

---

# Particle Swarm Optimization: a parallelized approach

Bortolotti Samuele, Izzo Federico



2022-11-22

## Contents


|   |           |
|---|-----------|
| <b>1 Abstract</b>                               | <b>4</b>  |
| <b>2 Introduction</b>                           | <b>5</b>  |
| 2.1 Particle Swarm Optimization . . . . .       | 5         |
| 2.1.1 Generalities . . . . .                    | 5         |
| 2.1.2 Parametrization . . . . .                 | 6         |
| 2.1.3 Continuous Optimization . . . . .         | 6         |
| 2.2 MPI . . . . .                               | 8         |
| 2.3 OpenMP . . . . .                            | 8         |
| 2.4 State-of-the art analysis . . . . .         | 8         |
| 2.5 Project generalities . . . . .              | 9         |
| 2.5.1 Libraries . . . . .                       | 9         |
| 2.5.2 Build . . . . .                           | 10        |
| 2.5.3 Execute . . . . .                         | 11        |
| <b>3 Problem Analysis</b>                       | <b>12</b> |
| 3.1 Serial version of the algorithm . . . . .   | 17        |
| 3.1.1 Serial algorithm optimization . . . . .   | 18        |
| 3.2 Parallel version of the algorithm . . . . . | 19        |
| 3.2.1 Architecture . . . . .                    | 20        |
| 3.2.2 Message . . . . .                         | 20        |
| 3.2.3 Communication pattern . . . . .           | 21        |
| 3.2.4 Logs . . . . .                            | 23        |
| 3.2.5 Output and SQLite . . . . .               | 24        |
| <b>4 Benchmarking</b>                           | <b>25</b> |
| 4.1 Problem configuration . . . . .             | 25        |
| 4.2 Cluster jobs . . . . .                      | 25        |
| 4.3 Results . . . . .                           | 26        |
| <b>5 DevOps</b>                                 | <b>32</b> |
| 5.1 Nix . . . . .                               | 32        |
| 5.2 Docker . . . . .                            | 33        |
| 5.3 GitHub actions . . . . .                    | 33        |
| 5.3.1 Container creation . . . . .              | 33        |
| 5.3.2 Documentation compilation . . . . .       | 33        |
| 5.4 Udocker . . . . .                           | 34        |
| 5.4.1 Build phase . . . . .                     | 34        |
| 5.4.2 OpenMPI communication . . . . .           | 34        |
| <b>6 Conclusion</b>                             | <b>36</b> |
| 6.1 Future Work . . . . .                       | 36        |

**References**

## 1 Abstract

*Particle Swarm Optimization (PSO)* is an optimization algorithm for nonlinear function based on birds swarm. It falls back into the sub-field of *Bio-Inspired Artificial Intelligence* and it was designed from a simplified social model inspired by the nature.

A key concept associated with PSO is the role of genetic algorithms and evolution, the functioning is based on several iterations that aim to identify the best possible position represented as a point in a landscape (figure 1).



**Figure 1:** Solution landscape with best possible solution represented in green

PSO is originally attributed to James Kennedy and Russell Eberhart and was first intended for simulating social behavior in 1995 [1].

The goal of this project is to design a parallelized implementation capable of exploring the solution space in a faster way. This is done through the usage of two main libraries for *High Performance Computing (HPC)*: *OpenMPI* and *OpenMP*.

The effectiveness of the proposed solution is tested using the HPC cluster of the University of Trento among other implementations found online.

## 2 Introduction

### 2.1 Particle Swarm Optimization

In order to deeply understand the reasons behind the report design choices, it is fundamental to understand comprehensively *Particle Swarm Optimization*.

#### 2.1.1 Generalities

PSO focuses on two main definitions: the notion of *particle* and the one of *particle perception*.

A particle can be seen as an entity which is characterized by:

- a position  $x$  depicting the *candidate solution* for our optimization problem;
- a velocity component  $v$ , which is used in order to *perturb* the particle;
- a performance measure  $f(x)$ , also called *fitness value*, which quantify the quality of the candidate solution.


The entire set of particles is referred as *swarm*.

Under the expression *particle perception*, we define how each particle communicate with each other. In practice, a particle needs to perceive the positions along with the associated performance measures of the *neighboring particles*. Thanks to this communication pattern, each particle remembers the position  $z$  associated to the best performance of all the particles within the neighborhood, as well as its own best performance so far  $y$ .

There are different structures of neighborhood which can be considered (figure 2), and they usually depend on the type of optimization problem one has to face.

The most relevant types of neighborhood are:

- *Global*: the best individual in the neighborhood is also the *global best* in the entire swarm;
- *Distance-based*: based on a proximity metric (e.g. euclidean distance);
- *List-based*: based on a predetermined topology arranging the solution indexes according to some order or structure, and a given neighborhood size.



**Figure 2:** Different neighborhood structures in PSO

This project implements a version of PSO considering *distance-based* neighborhood in a nearest neighbor fashion. In details, each particle has a fixed number of neighbors, which depend dynamically on the particle position on

the landscape. The program offers the user the possibility to modify the number of particles to consider within a particle neighborhood.

### 2.1.2 Parametrization

In order to assess a solution for an optimization problem, PSO requires the following parameters to be set:

- *Swarm size*: typically 20 particles for problems with dimensionality 2-200;
- *Neighborhood size*: typically 3 to 5, otherwise global neighborhood;
- *Velocity update factors*.

### 2.1.3 Continuous Optimization

Once the algorithm has been parametrized, a swarm of particles is initialized with random positions and velocity.

