO.

**Documento de Trabajo BILTOKI DT2012.01**

Editado por el Departamento de Economía Aplicada III (Econometría y Estadística)  
de la Universidad del País Vasco.

ISSN: 1134-8984

# MPI parallel programming of mixed integer optimization problems using CPLEX within COIN-OR

Unai Aldasoro<sup>1</sup>, María Araceli Garín<sup>2</sup>, María Merino<sup>3</sup> and Gloria Pérez<sup>3</sup>

<sup>1</sup>Dpto. Matemática Aplicada,  
Universidad del País Vasco, UPV/EHU  
Bilbao (Bizkaia), Spain  
e-mail: unai.aldasoro@ehu.es

<sup>2</sup>Dpto. de Economía Aplicada III,  
Universidad del País Vasco, UPV/EHU  
Bilbao (Bizkaia), Spain  
e-mail: mariaaraceli.garin@ehu.es

<sup>3</sup>Dpto. de Matemática Aplicada, Estadística e Investigación Operativa ,  
Universidad del País Vasco, UPV/EHU  
Leioa (Bizkaia), Spain  
e-mail: maria.merino@ehu.es, gloria.perez@ehu.es

## Abstract

The aim of this technical report is to present some detailed explanations in order to help to understand and use the Message Passing Interface (MPI) parallel programming for solving several mixed integer optimization problems. We have developed a C++ experimental code that uses the IBM ILOG CPLEX optimizer within the COmputational INfrastructure for Operations Research (*COIN-OR*) and MPI parallel computing for solving the optimization models under UNIX-like systems. The computational experience illustrates how can we solve 44 optimization problems which are asymmetric with respect to the number of integer and continuous variables and the number of constraints. We also report a comparative with the speedup and efficiency of several strategies implemented for some available number of threads.

**Keywords:** Optimization, Message Passing Interface, Parallel Computing, COIN-OR Open Solver Interface, CPLEX optimizer.

---

\*This research has been partially supported by the project ECO2008-00777 ECON from the Spanish Ministry of Education and Science, and Grupo de Investigación IT-347-10 from the Basque Government.

# 1 Introduction

*Parallel computing* is a Computer Science area that studies the necessary hardware and software aspects to perform simultaneous execution of tasks. Currently this discipline is the prominent paradigm on high performance computing and also on computer architecture due to the industry's shift to multicore processors.

At hardware level, the parallel era starts at the late 1950's in the form of shared memory multiprocessor supercomputers. The development of this type of computers continued until the early 1980s when a new paradigm of massively parallel multiprocessors (MPPs) arrived. MPPs showed a significantly better performance and became dominant of high performance computing. From late 1980's on computing clusters appear by linking a large numbers of off-the-shelf computers connected by an off-the-shelf network. Nowadays parallel computing is mainly based on clusters and multicore processors. For deeper information about parallel computing hardware consider [3] and [6]. The top 500 website gathers every six months the information concerning the top 500 super-computing sites at world level .

This cooperation among processors can be of different nature depending on the way processors exchange information. The basic parallel architectures correspond to *shared memory* and *message passing* (or distributed memory). As described on [8] "one processor of a shared memory machine can communicate with another by writing the information into a global shared memory location and having the second processor read directly from that location" using a bus. So the communication is carried out by *shared variables*. This makes inter-processor communication very easy and fast but produces problems in terms of simultaneous access of a unique memory location. On the other hand on message-passing paradigm each processor has its own local memory connected by a network with the rest of processors, the communication is based on message-passing.

On mathematical optimization, we often find large-scale problems or we need to solve many of them, which can not be done efficiently by a single processor, even if technology is constantly improving. Parallel computing aims to be a powerful tool to manage this type of problems, operating on the principle that large problems can often be divided into smaller ones. The idea is simple: to use  $p$  processors in cooperation in order to (ideally) be able to solve a problem, to solve a problem  $p$  times faster or to solve a  $p$  times bigger problem using the same amount of time.

In this technical report we show how to solve many problems by a message-passing paradigm code, because the solving processor of each model is completely independent from the rest of them, so the transfer of information is not significant. The computational experience will be carried out in




Figure 1: Shared memory and distributed memory architectures (P: processor, M: memory)

the cluster ARINA (see [1]), that provides an extensive environment for message-passing architecture.

The presented code have been programmed with the MPI library. As described on [10] “the Message-Passing Interface or MPI is a library of functions and macros that can be used in C, FORTRAN and C++ programs. As its name implies MPI is intended for use in programs that exploit the existence of multiple processors by message-passing”. Developed in 1993 is one of the first standards for parallel programming and it is the first that is based on message-passing.

This work aims to present a general environment for parallel solving of several optimization problems using directly the IBM ILOG CPLEX optimizer (see [7]) within the COmputational IN-frastructure for Operations Research (see [2] and [9]). The main objective of this work is to describe the MPI parallelization for solving optimization problems. In the computational experience, we show the evolution of the computing time depending on the number of processors used and the selected solver, where the combination of a general MPI parallelization with the internal ILOG CPLEX parallelization has been considered.

The remainder of the technical report is organized as follows. In Section 2 appears the context of the problem that we are interested in solving. Section 3 shows the main MPI syntaxis for parallel computing and Section 4 shows the paralleling strategy. The main ideas for MPI execution under UNIX-like systems are presented in Section 5. Section 6 reports the computational experience using COIN-OR and CPLEX. In the Appendix 1 we present the C++ code in detail for the illustrative example. An example of Makefile for compiling C++ code with MPI environment is given in the Appendix 2.

## 2 Context

Let us consider the following mixed 0-1 optimization problem

$$\begin{aligned}
 z = \min \quad & ax + by \\
 \text{s.t.} \quad & Ax + By = h \\
 & x \in \{0, 1\}^{nx}, \\
 & y \in \mathbb{R}^{ny},
 \end{aligned} \tag{1}$$

where  $x$  and  $y$  are the  $nx$  and  $ny$  dimensional vectors of the 0-1 and continuous variables,  $a$  and  $b$  are the vectors of the objective function coefficients,  $A$  and  $B$  are the constraint matrices, respectively, and  $h$  the right-hand-side.

As we have mentioned previously, we can consider two situations: to solve a large-scale optimization problem, that can be decomposed in  $q$  smaller subproblems or to solve a huge number  $q$  of independent mixed-integer problems. In any case, how can we solve  $z^1, z^2, \dots, z^q$  in non-sequential order?

We will use the MPI parallel programming for solving  $q$  optimization problems, that will be read in Mathematical Programming System (MPS), which is a standard format for linear and mixed integer optimization. We will compare the COIN-OR optimization solver and the CPLEX optimizer under COIN-OR, with sequential and parallel programming. It is important to notice that the ILOG

CPLEX optimizer also includes internally a parallel environment for solving an optimization model. The number of parallel threads (on shared memory) that CPLEX may use for any invocation of a parallel algorithm is controlled by the threads parameter CPX\_PARAM\_THREADS.



Figure 2: Strategies: MPI threads with CPLEX threads (8 threads available for solving 44 problems)

If we consider cores with 8 threads, we can solve the  $q$  problems in different ways depending on the number of threads for distributed memory (MPI threads) and the number of threads for shared memory (CPLEX threads). For example, we will consider  $q = 44$  problems, so-called  $P_1, P_2, \dots, P_{44}$ , we could use 8 MPI threads and sequential CPLEX; or 4 MPI threads and 2 CPLEX threads; or 2 MPI threads and 4 CPLEX threads; or 1 MPI thread and 8 CPLEX threads, among others. We illustrate the different executions in the Figure 2.

It is important to notice that not all thread have the same role. We denote **primary threads** to those who execute a copy of the executable, that is, the threads working on a distributed memory

paradigm and linked by MPI functions. One of the primary threads will be denoted **coordinator thread**, this specific thread will be the one who gather the distributed information. On the other hand **auxiliary threads** will take part on the process only when the CPLEX optimizer requires their participation for a shared memory solving. Each auxiliary thread is linked to a specific primary thread. The corresponding notation will be applied at the different diagrams of this technical report and it is summarized on the following lines.

Considering  $n + 1$  MPI threads and  $m + 1$  CPLEX threads:

Coordinator thread	<b>0.0</b>	
Primary threads	<b>i.0</b>	$i = 0, \dots, n$
Auxiliary threads	<b>i.j</b>	$i = 0, \dots, n, j = 1, \dots, m$

### 3 MPI sintaxis

Let us start the description of the MPI environment by underlining the three main phases that appear at machine level on a message-passing paradigm, see [10]. These phases are not directly executed by the user, they are automatically performed by the machine.

**Phase a:** The user issues a directive to the operating system which has the effect of placing a copy of the executable program on each processor.