At each step, each particle updates firstly its velocity (equation 1):

$$v' = w \cdot v + \phi_1 U_1 \cdot (y - x) + \phi_2 U_2 \cdot (z - x) \quad (1)$$

where:

- $x$  and  $v$  are the particle current position and velocity, respectively;
- $y$  and  $z$  are the personal and social/global best position, respectively;
- $w$  is the inertia (weighs the current velocity);
- $\phi_1, \phi_2$  are acceleration coefficients/learning rates (cognitive and social, respectively);
- $U_1$  and  $U_2$  are uniform random numbers in  $[0, 1]$ .

Finally, each particle updates its position (equation 2):

$$x' = x + v' \quad (2)$$

and in case of improvement, update  $y$  (and eventually  $z$ ).

The loop is iterated until a given stop condition is met.

The pseudocode of the algorithm is shown below:

**Algorithm 1** Initialize

---

```

1: procedure INITIALIZE( $S, \mathcal{D}, f, v, x, x_{min}, x_{max}, v_{max}$ )
2:   for each particle  $i \in S$  do
3:     for each dimension  $d \in \mathcal{D}$  do
4:        $x_{i,d} \leftarrow Rnd(x_{min}, x_{max})$  ▷ Initialize the particles' positions
5:        $v_{i,d} \leftarrow Rnd(-v_{max}/3, v_{max}/3)$  ▷ Initialize the particles' velocity
6:     end for
7:   end for
8:    $pb_i \leftarrow x_i$  ▷ Initialize the particle best position
9:    $gb_i \leftarrow x_i$  ▷ Update the particle's best position
10: end procedure

```

---

**Algorithm 2** Particle Swarm Optimization (Nearest Neighbors)

---

```

1: function PSO( $S, \mathcal{D}, MAX\_IT, n, f, v, x, x_{min}, x_{max}, v_{max}$ )
2:   INITIALIZE( $S, \mathcal{D}, f, v, x, x_{min}, x_{max}, v_{max}$ ) ▷ Initialize all the particles
3:    $it = 0$ 
4:   repeat
5:     for each particle  $i \in S$  do
6:       if  $f(x_i) < f(pb_i)$  then
7:          $pb_i \leftarrow x_i$  ▷ Update the particles' best position
8:       end if
9:     end for
10:     $S' = \text{COPY}(S)$  ▷ Copy the particle's vector
11:    for each particle  $i \in S$  do
12:       $S' = \text{SORT}(S', i)$  ▷ Sort the particles w.r.t.  $i$ th particle
13:      for each particle  $j \in S'$  do
14:        if  $f(x_j) < f(gb_i)$  then
15:           $gb_i \leftarrow x_j$ 
16:        end if
17:      end for
18:    end for
19:    for each particle  $i \in S$  do
20:      for each dimension  $d \in \mathcal{D}$  do
21:         $v_{i,d} = v_{i,d} + C_1 \cdot Rnd(0, 1) \cdot [pb_{i,d} - x_{i,d}] + C_2 \cdot Rnd(0, 1) \cdot [gb_d - x_{i,d}]$ 
22:         $x_{i,d} = x_{i,d} + v_{i,d}$  ▷ Update the velocity and positions
23:      end for
24:    end for
25:     $it \leftarrow it + 1$  ▷ Advance iteration
26:  until  $it < MAX\_ITERATIONS$ 
27:  return  $x$ 
28: end function

```

---

## 2.2 MPI

The MPI (Message Passing Interface) library is used to convey information across processes running on different nodes of a cluster.

In the scenario described by the application, the basic information unit is composed as a broadcast message shared over the whole network, in this way all particles of PSO are able to know all information associated to other members of the swarm.

## 2.3 OpenMP

OpenMP is an API which supports multi-platform shared memory programming.

In the program scenario, a process is delegated to handle the computing regarding one or more particles. The process job is divided in several threads which optimize the execution time of the process.

## 2.4 State-of-the art analysis

The first serial implementation of PSO was published in 1995 by Kennedy and Eberhat [1]. After that, many variations of PSO were proposed in the literature.

Based on what we have found, the approaches can be divided into three main categories:

1. those ones which aim to change the behavior of the algorithm introducing new features;
2. those ones which aim to solve a real world problem using PSO as main algorithm;
3. those ones which aim to optimize the runtime execution speed.

In our study we have decided to exclude the second category of PSO algorithms since these solutions are strictly problem dependent. Thus, a comparison would produce meaningless results.

On the other hand, all those approaches which belong to first category of problems can be employed as case studies for our benchmarking analysis. However, it is strictly required to change some implementation aspects by modifying directly the code. In some cases, this requires a deep understanding of others' code, most of the time a tough job due to the absence of documentation.

For instance, the topology used within the neighborhood selection highly influences the final performance of the algorithm. As an effective representation, we can think about the global neighborhood in comparison with the distance based one. In the first scenario, the global best solution is shared among each swarm unit, and it is given by the best value reached by each particle. Hence, a simple reduce without further computation will suffice. Instead, a distance based neighborhood requires to consider the best solutions within a certain range of particles, which implies that all the swarm units need to know all the others' positions. This inevitably imply a considerable increase in time complexity.

The third category is our perfect competitor, since they share our same objective. However, there are several cases in which different PSO version have been implemented. Hence, some hands on is still required.

In the following table we list some of the implementations we have decided to consider during the benchmarking phase.