**Phase b:** Each processor begins execution of its copy of the executable.

**Phase c:** Different processors can execute different statements by branching within the program. Typically branching will be based on processor ranks.

The main idea behind these phases is that every primary thread receives an exact copy of the executable. The code is executed in parallel but we would like different primary threads to execute different statements. This is achieved by a rank based branching, in other words, defining the rank of the processor that will perform a specific statement or by branching data vectors by rank. This allows to work on a Single Program Multiple Data paradigm.

The implementation of the corresponding executable supports different experimental codes. On the following lines we present the general structure of a MPI implementation in C++ (see [10]). This layout can be divided in five groups by considering the nature of the functions they use.

```

...
#include "mpi.h"
...
main(int argc, char **argv) {
...
    // Group 1: Declaring MPI variables.
    ...
    // Group 2: Beginning of the MPI environment (No MPI functions called before this).

```

```

...
// Group 3: Functions controlling the number of processes.
...
// Group 4: Communication functions.
...
// Group 5: End of the MPI environment (No MPI functions called after this).
...
...
}
...

```

The following subsections will describe the objective of each group and will use as practical example sentences of the C++ code added as Appendix 1.

### 3.1 Group 1: Declaring MPI variables

In this part the specific variables related to the MPI environment are declared. This section may not appear in every MPI code since a big part of MPI functions have as input or output common C++ variables.

Among the most frequent MPI variables we find **MPI\_Group** and **MPI\_Comm**. The first one creates *groups* of processor whereas the second creates *communicators*, that is, a group of processors plus a context of communication, in other words, a collection of processors that can send messages to each other.

By default, a *communicator* called MPI\_COMM\_WORLD is created at the beginning of each program and it consists of all the active processors at the execution. This turns very useful for global communications but frequently it is simpler to consider also subgroups of processors. That is the case of the considered code.

As it is shown at the Appendix 1 the following MPI variables are declared at *mainmpi.cpp*:

```

MPI_Group  orig_group,new_group;
MPI_Comm   new_comm;

```

For other MPI variables consider [10] and [12]

### 3.2 Group 2: Beginning of a MPI environment

```

MPI_Init(&argc,&argv);

```

**MPI\_Init** corresponds to the first MPI function that must be executed on a program. It allows the system to use the MPI library and therefore its features. It must be called just once and no MPI functions can be called before this.



### 3.3 Group 3: Functions controlling the number of processors

In order to have a functional point-to-point communication environment it is essential to identify each processor with a *rank* number, so that we can easily specify the sender and the receiver. Additionally, the *rank* allows the programmer to work on a *Single Program Multiple Data* (SPMD) paradigm, that implies that with a single program different processor execute different tasks. Schematically:

```
if (rank == 0) DO X
else if (rank == 1) DO Y
```

To do so, the MPI library provides the **MPI\_Comm\_rank** function:

```
MPI_Comm_rank(MPI_COMM_WORLD, &original_rank);
```

The first argument defines the *communicator* and the corresponding *rank* of the processor is stored on the second argument. Note that a processor has *rank*, in general different, for each *communicator* that is involved on.

Frequently, the tasks to be executed are divided among the available processors. The function **MPI\_Comm\_size** allows us to determine the number of processors in a communicator (stored on the second argument):

```
MPI_Comm_size(MPI_COMM_WORLD, &original_size);
```

How can we create a new communicator considering only active processors?

Let us consider the case where the number of processors is bigger than the number of submodels to solve, that is  $nmodel < original\_size$ . That will mean that some of the processor will not actually solve any subproblem. In order to simplify the processor of gathering the final solutions, we have created a new group considering only the active threads. The necessary steps will be:

Extract handle of global group from MPI\_COMM\_WORLD using **MPI\_Comm\_group**. The handle is stored on the second argument.

```
MPI_Comm_group(MPI_COMM_WORLD, &orig_group);
```

Form a new group as a subset of global group using **MPI\_Group\_incl**.

```
if (nmodel < original_size) {
MPI_Group_incl(orig_group, nmodel, ranks1, &new_group);
} else {
MPI_Group_incl(orig_group, original_size, ranks2, &new_group);
}
```

*ranks1* is an array of *nmodel* elements containing the *ranks* of the processors to be part of the new group. The handle of the new group is stored on *&new\_group*. On the other hand, *ranks2*

contains the ranks of the *MPI\_COMM\_WORLD*, so notice that if *nmodel*  $\neq$  *original\_size* the new group will correspond to the existing *MPI\_COMM\_WORLD* group.

Let us create a new communicator for the new group using **MPI\_Comm\_create**. This function creates the new communicator *&new\_comm* with the processor *new\_group* extracted from communicator *MPI\_COMM\_WORLD*.

```
MPI_Comm_create(MPI_COMM_WORLD, new_group, &new_comm);
```

### 3.4 Group 4: Communication functions

Message-passing communication is the core of a MPI environment. The basic functions to develop such an exchange correspond to *MPI\_Send* and *MPI\_Recv*. The first function sends a message to a designated processor whereas the second receives a message from a processor. As mentioned on [10] in order for the message to be successfully communicated the system must append some information to the data that the application program wishes to transmit. This information is called the **envelope of a MPI message**. It will contain the following basic information:

- The rank of the receiver
- The rank of the sender
- A tag
- A communicator

Every message-passing function defines as arguments the previous aspects. As example we will consider the specific functions used at the Appendix 1 code: **MPI\_Gatherv** and **MPI\_Allgatherv**.

Analysing the parallel computing strategy we will see that each submodel is assigned to a processor. After the solving task is finished we will gather the value of the objective function at the submodel ranked 0. The corresponding function will be defined as follows:

```
MPI_Gatherv(&zq_loc[0], assignment[pid], MPI_INT,  
&zq[0],assignment, inicial, MPI_INT, 0, new_comm);
```

Where **MPI\_Gatherv** is used to gather information of type vector in a specific processor (for scalar values use *MPI\_Gather*). The first three arguments define the vector to be sent, that is, a vector stored from *&zq\_loc[0]* of length *assignment[pid]* and of type *MPI\_INT* (in other words, a vector of integer values). The processor ranked 0 at the *new\_comm* communicator will receive a vector of type *MPI\_INT* from each processor of the same communicator. The size of the received vectors is defined at the vector *assignment* and their position at the new vector *&zq[0]* is established by the vector *inicial*.

Additionally we want all the processors to have access to the presolve status. To do so we will use the **MPI\_Allgatherv**. This function is an extension of **MPI\_Gatherv**, since the same structure is repeated for every processor, in other words, the gather vector will be stored in every processor. Note that the only difference at argument level is that this function does not ask to define the receiver rank.

```
MPI_Allgatherv(&istruef969_loc[0], assignment[pid], MPI_INT,
&istruef969[0],assignment, inicial, MPI_INT, new_comm);
```

Notice that the Appendix 1 code also includes the **MPI\_Allreduce** function. It has a double nature since it combines values from all processors and distributes the result back to all processors. In this particular case all processors will store at the *ncols* variable the value obtained by taking the maximum value (MPI\_MAX) among the local *ncols\_max\_loc* variables.

```
MPI_Allreduce(&ncols_max_loc,&ncols, 1, MPI_INT, MPI_MAX,new_comm);
```

The MPI library provides a rich variety of functions for message-passing tasks. For further information see [12].

### 3.5 Group 5: End of a MPI environment

```
MPI_Finalize();
```

**MPI\_Finalize** corresponds to the last MPI function that must be executed on a program. It cleans up unfinished tasks and closes the interaction with the MPI library. It must be called just once and no MPI functions can be called after this.

## 4 Parallelization Strategy

Before analysing the parallelization strategy, let us consider the serial version of the solving process. Figure 3 summarizes on a descriptive diagram the four tasks to be performed by the only thread. These tasks correspond to a data input, model creation, solving process and printing the numerical results. Each of them is applied to every single problem before starting the next step.

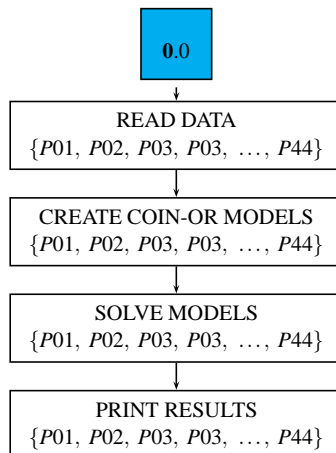


Figure 3: Basic steps of the serial execution

We can extend the descriptive diagram to a two thread parallel version as shown in Figure 4. In this case the first three tasks are performed by both threads sharing the total amount of problems, this will significantly reduce the total amount of time. Once the solving process is finished on both threads, the coordinator thread gathers the numerical results by massing passing functions and prints them.