**Table 1:** SOTA works.

| Authors   | Year | Type                            | Source code | Notes |
|---|------|---------------------------------|-------------|-------|
| Kennedy, Eberhart [1]                           | 1995 | Serial                          | No          |       |
| toddguant [2]                                   | 2019 | Serial                          | Yes         | 1     |
| sousouho [3]                                    | 2019 | Serial                          | Yes         | 1     |
| kkentzo [4]                                     | 2020 | Serial                          | Yes         | 1     |
| fisherling [5]                                  | 2020 | Serial                          | Yes         | 1     |
| Moraes, Mitre [6]                               | 2014 | Parallel<br>OpenMPI             | No          |       |
| Nedja, Moraes, Rogerio, Marcedo<br>Mourelle [7] | 2017 | Parallel<br>OpenMPI/MP          | No          |       |
| abhi4578 [8]                                    | 2019 | Parallel<br>OpenMPI/MP,<br>CUDA | Yes         | 1     |
| LaSEEB [9]                                      | 2020 | Parallel OpenMP                 | Yes         | 2     |
| pg443 [10]                                      | 2021 | Serial, Parallel<br>OpenMP      | Yes         | 1     |

The indexes in the notes refer to:

1. provides only global neighborhood implementation. Thus, the comparison would be untruthful as those implementations have a clear advantage in the execution time due to a favorable topology;
2. provides PSO with different neighborhood versions but without a distance based approach. Hence, the implication are the same as for the point 1.

According to the previous statements, we claim that we have implemented a PSO version which differ from the ones we have decided to consider since it has a different notion of neighborhood which makes it harder to parallelize.

## 2.5 Project generalities

In the following sections, the report addresses how to setup and run the program.

### 2.5.1 Libraries

The project requires few libraries in order to work properly. As it is mandatory for the course, OpenMP and MPI were employed. Along with the compulsory libraries, the following ones were exploited:

- `sqlite`: SQLite is a C-language library that provides a SQL database engine that is tiny, quick, self-contained, high-reliability, and full-featured. The choice of `sqlite` was made in order to save particles' information at each iteration in a simple and fast way, avoiding dealing with race conditions.
- `argp`: `argp` is a parsing interface for unix-style argument vectors. The `argp` features include, as defined in the GNU coding standards, automatically producing output in response to the `--help` and `--version` options and the possibility for programmers to explicitly define the input of the script. This library was employed in order to allow the user to explore the possible configurations made available by the software.
- `check`: `check` is a unit testing framework written in C. It has a straightforward interface for defining unit tests helping the developer to build robust software. This library was included in the application in order to perform unit-testing on the structure we have created. This choice implies a more robust software.

### 2.5.2 Build

In order to build the executable file of our project, as well as the binary file needed to run the project unit test, we have employed GNU Make.

GNU Make is a tool which manages the creation of executables and other non-source files from a program's source files. Make learns how to create your software using a file called the `Makefile`, which lists each non-source file and how to compute it from other files.

Thanks to the definitions of rules, Make enables the user to build and install packages without knowing the details on how that is done.

Moreover, thanks to wildcards, it is easy to automate the application building process. Indeed, it first allotraduttorews to assemble each C source file in order to create the object files. Then, all of the object files are linked together, along with other libraries, in order to produce the final executable file. If the building rule is called multiple times, Make is smart enough to understand whether an object file needs to be recreated or not, making use of the already assembled objects, thus speeding up the building process.

Furthermore, Make can do much more than compiling software, for instance, the project contains rules which allow to build and open the code documentation written by the means of Doxygen.

In order to get the right flag for linking the needed external, the project employs `pkg-config`. This package collects metadata about the installed libraries on the system and easily provides it to the user. Hence, `pkg-config` takes care of where a library is located regardless of the distribution simplifying the application building process.

**2.5.2.1 Compile** To compile the project, it is possible to call the Makefile by typing:

```
1 make build
```

In this way, the executable `bin/particle-swarm-optimization` is ready to be launched.

Instead, to build the unittest, it is possible to execute the following command.

```
1 make test
```

The artifact is located in the `bin` directory and it is called `test`.

Along with the executable files, there are also scripts used to run the program within the University cluster. Each job in the cluster is handled by *PBS (Portable Bash Script)* which submits them to the scheduler. By means of a script, it is possible to tell the scheduler what resources the job requires in order to complete (e.g. number of processors, amount of memory, time to complete etc.) and the application the user wants to run.

The `run.sh` file in the `scripts` folder of the repository allows the user to submit the application to the cluster. The script has three parameters: number of processes, path of the ini file containing the program configuration and the number of threads. Once submitted with the `qsub` command, the script generates a number of docker containers equal to the number of specified processes thanks to the `mpirexec` binary. Each container runs the application in a shared network, therefore each process is able to communicate with each other. The details of the program deployment is discussed in the section 5.

The `generate_cluster_run.sh` file, contained in the `scripts` folder, is employed to generate specific runs in order to benchmark the application. In details, the shell file considers several combinations of processes, threads, nodes and places. More details are provided in section 4.

### 2.5.3 Execute

The executable file can be invoked with or without `OpenMP` and with or without `OpenMPI`. However, to fully exploit `OpenMPI`, it is recommended to execute the program `mpirexec` to spawn multiple processes of the multi-process application.

The executable file requires several arguments. Below there is an excerpt of the program output when the `--help` flag is called.

```
1 A Cooperating parallelized solution for Genetic Algorithm. A tool that takes a
2 set of continuous or discrete variables and an optimization
3 problem designed to work with them. The goal is to find the optimal solution by
4 exploiting Genetic Algorithms and the computational power offered by the
5 cluster
6
7 -m, --number-of-threads[=COUNT]      Number of threads for process
8 -u, --use-openmpi                      Use OpenMPI
9 -?, --help                            Give this help list
10 --usage                               Give a short usage message
11 -V, --version                          Print program version
```

In order to run, the application requires three parameters, two of them are optional, while one is mandatory.