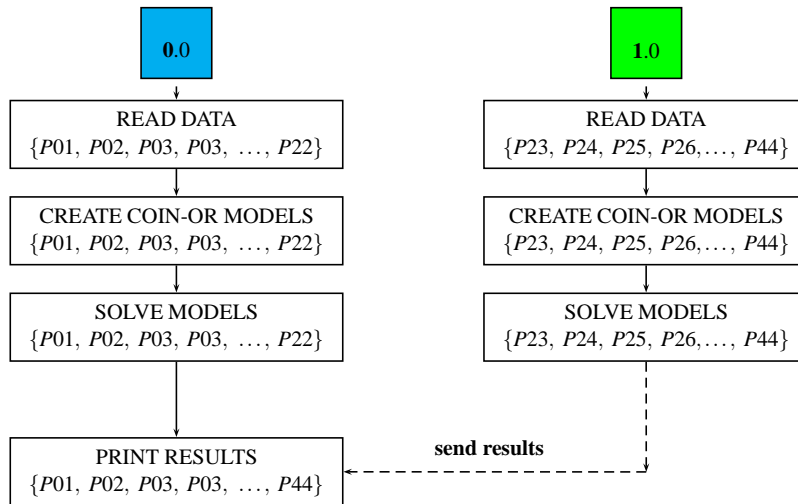


Figure 4: Basic steps of the MPI parallel execution

The user can define the branching cited on the third block according to the desired parallel programming strategy. The following procedure shows the parallel computing idea behind the example code provided in the Appendix 1 and a parallel diagram is illustrated in Figure 5.

**Step 0: Declaring optimization and MPI variables.** [All primary processors]

The correspondence at variable level between the optimization model (1) and appended C++ code is described at Table 1. Declare MPI variables by using functions of Group 1 described at Subsection 3.1.

Table 1: Correspondence between the optimization model and the implemented C++ code

Model	q	z	(x,y)	nx+ny
C++	nmodel	zq	x0	ncols

**Step 1: Definition of the global environment.** [All primary processors]

★ Beginning of the MPI environment using functions from Group 2, Subsection 3.2

Definition of the total number of processors, the rank of each of them and creation of a new communicator by using functions of Group 3 described at Subsection 3.3. By default, processor ranked 0 is defined as coordinator.

**Step 2: Presolve assigned models.** [All primary and auxiliary processors]

Every processor reads the corresponding MPS files and creates the associated CPLEX within COIN-OR optimization models.

- If selected optimizer = CPLEX: Each processor starts a shared memory parallel computing environment with a chosen number of auxiliary processors. In Figure 5 we show an illustrative diagram that corresponds to use 2 MPI threads and 4 CPLEX threads. Assigned models are presolved.
- If selected optimizer = COIN-OR: Each processor presolves assigned models.

**Step 3: Global MPI communication.** [All primary processors]

The coordinator processor gathers the presolve information of all models and all primary processors gather presolve status by using functions from Group 4 described at Subsection 3.4

**Step 4: Check the presolve status** [Coordinator processor]

- If all models are feasible and bounded: Go to Step 5.
- Else: Go to Step 8.

**Step 5: Solve assigned models.** [All primary and auxiliary processors]

- If selected optimizer = CPLEX: Each processor starts a shared memory parallel programming environment with a chosen number of auxiliary processors (see Figure 5). Assigned models are solved.
- If selected optimizer = COIN-OR: Each processor solves assigned models.

**Step 6: Global MPI communication.** [All primary processors]

The coordinator processor gathers the presolve information of all models by using functions form Group 4 described at Subsection 3.4

**Step 7: Print results** [Coordinator processor]

The coordinator processor prints all results.

- ★ End of the MPI environment using functions from Group 5 described at Subsection 3.5

**Step 8: Execution ends in all processors.** [All primary processors]

## 5 Basic installations and executions under UNIX-like systems

### 5.1 Basic installations: COIN-OR, CPLEX, C++, MPI

Firstly, its necessary to install COIN-OR, which stands for **CO**mputational **IN**frastructure for **O**perations **R**esearch, a collection of open source software for optimization, see [2].

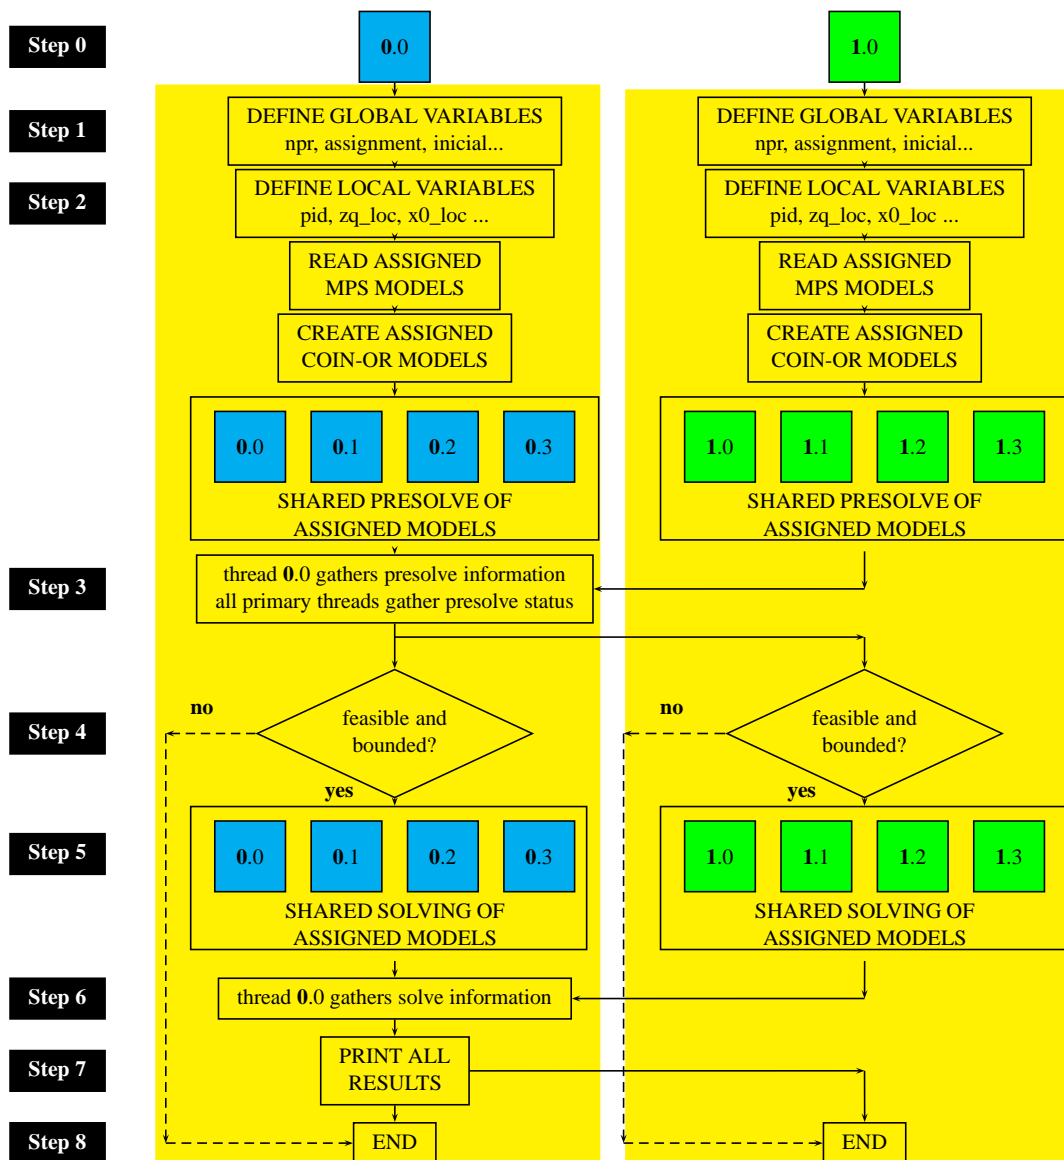


Figure 5: Parallelization diagram (2 MPI threads  $\times$  4 CPLEX threads)

The first step is to download the codes. For downloading *COIN-OR*, you must click on [Download/Use](http://www.coin-or.org) in the left hand side of the home page <http://www.coin-or.org>. Then in the second Section titled Source Code, you must click on [here](#) to download the source code for the latest stable release. You can observe an index for the source of a number of COIN projects. You must click on [CoinAll/](#) to obtain the list of last versions. In this case you will click and select [CoinAll-1.3.1.zip](#). Alternatively you can go directly to the corresponding home page [CoinAll-1.3.1.zip](#). After downloading the tarball you must extract the code into a new subdirectory, `/software/CoinAll`. For a step by step explanation on how to download and use COIN-OR consider [11].

Secondly, you must install the IBM ILOG CPLEX software for optimization with academic license, see [7]. For downloading CPLEX, you must click on

[www-01.ibm.com/software/websphere/products/optimization/academic-initiative/index.html](http://www-01.ibm.com/software/websphere/products/optimization/academic-initiative/index.html)

You have to register "join now" in the right of screen and go to step 2. You will see the screen, IBM Academic Initiative (the ILOG Optimization Key Request) page, or click on

[https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?lang=en\\_US&source=ai-ILOGKEY](https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?lang=en_US&source=ai-ILOGKEY), give your institutional (for example, UPV/EHU) e-mail address. When you are registered click on

[academic-initiative](#) you can do "Get Full version software". Select "download from the software catalogue", give your institutional e-mail address and password, click by email. In **Find by search text** put CPLEX and click on **Search** you have a list of programs, *IBM ILOG CPLEX Optimization Studio Academic Research Edition V12.2*, for several platforms. Select one of them and click on "I agree" at the end of the screen and **Download now**. Select the directory to install `/software/C-PLEX`. Also, save the license key file [access.ilm](#) in the location of your choice, for downloading the license key, you can follow the instructions of [ILOGQuickStart.pdf](#), to obtain this file, please click on [ILOGQuickStart.pdf](#). For a step by step explanation on how to download and use CPLEX within COIN-OR consider [11]. Finally, you can install the mpic++ compiler and the [mpi.h](#) library for MPI parallel computing from the webpage of GNU <http://gcc.gnu.org/> (see [5]).

## 5.2 Basic executions

Once you have downloaded and installed all the software described in the previous subsection, you can compile, link and execute your own C++ code with the previous ones.

For the compilation in the machine `guinness.lgp.ehu.es` of the so-called `mainmpi.cpp` file provided in the Appendix 1 of this technical report, you will need a script that contains all the sentences to compile the C++ file under MPI, linking with the COIN-OR libraries and the optimization software CPLEX. We will use the following script so-called `compile`, which has been written by the technicians of the SGI/IZO-SGIker Computing Service at the UPV/EHU, see [1]. For an alternative unix machine, you can consult the Appendix 2.

```

1  #!/bin/bash
   [ $# != 1 ] && echo -e "\n Provide one file name\n" && exit
3  #compiler=g++
   #compiler=icpc
5  compiler=mpicpc
   extra_empt64=""
7  if [ $compiler == "g++" ]; then
   gnu_CXXFLAGS="-O3 -fomit-frame-pointer -pipe -fexceptions -DIL_STD -DNDEBUG"
9  gnu_warnings="-pedantic-errors -Wimplicit -Wparentheses -Wreturn-type -Wcast-
   qual -Wall -Wpointer-arith -Wwrite-strings -Wconversion -Wno-unknown-
   pragmas"
   else
11 gnu_CXXFLAGS="-O3 -fomit-frame-pointer -pipe -fexceptions -DIL_STD -DNDEBUG"
   [ $(arch) == "x86_64" ] && extra_empt64="-Wparentheses "
13 gnu_warnings="$extra_empt64 -Wreturn-type -Wcast-qual -Wall -Wpointer-arith -
   Wwrite-strings -Wconversion -Wno-unknown-pragmas "
   fi
15 coin_dir="/software/CoinAll"
   coin_libdir="$coin_dir/lib"
17 coin_liblink="-L$coin_libdir"
   #all libraries of coin
19 coin_libs=$(for i in $(ls $coin_libdir/lib*.la);do basename $i .la;done |sed 's
   /lib/-l/g' |tr -s "\n" " ")
   coin_incs="-I$coin_dir/include/coin"
21 cplex_incs="-I/software/CPLEX/cplex/include/ilcplex/"
   mkl_libs="-L/opt/intel/composerxe/mkl/lib/intel64/ -lmkl_intel_lp64 -
   lmkl_sequential -lmkl_core -lpthread"

```

```

23 cplex_libs="-L/software/CPLEX/cplex/lib/x86-64_sles10_4.1/static_pic/ -lcplex -
    lilocplex "
    [ $(arch) == "ia64" ] && cplex_libs=""
25 [ -f $_$(arch) ] && echo -e "\n ===== Deleting old $_$(arch)\n" && rm
    $_$(arch)
    $compiler $gnu_CXXFLAGS $gnu_warnings $coin_incs $cplex_incs $coin_liblink
    $coin_libs $cplex_libs $mkl_libs $_.cpp -o $_$(arch)
27 echo ""
    if [ -f $_$(arch) ] ;then
29     echo ===== Done!!, $_$(arch) binary has been created
    else
31     echo ===== ERROR!!, $_$(arch) binary could not be created
    fi
33 echo ""

```

compile.dat

If you import this script from non UNIX-like systems, you should edit it

```
$ vi compile
```

and write

```
Esc : set fileformat=unix
```

```
Esc : wq
```

copy compile.dat in compile

```
$ cp compile.dat compile
```

Then, you can give execution permission with the following sentence:

```
$ chmod u+x compile
```

Now, you can execute the main C++ file with the script described previously:

```
$ ./compile mainmpi
```

The binary will be created with the mainmpi\_x86\_64 name and it can be executed from the prompt with the command mpirun or mpiexec for mpi execution with 8 threads:

```
$ mpirun -n 8 ./mainmpi_x86_64
```

Or alternatively, the binary can be sent to the queue of the system using a script\_file with the specific sentences to determine, among others,

- the maximum real time. For reserving 12 hours: #PBS -l walltime=06:00:00
- the memory allowed. For reserving 1 gb: #PBS -l mem=40gb
- the number of threads. For reserving 8 threads: #PBS -l nodes=1:ppn=8:xeon8

You can send the executable file to the queue with the sentence:

```
$ qsub_arina script_file
```

For more information about sending files to ARINA, you can see

<http://www.ehu.es/sgi/category/mandar-trabajos>.



## 6 Computational experience

The proposed main program has been implemented in a C++ experimental code. It uses the open source optimization engine *COIN-OR* for solving the mixed 0-1 optimization problems, in particular, we have used the functions: Clp (LP solver), Cbc (MIP solver), Osi, OsiClp, OsiCbc, OsiCpx and CoinUtils. Additionally, it uses one of the state-of-the-art commercial optimization engines, in particular CPLEX, see [7], within the open source engine *COIN-OR*.

The computational experiments were conducted at the ARINA computational cluster provided by the SGI/IZO-SGIker at the UPV/EHU. ARINA provides 1400 cores divided as follows: 1112 xeon cores, 248 Itanium2 cores and 40 opteron cores. All calculation nodes are connected by an Infiniband network with high bandwidth and low latency. For the present experiments the xeon x86\_64 architecture (Xeon Nehalem-EP E5520 @ 2.27GHz) type nodes have been used, consisting on 8 cores with 24 Gb of RAM with an QDR infiniband interconnection. Whereas for the calculation data storage, a 22 Tb high performance file system based on Lustre was used.

The list of optimization problems that we are going to solve is based on the scenario-cluster submodels obtained from the multistage stochastic mixed 0-1 problem P4 taken from the computational experience reported in [4]. The full model is a large-scale problem with  $|\Omega| = 217$  scenarios (nonsymmetric scenario tree) and  $|\mathcal{S}| = 4$  stages, where the nonanticipativity constraints has been relaxed for the first and second stage, so  $q = 44$  subproblems have been obtained. The main dimensions of the original large-scale optimization problem and the subproblems average and standard deviation dimensions are shown in Tables 2 and 3, respectively and Table 4 shows the value of the optimal objective function for the 44 mixed integer problems. The 44 MPS files can be downloaded as well as the `mainmpi.cpp` code.