The compulsory parameter is the configuration file, which needs to be provided in an `INI` file. This file, takes care of all the parameters which are needed by the Particle Swarm Optimization algorithm to run and which have been fully discussed during the introduction to the problem. The repository provides a standard `INI` file, called `ps0-data.ini`, which can be modified in order to configure algorithm so as to solve the target problem.


It is possible to specify the number of threads the program is allowed to spawn for each process with the `-m` flag, and whether to employ `MPI` primitives or not with the `-u` flag.

### 3 Problem Analysis

As explained during the introductory part, the main focus of the PSO algorithm is to find an approximated solution of a continuous optimization problem. Therefore, we have relied on some of the most relevant benchmark functions for continuous optimization. The experiments done focus mostly on six of them, which are listed below:

- **sphere function**: unimodal function suitable for single objective optimization. The single optimum is located in  $\vec{x} = \vec{0}$ . The sphere function is defined as follows:


$$\vec{x} \operatorname{argmin} f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n x_i^2 \quad (3)$$



**Figure 3:** sphere function

- **wave function**: unimodal function suitable for single objective optimization. The function does not admit a single optimum, however the fitness decreases as long as  $x$  approaches to  $-\infty$ . The wave function is defined as follows:


$$f(x, y) = x^3 + y^2 \quad (4)$$



**Figure 4:** Wave function

- **Easom function:** multimodal function suitable for single objective optimization. The single optimum is located in  $\vec{x} = \vec{\pi}$ . The Easom function is defined as:


$$f(x) = -\cos(x_1) \cos(x_2) \exp(-(x_1 - \pi)^2 - (x_2 - \pi)^2) \quad (5)$$



**Figure 5:** Easom function

- **Ackley function:** unimodal function suitable for single objective optimization. The single optimum is located in  $\vec{x} = \vec{0}$ . The Ackley function is defined as:


$$f(x) = -a \exp \left( -b \sqrt{\frac{1}{d} \sum_{i=1}^d x_i^2} \right) - \exp \left( \frac{1}{d} \sum_{i=1}^d d \cos(cx_i) \right) + a + \exp(1) \quad (6)$$



**Figure 6:** Ackley function

- **Himmelblau function:** multimodal function suitable for single objective optimization. The function presents four identical minima:  $f(\vec{x}) = f(3.0, 2.0) = f(-2.805118, 3.131312) = (-3.779319, -3.283186) = f(3.584428, -1.848126) = 0.0$  The function is defined as:

$$f(x, y) = (x^2 + y - 11)^2 + (x + y^2 - 7)^2 \quad (7)$$




**Figure 7:** Himmelblau function

- **Holder table function:** multimodal function suitable for single objective optimization. The function presents four identical minima:  $f(\vec{x}) = f(8.05502, 9.66459) = f(8.05502, 9.66459) = (8.05502, 9.66459) = f(8.05502, 9.66459) = -19.2085$ . The function is defined as:

$$f(x) = - \left| \sin(x_1) \cos(x_2) \exp \left( 1 - \frac{\sqrt{x_1^2 + x_2^2}}{\pi} \right) \right| \quad (8)$$





**Figure 8:** Holder table function

For the sake of the explainability, the functions presented above shows the two-dimensional case. Of course, such functions can scale up to as many dimensions as one desires.

Moreover, our particle swarm optimization implementation can handle also other functions. Indeed, it is possible to define the proper single objective function to optimize in the `problems.h` and `problems.c` files, specifying it in the configuration along with the `fitnessGoal`, namely, whether the function needs to be maximized or minimized.

### 3.1 Serial version of the algorithm

As can be seen from the PSO pseudocode shown in the introduction, the main steps the algorithm has to face are:

1. initialize the particles in the swarm according to the problem dimensionality;
2. exchange particles' positions among within the swarm;
3. sort the particles according to a distance measure (euclidean distance) in ascending order;
4. update the particles position and velocity.

### 3.1.1 Serial algorithm optimization

As mentioned in section 2, the program provides the possibility to either run on a single thread or on multiple threads. The multithreading functionalities exploited by the application are offered by the OpenMP library.

As a first approach, we have tried to use OpenMP directives in order to generate a thread for each loop iteration whenever it was possible.

To prevent contaminating the spaces of other threads, the actions inside the *for* loop worked on distinct data structures and variables.

For example, we inserted the relative OpenMP directive every time there was the need to loop though all the particles.

However, OpenMP *fork-join* model requires a non negligible overhead so as to spawn multiple threads which are eventually joined into the master at the end of the OpenMP block. For relatively small problems, this operation was a time-consuming procedure which leads to a significant rise in execution time with respect to the single thread model. Moreover, during the experiments we have not been able to observe the threads advantage we were hoping for. We assume that the main reason behind this non-tangible advantage are the optimization provided by `gcc` at compile time and the non-optimal thread allocation patterns performed on the cluster. Indeed, it is not rare to observe different threads being executed on the same computational unit, which clearly slows down the computation due to the overhead required by the context switching operation.

In the final version of the application, we have included the OpenMP directives only in the portion of the code where we thought it was needed, even if the advantage in terms of time were not satisfactory compared to the single threaded application.