Table 2: Large-scale problem dimensions

# rows	nx	ny	# non-zeros	density (%)	$q$	$ \Omega $	$\mathcal{G}$	$\mathcal{S}$
9248	2176	4792	515768	0.64	44	217	272	4

Table 3: Testbed 44 subproblem average (standard deviation) dimensions

	# rows	nx	ny	# non-zeros	density (%)
average	270	63	151	14263	7.93
(st. deviation)	(53.46)	(12.58)	(25.16)	(2981.38)	(1.57)

Tables 5 and 6 give the main execution times for COIN-OR optimizer and for CPLEX solver within COIN-OR (allowing one single thread for CPLEX), respectively. The headings are as follows: *Available threads*, the number of threads for the serial and parallel programming; *Maximum # prob/thread*, the bottleneck or maximum number of problems to be solved by thread, that is,  $\left\lceil \frac{q}{npr} \right\rceil$ ; *Min, Average, Max*, the minimum, average and maximum CPU time by thread (in seconds), respectively; *Average, St. dev.*, the average and standard deviation (in seconds) for the wall clock time after 15 executions of the same code. The CPLEX within COIN-OR solver is between 3 and 5 five times

Table 4: Testbed 44 MIP solutions

P1	-2953.46	P12	-3215.94	P23	-8945.76	P34	-2915.80
P2	-5067.97	P13	-6662.71	P24	-7271.52	P35	-4291.44
P3	-3957.70	P14	-4078.15	P25	-7136.25	P36	-9412.94
P4	-5310.86	P15	-5315.24	P26	-7046.85	P37	-7899.80
P5	-9219.36	P16	-8504.21	P27	-5561.61	P38	-4851.96
P6	-6539.62	P17	-6592.45	P28	-5327.66	P39	-7735.64
P7	-8196.80	P18	-5186.78	P29	-3991.81	P40	-7764.66
P8	-6582.53	P19	-6291.11	P30	-8051.70	P41	-9828.10
P9	-7811.68	P20	-7378.17	P31	-4887.34	P42	-2786.57
P10	-8126.43	P21	-8450.21	P32	-5434.41	P43	-4311.35
P11	-8420.17	P22	-8558.38	P33	-9432.48	P44	-8135.21

faster than the plain use of COIN-OR (without any presolving or cut generation). If we compare CPU time versus real time, we can conclude that communication time is not very relevant.

Table 5: Computing times under COIN-OR

	Available threads	Maximum # prob/thread	CPU time			Real time	
			Min.	Average	Max.	Average	St. dev.
<b>Serial</b>	<b>1</b>	44	17.651	17.651	17.651	17.780	0.013
<b>MPI</b>	<b>1</b>	44	18.013	18.013	18.013	18.239	0.093
	<b>2</b>	22	8.823	8.866	8.779	8.972	0.010
	<b>4</b>	11	5.018	5.165	5.271	5.373	0.029
	<b>8</b>	6	3.118	3.647	3.960	4.508	0.253
	<b>16</b>	3	1.735	2.077	2.272	2.632	0.270
	<b>32</b>	2	1.239	1.857	1.996	2.054	0.005
	<b>64</b>	1	0.839	1.403	1.517	1.590	0.260

Table 6: Computing times under COIN-OR&amp;CPLEX - single thread for CPLEX

	Available threads	Maximum # prob/thread	CPU time			Real time	
			Min.	Average	Max.	Average	St. dev.
<b>Serial</b>	<b>1</b>	44	6.107	6.107	6.107	6.163	0.019
<b>MPI</b>	<b>1</b>	44	6.344	6.344	6.344	6.401	0.005
	<b>2</b>	22	3.161	3.165	3.168	3.201	0.002
	<b>4</b>	11	1.647	1.667	1.685	1.706	0.005
	<b>8</b>	6	0.992	1.025	1.074	1.120	0.030
	<b>16</b>	3	0.521	0.546	0.564	0.604	0.058
	<b>32</b>	2	0.349	0.390	0.405	0.426	0.0175
	<b>64</b>	1	0.227	0.259	0.276	0.297	0.0857

Table 7 gives the main execution times for CPLEX solver within COIN-OR considering a number of available threads and several combinations between the number of threads for distributed memory (MPI threads) and shared memory (CPLEX threads). The rest of headings are as defined in the previous tables. For 8 available threads the average real time for computing the serial execution is 7.699 seconds, which corresponds to the worst strategy. The Figure 6 shows the real time for COIN-OR in horizontal red line and the real time for CPLEX in vertical blue bars, it is classified according to the available number of threads, as detailed in Tables 5 and 7. We can observe that for a fixed number of available threads the real time for solving the  $q$  problems (a big number of them with small dimensions) is decreasing with the bigger number of MPI threads. So, the fastest case corresponds to the biggest number of MPI threads and a single thread for CPLEX, which highlights the interest of MPI parallel computing.

Table 7: Computing times under COIN-OR&CPLEX - multiple threads for CPLEX

Available threads	MPI $\times$ CPLEX threads	CPU time			Real time	
		Min.	Average	Max.	Average	St. dev.
1	1 $\times$ 1	6.107	6.107	6.107	6.163	0.019
2	1 $\times$ 2	4.328	4.328	4.328	5.683	0.046
2	2 $\times$ 1	3.161	3.165	3.168	3.201	0.002
4	1 $\times$ 4	2.382	2.382	2.382	6.076	0.055
4	2 $\times$ 2	2.141	2.184	2.227	2.986	0.022
4	4 $\times$ 1	1.647	1.667	1.685	1.706	0.005
8	1 $\times$ 8	2.590	2.590	2.590	7.699	0.222
8	2 $\times$ 4	1.184	1.229	1.258	3.229	0.245
8	4 $\times$ 2	1.095	1.124	1.176	1.632	0.032
8	8 $\times$ 1	0.992	1.025	1.074	1.120	0.030

MPI= 1 corresponds to serial execution

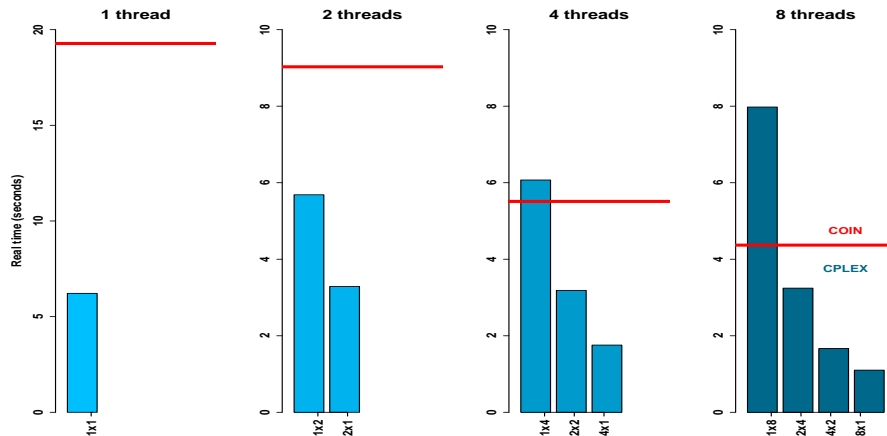


Figure 6: CPLEX within COIN-OR real time, classification for (MPI  $\times$  CPLEX) threads

Finally, the Table 8 summarizes the best real times for COIN-OR and CPLEX, with all the available threads considered. The headings are as follows: *Real Time*, the wall clock time in seconds (average time over 15 realizations); *Speedup*, serial real time over parallel real time; and *Efficiency*, parallel real time over number of threads, in percentage. The speedup and efficiency is high and similar for 2 and 4 available threads, but is better for CPLEX within COIN-OR optimization solver for 8 or more threads. These results are illustrated in Figure 7.

Table 8: Efficiency COIN-OR versus COIN-OR&CPLEX

Available threads	COIN-OR			COIN-OR&CPLEX			
	Real time	Speedup	Efficiency	Real time	Speedup	Efficiency	
Serial	1	17.780	1.000	100.0	6.163	1.000	100.0
MPI	2	8.972	1.982	99.1	3.201	1.925	96.3
	4	5.373	3.309	82.7	1.706	3.613	90.3
	8	4.508	3.944	49.3	1.120	5.503	68.8
	16	2.632	6.755	42.2	0.604	10.204	63.8
	32	2.054	8.656	27.1	0.426	14.670	45.8
	64	1.590	11.182	17.5	0.297	20.751	32.4

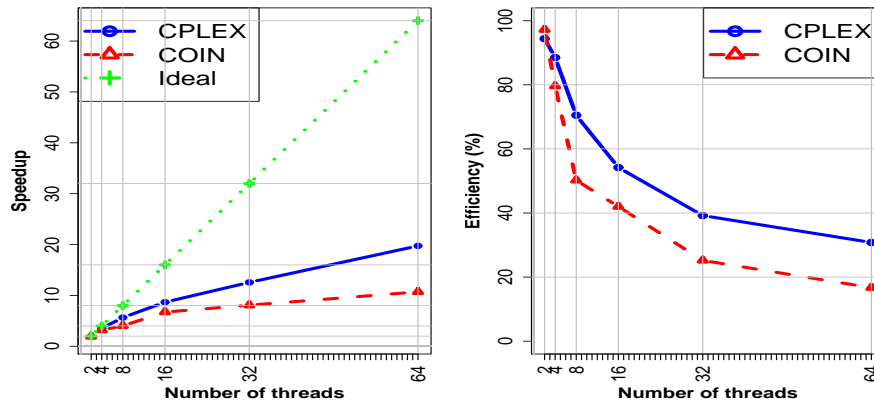


Figure 7: Speedup and efficiency vs number of threads

## Appendix 1

The main program `mainmpi.cpp` is detailed below (C++ keywords are shown in blue, comments in green and MPI elements in red).

```

1 // This is the last version of the code mainmpi.cpp. All rights reserved.
2 // Copyright (C) 2012 by U. Aldasoro, M.A. Garin, M. Merino and G. Perez,
3 // from UPV/EHU.
// No part of this code may be reproduced, modified or transmitted, in any

```

```

5 // form or by any means without the prior written permission of the authors.
7 // Cluster models are in mps format, as external files
9 #include "itoa.h"
  #include <mpi.h>
11 #include <string.h>
  #include <stdlib.h>
13 #include <math.h>
  #include <assert.h>
15 #include <stdio.h>
  #include <iostream>
17 #include <fstream>
  using namespace std;
19 #include "ClpSimplex.hpp"
  #include "CoinHelperFunctions.hpp"
21 #include "CoinBuild.hpp"
  #include "CoinModel.hpp"
23 #include <iomanip>
  #include <cassert>
25 #include "OsiClpSolverInterface.hpp"
  #include "CbcModel.hpp"
27
  #define rmaxmin (1) // (1): Maximum; (-1): Minimum
29 #define nmodel 44 // Number of cluster models
  #define useCplex 1 // Select solver. 0: COIN-OR; 1: Cplex
31 #define threadsCplex 2 // Number of parallel threads used by each MPI thread
33 #ifdef useCplex
  #include "OsiCpxSolverInterface.hpp"
35 #endif
37 int main(int argc, char **argv)
  {
39 // STEP 0 – DECLARING OPTIMIZATION AND MPI VARIABLES
  int imod,i,j,loc,pid,npr,lp,original_rank,original_size,error,ncols;
41 int assignment[nmodel],inicial[nmodel],nonzero[nmodel],numres[nmodel],
  numvar[nmodel],numvarint[nmodel],assignmentX0[nmodel],iniX0[nmodel],
43 istruef969[nmodel];
  double zq[nmodel];
45
  MPI_Group orig_group,new_group; //MPI groups
47 MPI_Comm new_comm; //MPI communicator
49
  // STEP 1 – DEFINITION OF THE GLOBAL ENVIRONMENT
51 MPI_Init(&argc,&argv); // Beginning of the MPI environment
  MPI_Comm_size(MPI_COMM_WORLD, &original_size); // Total number of threads
53 MPI_Comm_rank(MPI_COMM_WORLD, &original_rank); // Who am I?
55
  // 1.1 – Creating a new communicator considering active threads
57
  // 1.1.1– Basic variables
  int ranks1[nmodel];

```

```

59  int *ranks2; ranks2=new int[original_size];
60  for(i=0;i<nmodel;i++) {ranks1[i]=i;}
61  for(i=0;i<original_size;i++) {ranks2[i]=i;}
62
63  // 1.1.2- Extract handle of global group from MPI_COMM_WORLD using
        MPI_Comm_group
64  MPI_Comm_group(MPI_COMM_WORLD, &orig_group);
65
66  // 1.1.3- Form new group as a subset of global group using MPI_Group_incl
67  if (nmodel < original_size){MPI_Group_incl(orig_group, nmodel, ranks1, &
        new_group);
        } else {MPI_Group_incl(orig_group, original_size, ranks2, &new_group);}
68
69  // 1.1.4- Create new communicator for new group using MPI_Comm_create
70  MPI_Comm_create(MPI_COMM_WORLD, new_group, &new_comm);
71
72  // 1.1.5- Determine new rank in new communicator using MPI_Comm_rank
73  MPI_Group_rank (new_group, &pid);
74  MPI_Group_size (new_group, &npr);
75
76  // 1.1.6 - Creating the output file
77  ofstream solution("CPLEX_COIN_MPI_solution.dat", ios::out); //Output file
78
79  // 1.2 - Only active threads follow the solving process
80
81  if (original_rank < npr) {
82
83      // 1.2.1 - Assigning the work load
84      for(i=0;i<npr;i++) assignment[i]=int(nmodel/npr);
85      for(i=0;i<(nmodel-npr*int(nmodel/npr));i++) assignment[i]=assignment[i]+1;
86
87      inicial[0]=0;
88      for(i=1;i<npr;i++) inicial[i]=inicial[i-1]+assignment[i-1];
89
90      // 1.2.2 - Declaring local variables
91      int ncols_max_loc=0;
92      double *zq_loc; zq_loc=new double[assignment[pid]];
93      int *nonzero_loc; nonzero_loc=new int[assignment[pid]];
94      int *numres_loc; numres_loc=new int[assignment[pid]];
95      int *numvar_loc; numvar_loc=new int[assignment[pid]];
96      int *numvarint_loc; numvarint_loc=new int[assignment[pid]];
97      int *istruef969_loc; istruef969_loc=new int[assignment[pid]];
98
99      // 1.2.3 - Creating the output file
100     if (pid == 0) {
101         if (useCplex == 0) {solution<<"*** SOLVER: COIN-OR ***"<<"\n\n";}
102         else {solution<<"*** SOLVER: CPLEX *** \n\n +++ NUMBER OF " <<
            "PARALLEL THREADS ON CPLEX = "<<threadsCplex<<" +++ \n\n";}
103
104         solution<<"Beginning of CPLEX_COIN_MPI_solution.dat (output of mainmpi.
            cpp) \n Number of threads = "<<original_size<<"\n Number of
            submodels = "<<nmodel<<"\n Number of active threads = "<<npr<<"\n";
105         for(i=0;i<npr;i++) {
106             solution<<" Number of submodels solved by thread "<<i<<" = "<<

```

```

109         assignment[i]<<" , starting by submodel = "<< inicial[i]+1<<"\n";}
110     }
111
112     // 1.2.4 – Creating the COIN-OR or CPLEX/COIN-OR models
113     OsiClpSolverInterface *sol0;
114     sol0=new OsiClpSolverInterface [ assignment [ pid ] ];
115
116     OsiCpxSolverInterface *sol1;
117     sol1=new OsiCpxSolverInterface [ assignment [ pid ] ];
118     for ( loc=0;loc<assignment [ pid ];loc++){
119         CPXENVptr env = sol1 [ loc ].getEnvironmentPtr();
120         CPXsetintparam(env, CPX_PARAM_THREADS, threadsCplex);
121     }
122
123     for ( loc=0;loc<assignment [ pid ];loc++) {
124
125         imod=inicial [ pid ]+loc;         const char* final;
126         final="";                         char buffer[33];
127         final=itoa(imod+1, buffer,10); //convert in char
128         char model[80];                     strcpy (model,"Cluster");
129         strcat (model,final);               puts (model);
130
131         if (useCplex == 0) {
132             sol0 [ loc ].setObjSense(rmaxmin);
133             sol0 [ loc ].readMps(model); //Read cluster submodel
134             nonzero_loc[ loc ] = sol0 [ loc ].getNumElements();
135             numres_loc[ loc ] = sol0 [ loc ].getNumRows();
136             numvar_loc[ loc ] = sol0 [ loc ].getNumCols();
137             for (j=0;j<sol0 [ loc ].getNumCols();j++)
138                 if(sol0 [ loc ].isInteger(j)) numvarint_loc[ loc ]++;
139             if(sol0 [ loc ].getNumCols() > ncols_max_loc)
140                 {ncols_max_loc = sol0 [ loc ].getNumCols();}
141         } else {
142             sol1 [ loc ].setObjSense(rmaxmin);
143             sol1 [ loc ].readMps(model); //Read cluster submodel
144             nonzero_loc[ loc ] = sol1 [ loc ].getNumElements();
145             numres_loc[ loc ] = sol1 [ loc ].getNumRows();
146             numvar_loc[ loc ] = sol1 [ loc ].getNumCols();
147             for (j=0;j<sol1 [ loc ].getNumCols();j++)
148                 if(sol1 [ loc ].isInteger(j)) numvarint_loc[ loc ]++;
149             if(sol1 [ loc ].getNumCols() > ncols_max_loc)
150                 {ncols_max_loc = sol1 [ loc ].getNumCols();}
151         }
152     }
153
154     if (npr > 1) {
155         MPI_Allreduce(&ncols_max_loc,&ncols , 1, MPI_INT, MPI_MAX,new_comm);
156     } else {
157         ncols=ncols_max_loc;}
158
159     double **x0;x0=new double *[nmodel];
160         for(i=0;i<nmodel;i++) x0[i]=new double [ncols];
161
162     double *x0_loc; x0_loc=new double [assignment [pid]*ncols];