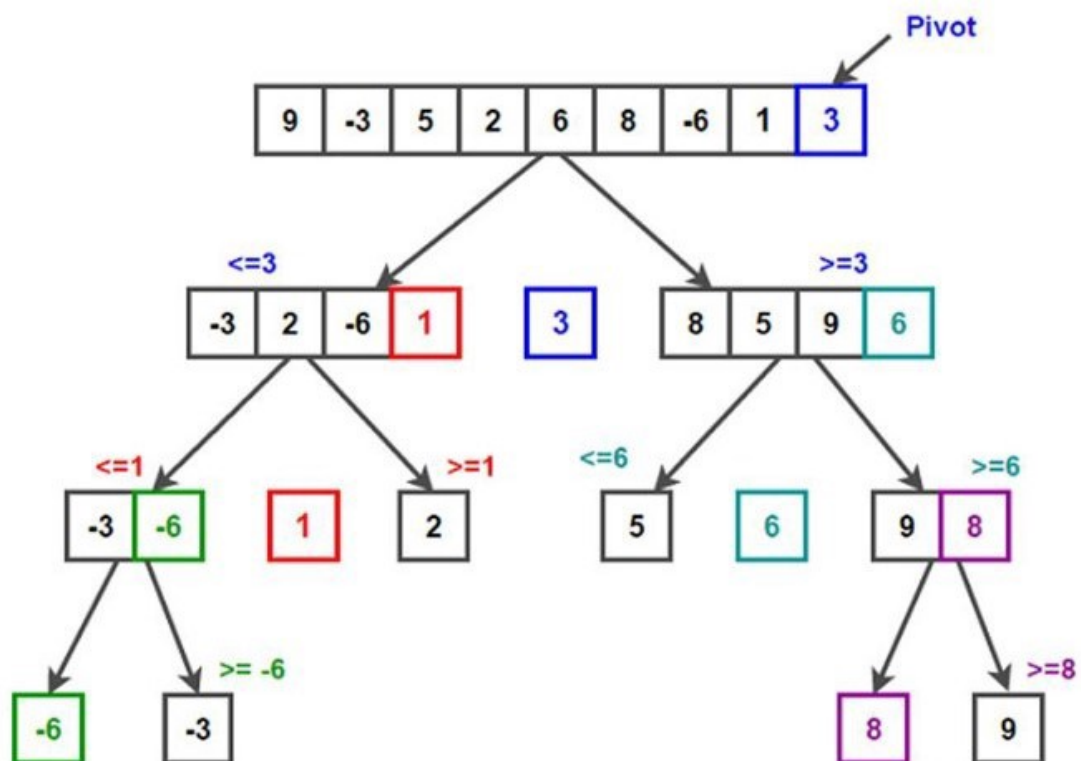
**3.1.1.1 Particles' initialization** Since only one process is involved, the initialization step is trivial, as it requires to create a given number of particles sequentially, which are immediately stored within an array data structure.

**3.1.1.2 Exchange particles' information** In the serial version of the application, each particle is stored within an array data structure, therefore with a two-level nested loop, it is possible to make every particle contribution to each other. In this case, as each particle handles a different portion of the memory, a *pragma for* directive is included.

**3.1.1.3 Sorting algorithm** Concerning the sorting algorithm, the program relies on *quicksort* (figure 9). The main reason behind this choice is the amount of parallelization this algorithm can provide. Indeed, merge sort has a better worst-case performance  $\mathcal{O}(n \log n)$  with respect to quicksort having  $\mathcal{O}(n^2)$  but requires synchronization in order to merge the partial solutions which is not required in quicksort. Moreover, its average performance is  $\mathcal{O}(n \log n)$  as for merge sort.

The parallel quicksort main working loop can be described as follows. Before going into the implementation details, in the base settings we have a pool of threads provided by the OpenMP library which can be called whenever is needed, and an array of items that needs to be sorted.

Initially, one thread selects a pivot and moves the elements of the array which are smaller than the pivot value to the left and the elements which are bigger than the pivot value to the right. The resulting two portions of the array wait until there is a free thread ready to process it.



**Figure 9:** Quicksort

**3.1.1.4 Velocity and position update** As for the initialization step, the process consists in a for loop which scans all the particles' data, and applies the formula according to the algorithm list in section 2.

## 3.2 Parallel version of the algorithm

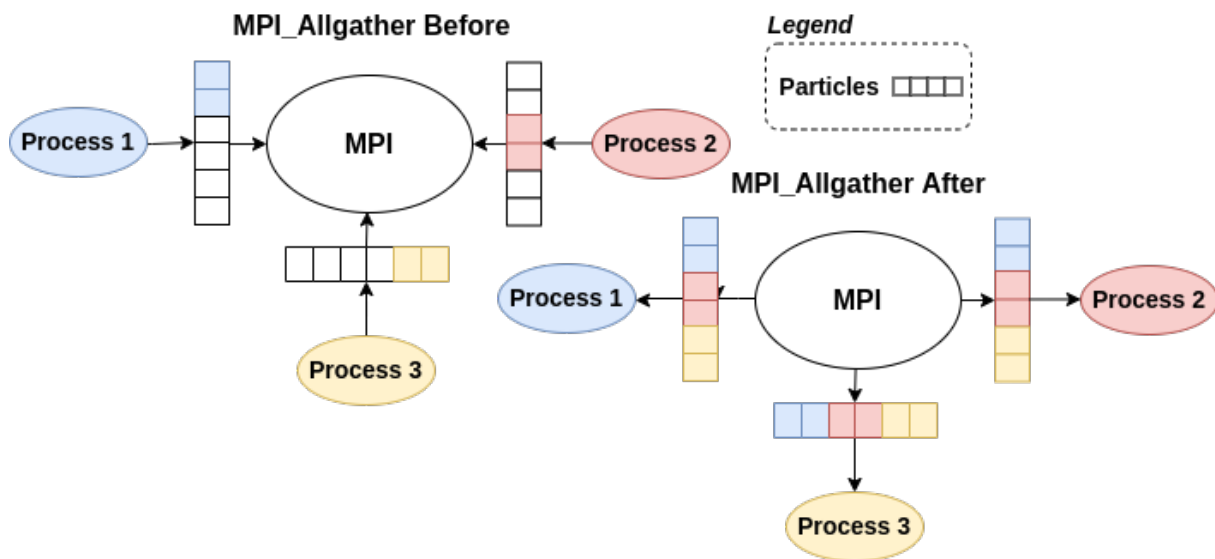
This section of the report presents how we have parallelized the algorithm in order to speed up the computation time.