```

```

163     double *x0vector;  x0vector=new double [nmodel*ncols];
165     for (imod=0;imod<assignment [ pid ];imod++) {
167         for (j=0;j<nmodel;j++) {
169             x0_loc[imod*ncols+j]=0;
171         }
172     }
173     for (i=0;i<npr;i++) {
174         iniX0[i]=inicial [ i ]*ncols;  assignmentX0 [ i ]=assignment [ i ]*ncols;}
175
176     // STEP 2 – PRESOLVE ASSIGNED MODELS
177     error=0;
178
179     for (loc=0;loc<assignment [ pid ];loc++)
180     {
181         if (useCplex == 0) {
182             imod=inicial [ pid ]+loc;
183             CbcModel pm0(sol0 [ loc ]);
184             istruef969 _loc[ loc ]=0;          sol0 [ loc ].initialSolve ();
185
186             if (sol0 [ loc ].isProvenPrimalInfeasible ()) {zq_loc [ loc ]=rmaxmin*1.e25;
187                 istruef969 _loc[ loc ]=1;}
188
189             if (istruef969 _loc[ loc ]==0){
190                 if (sol0 [ loc ].isProvenDualInfeasible () ){
191                     zq_loc [ loc ]=-1.0*rmaxmin*1.e25; istruef969 _loc[ loc ]=2;}
192
193                 if (!sol0 [ loc ].isProvenOptimal () ) {
194                     zq_loc [ loc ]=-1.0*rmaxmin*1.e25; istruef969 _loc[ loc ]=2;}
195             }
196         } else {
197             imod=inicial [ pid ]+loc;
198             istruef969 _loc[ loc ]=0;          sol1 [ loc ].initialSolve ();
199
200             if (sol1 [ loc ].isProvenPrimalInfeasible ()) {zq_loc [ loc ]=rmaxmin*1.e25;
201                 istruef969 _loc[ loc ]=1;}
202
203             if (istruef969 _loc[ loc ]==0){
204                 if (sol1 [ loc ].isProvenDualInfeasible () ){
205                     zq_loc [ loc ]=-1.0*rmaxmin*1.e25; istruef969 _loc[ loc ]=2;}
206
207                 if (!sol1 [ loc ].isProvenOptimal () ) {
208                     zq_loc [ loc ]=-1.0*rmaxmin*1.e25; istruef969 _loc[ loc ]=2;}
209             }
210         }
211     }
212
213     // STEP 3 – GLOBAL MPI COMMUNICATION (GATHERING PRESOLVE INFORMATION)
214     if (npr > 1) {
215         MPI_Allgatherv(&istruef969_loc [0], assignment [ pid ], MPI_INT,
216             &istruef969 [0],assignment , inicial , MPI_INT, new_comm);
217         MPI_Gatherv(&zq_loc [0], assignment [ pid ], MPI_INT, &zq [0],
218             assignment , inicial , MPI_INT, 0, new_comm);

```



```

217     } else {
218         for(imod=0;imod<nmodel;imod++) {
219             istruef969 [imod]=istruef969 _loc[imod];    zq [imod]=zq _loc[imod];
220         }
221     }
222
223     for(imod=0;imod<nmodel;imod++) {
224         if(istruef969 [imod] == 1){
225             solution<<"\n Cluster submodel "<<imod+1
226                 <<" is primal infeasible. ObjValue = "<<zq[imod]; error=1;}
227         if(istruef969 [imod] == 2){
228             solution<<"\n Cluster submodel "<<imod+1
229                 <<" is unbounded. ObjValue = "<<zq[imod]; error=1;}
230     }
231
232     // STEP 4 – CHECK PRESOLVE STATUS
233     if (error == 0){
234
235         // STEP 5– SOLVE ASSIGNED MODELS
236         if (useCplex == 0) {
237             for (loc=0;loc<assignment [pid];loc++)
238             {
239                 imod=inicial [pid]+loc;    CbcModel pm0(sol1 [loc]);
240                 pm0.branchAndBound ();
241
242                 if (!pm0.isProvenOptimal ()) {
243                     zq _loc[loc]=-1.0*rmaxmin*1.e25;    istruef969 _loc[loc]=3;}
244
245                 if (fabs(pm0.getBestPossibleObjValue ())>1.e24){
246                     zq _loc[loc]=-1.0*rmaxmin*1.e25;    istruef969 _loc[loc]=4;}
247
248                 if(istruef969 _loc[loc] == 0) {
249                     zq _loc[loc]=pm0.getObjValue ();
250                     for (j=0;j<pm0.getNumCols ();j++) {
251                         x0 _loc[loc*ncols+j]=pm0.getColSolution () [j];}
252                 }
253             } else {
254                 for (loc=0;loc<assignment [pid];loc++)
255                 {
256                     imod=inicial [pid]+loc;    sol1 [loc].branchAndBound ();
257
258                     if (!sol1 [loc].isProvenOptimal ()) {
259                         zq _loc[loc]=-1.0*rmaxmin*1.e25;    istruef969 _loc[loc]=3;}
260
261                     if (fabs(sol1 [loc].getObjValue ())>1.e24){
262                         zq _loc[loc]=-1.0*rmaxmin*1.e25;    istruef969 _loc[loc]=4;}
263
264                     if(istruef969 _loc[loc] == 0) {
265                         zq _loc[loc]=sol1 [loc].getObjValue ();
266                         for (j=0;j<sol1 [loc].getNumCols ();j++)
267                             x0 _loc[loc*ncols+j]=sol1 [loc].getColSolution () [j];}
268                 }
269             }

```

```

271 // STEP 6– GLOBAL MPI COMMUNICATION (GATHERING SOLVE INFORMATION)
272 if (npr > 1) {
273     MPI_Gatherv(&zq_loc [0], assignment [ pid ],MPI_DOUBLE,&zq [0] ,
274               assignment , inicial ,MPI_DOUBLE,0 ,new_comm);
275     MPI_Gatherv(&nonzero_loc [0], assignment [ pid ], MPI_INT,&nonzero [0] ,
276               assignment , inicial ,MPI_INT,0 ,new_comm);
277     MPI_Gatherv(&numvar_loc [0], assignment [ pid ],MPI_INT,&numvar [0] ,
278               assignment , inicial ,MPI_INT,0 ,new_comm);
279     MPI_Gatherv(&numvarint_loc [0], assignment [ pid ],MPI_INT,&numvarint [0] ,
280               assignment , inicial ,MPI_INT,0 ,new_comm);
281     MPI_Gatherv(&numres_loc [0], assignment [ pid ],MPI_INT,&numres [0] ,
282               assignment , inicial ,MPI_INT,0 ,new_comm);
283     MPI_Gatherv(x0_loc , assignmentX0 [ pid ],MPI_DOUBLE,x0vector ,
284               assignmentX0 ,iniX0 ,MPI_DOUBLE,0 ,new_comm);
285     MPI_Gatherv(&istruef969_loc [0], assignment [ pid ],MPI_INT,&istruef969 [0] ,
286               assignment , inicial ,MPI_INT,0 ,new_comm);
287 } else {
288     for (imod=0;imod<nmodel;imod++) {
289         zq [imod]=zq_loc [imod];           nonzero [imod]=nonzero_loc [imod];
290         numvar [imod]=numvar_loc [imod];   numres [imod]=numres_loc [imod];
291         istruef969 [imod]=istruef969_loc [imod];
292         for (j=0;j<numvar [imod];j++){
293             x0vector [imod*ncols+j]=x0_loc [imod*ncols+j];
294         }
295     }
296 }
297
298 for (imod=0;imod<nmodel;imod++) {
299     if ((istruef969 [imod] == 3) | (istruef969 [imod] == 4)){
300         solution << "\n Cluster submodel "<<imod+1 << " is unbounded ";}
301 }
302
303 // STEP 7 – PRINT RESULTS (if pid == 0)
304 if (pid == 0) {
305
306     solution << "\n\n [RESULTS 1:] PROPERTIES OF SOLVED MODELS ";
307     for (imod=0;imod<nmodel;imod++){
308         solution << "\n \n Cluster submodel "<<imod+1<< " of "<<nmodel;
309         solution << "\n Number of variables: "; solution << numvar [imod];
310         solution << "\n Number of constraints: "; solution << numres [imod];
311         solution << "\n Number of nonzero elements: "<<nonzero [imod];
312         for (j=0;j<numvar [imod];j++){
313             x0 [imod][j]=x0vector [imod*ncols+j];
314             // solution << "\n x[ "<<imod<<","<<j<<"]"<<x0 [imod][j]<< " \n"; // To
315             print x0
316         }
317     }
318
319     solution << "\n\n\n [RESULTS 2:] OBJECTIVE FUNCTION VALUES \n";
320
321     for (imod=0;imod<nmodel;imod++) {
322         solution << "\n Optimal objective value Cluster submodel "<<imod+1
323         << " of "<<nmodel << " is "<<zq [imod];}

```

```

    }
325 }
327 // STEP 8 – EXECUTION ENDS IN ALL PROCESSORS
329 solution.close();
331 MPI_Finalize(); //End of the MPI environment
    return 0;
}

```

mainmpi.cpp

## Appendix 2

To compile and link the code, you can also use a Makefile similar to the described below in this Appendix, where are used additionally the libraries `-lpthread`. As summary to generate and run the executable you must to install the packages,

lam-runtime	libgomp1
liblam4	libmpich1.0gf
libopenmpi-dev	libopenmpi1.3
mpi-default-bin	mpi-default-dev
mpich-bin	openmpi-bin
openmpi-checkpoint	openmpi-common
openmpi-doc	

and their dependences. All of them have been installed in the machine `b012526.bs.ehu.es` of the Laboratory of Quantitative Economics from the University of the Basque Country (UPV/EHU, Bilbao, Spain). To run the executable of name `mainmpi` in this machine with eight cores, you must type from the prompt,

```
$ mpirun -v -np 8 -host b012526 mainmpi
```

```

2 # Copyright (C) 2006 International Business Machines and others.
# All Rights Reserved.
# This file is distributed under the Common Public License.
4 # $Id: Makefile.in 726 2006-04-17 04:16:00Z andreasw $
#####
6 # You can modify this example makefile to fit for your own program. #
# Usually, you only need to change the five CHANGEME entries below. #
8 #####
# To compile other examples, either changed the following line, or
10 # add the argument DRIVER=problem_name to make
DRIVER = mainmpi
12
# CHANGEME: This should be the name of your executable
14 EXE = $(DRIVER)

16 # CHANGEME: Here is the name of all object files corresponding to the source
# code that you wrote in order to define the problem statement
18 OBJS= mainmpi.o

```

```

SYSTEM = x86-64_sles10_4.1
20 # Directory with COIN header files
COININCDIR = /opt/coin-1.3.1/include/coin
22 # Directory with COIN libraries
COINLIBDIR = /opt/coin-1.3.1/lib
24 COINLIBDIROSI = /opt/coin-1.3.1/Osi/lib
# Directory with CPLEX
26 CPLEXDIR= /opt/ILOG/CPLEX_Studio_AcademicResearch122/cplex
# Directory with CPLEX header files
28 CPLEXINCDIR = $(CPLEXDIR)/include/ilcplex
# Directory with CPLEX libraries
30 CPLEXLIBDIR=$(CPLEXDIR)/lib/$(SYSTEM)/static_pic/libcplex.a

32 CPLEXBINDIR=$(CPLEXDIR)/bin/$(SYSTEM)

34 # CHANGEME: Additional libraries
ADDLIBS =
36
# CHANGEME: Additional flags for compilation (e.g., include flags)
38 ADDINCFLAGS =
#####
40 # Usually, you don't have to change anything below. Note that if you #
# change certain compiler options, you might have to recompile the #
42 # package. #
#####
44 # C++ Compiler command
#CXX = g++
46 CXX=mpic++

48 # C++ Compiler options for g++
#CXXFLAGS = -O3 -fomit-frame-pointer -pipe -DNDEBUG -pedantic-errors -Wimplicit
#-Wparentheses -Wreturn-type -Wcast-qual -Wall -Wpointer-arith -Wwrite-
strings
50 #-Wconversion -Wno-unknown-pragmas -O -fPIC -fexceptions -DIL_STD

52 # C++ Compiler options for mpic++
CXXFLAGS = -O3 -fomit-frame-pointer -pipe -DNDEBUG -Wimplicit -Wparentheses -
Wreturn-type -Wcast-qual -Wall -Wpointer-arith -Wwrite-strings -Wconversion
-Wno-unknown-pragmas -O -fPIC -fexceptions -DIL_STD
54
# additional C++ Compiler options for linking
56 CXXLINKFLAGS = -Wl,--rpath -Wl,$(COINLIBDIR) -Wl,$(CPLEXLIBDIR)

58 # Libraries necessary to link with Clp
LIBS = -L$(COINLIBDIR) -lCbcSolver -lCbc -lCgl -lOsiClp -lOsiCbc -lOsi -lClp -
lCoinUtils -L$(COINLIBDIROSI) -lOsiCpx -L$(CPLEXBINDIR) -lcplex122 -lpthread
\
60 -lm

62 # Necessary Include dirs (we use the CYGPATH_W variables to allow
# compilation with Windows compilers)
64 INCL = -I$(CYGPATH_W) $(COININCDIR) $(CPLEXINCDIR) $(ADDINCFLAGS)

66 # The following is necessary under cygwin, if native compilers are used

```

```

68 CYGPATH_W = echo
# Here we list all possible generated objects or executables to delete them
70 CLEANFILES =
72 all: $(EXE)
74 .SUFFIXES: .cpp .c .o .obj
76 $(EXE): $(OBJS)
    bla =;\
78   for file in $(OBJS); do bla="$$bla '$(CYGPATH_W) $$file '"; done; \
    $(CXX) $(CXXLINKFLAGS) $(CXXFLAGS) -o $@ $$bla $(ADDLIBS) $(LIBS)
80 clean:
    rm -rf $(CLEANFILES) $(OBJS)
82 .cpp.o:
    $(CXX) $(CXXFLAGS) $(INCL) -c -o $@ 'test -f '$<' || echo '$(SRCDIR)/'$<
84 .cpp.obj:
    $(CXX) $(CXXFLAGS) $(INCL) -c -o $@ 'if test -f '$<; then $(CYGPATH_W) '$<;
    else $(CYGPATH_W) '$(SRCDIR)/$<; fi '
86 .c.o:
    $(CC) $(CFLAGS) $(INCL) -c -o $@ 'test -f '$<' || echo '$(SRCDIR)/'$<
88 .c.obj:
    $(CC) $(CFLAGS) $(INCL) -c -o $@ 'if test -f '$<; then $(CYGPATH_W) '$<;
    else $(CYGPATH_W) '$(SRCDIR)/$<; fi '

```

Makefile

## 7 Acknowledgements

This research has been partially supported by the projects ECO2008-00777 ECON from the Ministry of Education and Science, Grupo de Investigación IT-347-10 from the Basque Government, Spain. Technical and human support provided by Dr. E. Ogando and Dr. T. Mercero from IZO-SGI, SGIker (UPV/EHU, MICINN, GV/EJ, ERDF and ESF) is gratefully acknowledged. We would like to express our gratefulness to Prof. Dr. F. Tusell for making easier the access to the Laboratory of Quantitative Economics from the University of the Basque Country (UPV/EHU, Bilbao, Spain) to perform and check some of our computational experiments.

## References

- [1] ARINA cluster. IZO-SGI, SGIker (UPV/EHU). Website.  
<http://www.ehu.es/sgi/recursos/cluster-arina-2>.
- [2] COIN-OR. COmputational INfrastructure for Operations Research. Website.  
<http://www.coin-or.org/>.

- [3] David E. Culler, Anoop Gupta, and Jaswinder Pal Singh. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 1997.
- [4] Laureano F. Escudero, María Araceli Garín, María Merino, and Gloria Pérez. An algorithmic framework for solving large-scale multistage stochastic mixed 0-1 problems with nonsymmetric scenario trees. *Comput. Oper. Res.*, 39:1133–1144, May 2012.
- [5] GCC. GNU compiler collection. Website. <http://gcc.gnu.org/>.
- [6] John L. Hennessy and David A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3 edition, 2003.
- [7] IBM. ILOG CPLEX optimizer. Website  
. <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [8] J T Linderoth. *Topics in Parallel Integer Optimization*. PhD thesis, Georgia Institute of Technology, 1998.
- [9] R Lougee-Heimer. The common optimization Interface for operations research: Promoting open-source software in the operations research community. *IBM JOURNAL OF RESEARCH AND DEVELOPMENT*, 47(1):57–66, JAN 2003.
- [10] Peter S. Pacheco. *Parallel programming with MPI*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.
- [11] Gloria Pérez Sainz de Rozas and María Araceli Garín Martín. On Downloading and Using CPLEX within COIN-OR for Solving Linear/Integer Optimization Problems. Biltoki, Universidad del País Vasco - Departamento de Economía Aplicada III (Econometría y Estadística), 2011.
- [12] Marc Snir, Steve W. Otto, David W. Walker, Jack Dongarra, and Steven Huss-Lederman. *MPI: The Complete Reference*. MIT Press, Cambridge, MA, USA, 1995.