In practice, we have distributed the workload among  $N$  different processes in the cluster using the *MPI* library and we have exploited multiprocessing via OpenMP for a couple of different shared-memory tasks.

### 3.2.1 Architecture

To subdivide the work and to carry out the final computation, the architecture proposed by the report focuses on the *all-to-all* parallel computational pattern (figure 10).

*All-to-all* parallel pattern, implemented using `MPI_Allgather` function, is characterized by the exchange of individual messages from every process to any other process. In this way, the program effectively uses all the computational units in order to carry out the computation, as the coordination operations are handled by MPI.



**Figure 10:** Communication schema.

### 3.2.2 Message

To send a message between different processes we created a custom MPI data type called `broadcastMessage_t`. Its purpose is to inform the receiver process about the particles' position and fitness of the sender. The structure is composed by:

- a timestamp, which is needed for logging purposes on the sqlite database;
- the current iteration of the algorithm;
- the identifier of the particle;
- the sender rank and the current solution.

On the other hand, solution is another structured datatype, which has been defined in order to convey a specific particle information. This structured datatype contains:

- the problem dimensionality;
- the fitness value of the candidate solution;
- the vector depicting the current particle position within the fitness landscape.

So as to create a variable of the previously mentioned message, we have defined a proper function called `define_datatype_broadcast_message`. This function, in turn, calls the procedure required to define a message carrying the solution type variable and a timestamp type variable.

The stratification above has simplified the `MPI_Datatype` definition process.

### 3.2.3 Communication pattern

The communication between the different processes is synchronous.

Firstly, each process takes charge of a given number of particles. In details, let  $N$  be the number of particles the user has requested to program to manage and let  $p$  be the number of processes available to `MPI`. Without the need of synchronization nor of message exchange, each process creates  $N/p$  particles and the remaining  $N \% p$  ones are split among the remaining processes. This was possible by exploiting the number of processes and the process rank in the following way:

```
1 int particlesNumberPerProcess = particlesNumber / n_processes;
2 int particlesNumberReminder = particlesNumber % n_processes;
3 int processToNumberOfParticles[n_processes];
4 for (int i = 0; i < n_processes; i++) {
5     processToNumberOfParticles[i] = particlesNumberPerProcess;
6     if (i != 0 && i <= particlesNumberReminder)
7         processToNumberOfParticles[i]++;
8 }
```


In this way, the  $i$ -th rank process has `processToNumberOfParticles[i]` particles to handle and, at the same time, knows how many particles the other processes have to manage which will be useful later when handling message communication. The possibility of coordinate the processes without synchronization calls or messages provides a relevant improvement in the application performances.

The most interesting part in the algorithm parallelization is the program segment related to the message exchange among multiple processes.

To accomplish this operation, each process embeds its own particles in an array of `define_datatype_broadcast_message`. Then, the particle information exchange happens with an `MPI_Allgather` communication primitive.

In principle, at the beginning of the algorithm execution, the set of all the particles have been distributed across all processes. However, this operation has been carried out by each process alone without the need for a `MPI_Scatter` call.

`MPI_Allgather` (figure 11) primitive is suitable for the problem since it is an *all-to-all* communication primitive and since it allow to reunite all the particles of each process into a single vector, which, at the end of the communication, will be equal for each process. A scheme illustrating the working behavior of the communication primitive employed is shown below:



**Figure 11:** Allgather

Once each process knows everything about the others, the application needs to consider the neighbor contributions in order to update the process particles' position and velocity.

At this point, each process can sort all the particles, whose position is known thanks to the `MPI_Allgather` communication, with respect to all particles proper to the process, according to the euclidean distance. In this way, for each process particle is possible to identify the  $k$ -th nearest neighbors.

Finally, by applying the position and velocity update formulas listed in equations 2 and 1, it is possible to evolve the algorithm and approach the target function optima.

Moreover, with the help of OpenMP we have parallelized the computation of the sorting algorithm as well as the loop needed in order to update the algorithm variables.

We have observed that the workload split implied by the algorithm enhance the PSO performances. A first, and trivial way to observe the advantage offered by `MPI` and optimization is to give a look at some of the expensive computation the serial algorithm has to carry out, in particular, let  $n$  be the number of particles and let  $m$  be the problem dimensionality. Moreover, during this reasoning process, we considered the average performance case and the single threaded scenario.

1. the distance computation is quadratic in the number of particles, while the euclidean distance is linear in the vectors' number of dimensions, hence the complexity is  $\mathcal{O}(mn^2)$ . In the application scenario, the distance needs to be computed a  $\mathcal{O}(n^2)$  number of times, as it is required to know the relative position of each particle with respect to all the others. Hence, the complexity grows to  $\mathcal{O}(mn^4)$ ;
2. the sorting algorithm, in the average performance scenario has a complexity which is given by  $\mathcal{O}(n \log n)$ , where  $n$  is the number of particles. This operation has to be performed with respect to all the particle in the swarm, increasing the complexity to  $\mathcal{O}(n^2 \log n)$ ;
3. finally, the particle's update is linear in the dimensionality of the problem, hence the complexity is  $\mathcal{O}(m)$ , which has to be performed for all the particle in the swarm, bringing the complexity to  $\mathcal{O}(nm)$ .

All those operations, are quite expensive in terms of time-complexity. Thanks to the workload sharing we have set up, the number each process has to manage is drastically reduced, and decreases the more processes `MPI`

has at its disposal. Let  $p$  be the number of processes; on average, each process has to manage  $n/p$  particles, while the problem dimensionality remains untouched.

The complexity then decreases as follows:

1.  $\mathcal{O}(m * (n/p)^4)$  for the distance computation;
2.  $\mathcal{O}((n/q)^2 \log(n/q))$  for the sorting operation;
3.  $\mathcal{O}((n/q)m)$  for the particles' variables update.

Despite this results being positive, we have to consider the time needed for each process to exchange their particle to each process. However, most HPC systems use *InfiniBand* interconnection, an high throughput, low latency connection among nodes in the cluster, therefore we claim that the advantage in terms of computational complexity remains legitimate since, in this scenario, the network has little impact on the application performance.

A visual proof of this statement is deeply discussed in section 4 of the report.

All the previously described operations are executed for a each iteration of the algorithm.

### 3.2.4 Logs

In order to provide a more effective way of visualizing the program behavior, we have employed a thread-safe logging utility library.

In this way, we were able to always know each process state. The logging library provides six different logging level: trace, debug, info, warn, error and fatal.

An excerpt of the logging library output is shown below:

```

1 15:27:58 DEBUG : PSODATA      :: problem dimension :: 2
2 15:27:58 DEBUG : PSODATA      :: particles number  :: 10
3 15:27:58 DEBUG : PSODATA      :: iterations number :: 10
4 15:27:58 DEBUG : PSODATA      :: neighborhood      :: 10
5 15:27:58 DEBUG : PSODATA      :: w                  :: 0.800000
6 15:27:58 DEBUG : PSODATA      :: phi 1              :: 0.300000
7 15:27:58 DEBUG : PSODATA      :: phi 2              :: 0.300000
8 15:27:58 DEBUG : PSODATA      :: max position       :: 10.000000
9 15:27:58 DEBUG : PSODATA      :: min position       :: -10.000000
10 15:27:58 DEBUG : PSODATA      :: max velocity       :: 100.000000
11 15:27:58 DEBUG : PSODATA      :: min velocity       :: -100.000000
12 15:27:58 DEBUG : PSODATA      :: fitness function   :: sphere
13 15:27:58 DEBUG : PSODATA      :: distance function  :: euclidean
14 15:27:58 DEBUG : PSODATA      :: fitness goal       :: minimum
15 15:27:58 INFO  : INIT             :: Using serial version
16 15:27:58 DEBUG : INIT             :: Initializing particle 1 in range 0-10
17 ...
18 15:27:58 DEBUG : INIT             :: Initializing particle 10 in range 0-10
19 15:27:58 DEBUG : New best global solution found with fitness value 20.918734
20 15:27:58 DEBUG : New best global solution found with fitness value 7.385129
21 15:27:58 DEBUG : COMPUTING       :: Update social best solution of particle 0
22 15:27:58 DEBUG : COMPUTING       :: Iteration 1/10 particle 1/10 velocity updated
23 15:27:58 DEBUG : COMPUTING       :: Iteration 1/10 particle 1/10 position updated
24 ...

```

```
25 15:27:58 DEBUG : New best global solution found with fitness value 4.690962
26 ...
27 15:27:58 DEBUG : COMPUTING :: Iteration 10/10 particle 10/10 position updated
28 15:27:58 INFO  : COMPUTING :: iteration 10/10 with best fitness 4.690962
29 Best fitness 4.690962
```

Moreover, all the logs have been formatted in order to comply to a common standard. In this way, during the benchmarking phase, it was possible to extract and manage logs information.

### 3.2.5 Output and SQLite

The final output of the program are best particle fitness values found by every process in the system.

In a real case scenario, one would be more interested in the candidate solution found by the application rather than in the fitness value. However, the fitness result is a suitable value to analyze in order to have a clear understanding on whether the model is improving its solution or not. Moreover, the particles position at every iteration can be stored within the SQLite database.

In this way, it is possible to recover the path the program has followed in order to build the final solution, and eventually choosing the most suitable point for the user application.



## 4 Benchmarking

After implementing the algorithm, we wanted to understand how our parallelization impacts the performance of the proposed solution.

Hence, we devised a full analysis of our algorithm performance changing the number of threads, the number of processes and the PBS process allocation pattern, in order to understand how the running time would have been affected.

### 4.1 Problem configuration

We devised a configuration file which is the same for every run, so as to have a common baseline.

The configuration is listed below:

```
1 [global]
2 problemDimension = 50
3 particlesNumber = 5000
4 iterationsNumber = 500
5 neighborhoodPopulation = 5000
6
7 [velocity]
8 w = 0.8
9 phi_1 = 0.3
10 phi_2 = 0.3
11
12 [randomBounds]
13 initMaxPosition = 500.0
14 initMinPosition = -500.0
15 initMaxVelocity = 100.0
16 initMinVelocity = -100.0
17
18 [functions]
19 fitness = sphere
20 distance = euclidean
21 fitnessGoal = minimum
```

The amount of particles and the neighborhood population are unreasonable for any known problem but they were chosen to show clearly the advantages brought by a multi process solution.

### 4.2 Cluster jobs

In order to have high-quality and trustworthy results to examine, as indicated in the repository structure, we created a script that allowed us to send thousands of tasks to the University's HPC cluster over the course of many days. The script keeps a limit on the user current submitted job in order to do not monopolize the cluster and comply with policies. Precisely, we have set a 15 process limit and every 10 seconds the script checks whether the user has 15 or more running jobs in the cluster. If it is the case, then the script waits, otherwise it submits a new job to the scheduler.

The number of tests we have ran in total is around 1280, in particular we tried every possible combination of different parameters:

- processes: chosen between [1 2 4 8 16 32 64];
- threads: chosen between [1 2 4 8 16 32 64];
- select: chosen between [1 2 3 4 5];
- places: chosen between [`pack` `scatter` `pack:excl` `scatter:excl`];

Just to provide a common baseline, a PBS select option looks like this:

```
1 select=1:ncpus=5:mem=1mb
```

Where `select` refer to the number of chunks requested to the PBS. A chunk can be seen as group of sockets, which are not necessarily on the same machine. `ncpus` instead are the number of cores per chunk the user wants to have.

As for the `select` option, here we list the different meanings for the PBS places. The option used to specify the PBS place is the following:

```
1 place=$place
```

Where `$place` can take one of the following values:

- `pack`: all the chunks are allocated within the same nodes;
- `scatter`: the chunks are located on different nodes;
- `pack:excl`: the scenario is the same as `pack`, but the chunks are reserved entirely for the user application;
- `scatter:excl`: the scenario is the same as `scatter`, but the chunks are reserved entirely fo the user application.

Since it takes 50 minutes on average for each run, we had to stretch out the submission of jobs across many days (three weeks).


### 4.3 Results

During the execution of the benchmarking phase two problems occurred:

- one student has saturated the available space in the the cluster home. For this reason, some of our runs were not able to store the results on disk at the end of the execution, hence, some computations got lost;
- we decided to submit all the jobs to the short job queue of the cluster. However, due to the problem difficulty and to the thread scheduling issues, some jobs failed as a result of *time exceed* or missing resources.

All the job configurations were tested by both members of the group in order to validate and reduce possible noise of the results.

Figure 12 shows the amount of jobs we have run and the associated time exceeded rate.




**Figure 12:** Number of time exceeded runs per process.

The presented figure highlights a correlation between the failure rate and the number of processes. Thus, we have tried to investigate the main reason behind this weird behavior.

To begin with, we have kept constant the number of processes and we have increased the number of chunks for our jobs.


Figure 13 shows the number of failed runs associated with the corresponding number of threads. As a matter of fact, the more the requested chunks, the more the cores for the job are. Hence, since the number of MPI processes is always the same, unused cores can host threads, which could be a reasonable explanation for the low amount of failed jobs in higher chunks requests.

This proof of concept highlights how the overhead paid for a continuous context switch introduced by OpenMP is higher than the performance gain due to the parallelization. Therefore, we came to the conclusion that since several optimizations are already included within modern compilers such as gcc, OpenMP introduces only an unwanted overhead for the problem that we are facing. Hence, the optimal scenario is represented by the single threaded multi-process case.



**Figure 13:** Thread and time exceed failures correlation

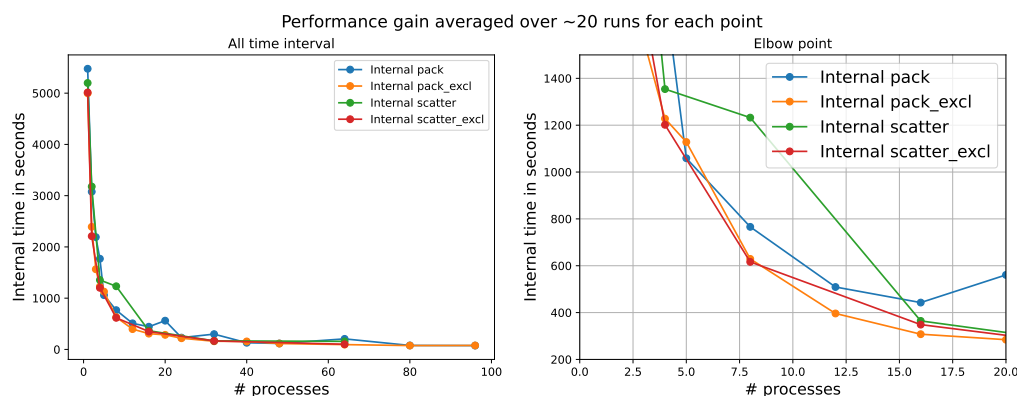
The previously described phenomena is also observable from figure 14. From there, it is possible to see that the execution time increases when the number of threads increases. Specifically, the dots in the plot represent the executed jobs while the size of the dots expresses the number of correctly terminated runs used to compute the average.



**Figure 14:** Thread and time exceed failures correlation

A part from the efficiency issues in the the multi-threaded scenario, we have deepened our benchmark analysis by considering single threaded jobs. Figure 15 shows the execution time difference between various configuration places, where *excl* means exclusive.

The plot suggests that the difference in terms of execution time between exclusive chunks and shared ones is marginal with respect to the entire time needed for the job execution. This implies that the computation time of the job is markedly higher then the scheduling time between other users' processes. For the same reason, we can claim that the overall computation time does not suffer from the network overhead. The last statement can be appreciated from the small differences in term of execution time between pack and scatter jobs executions.



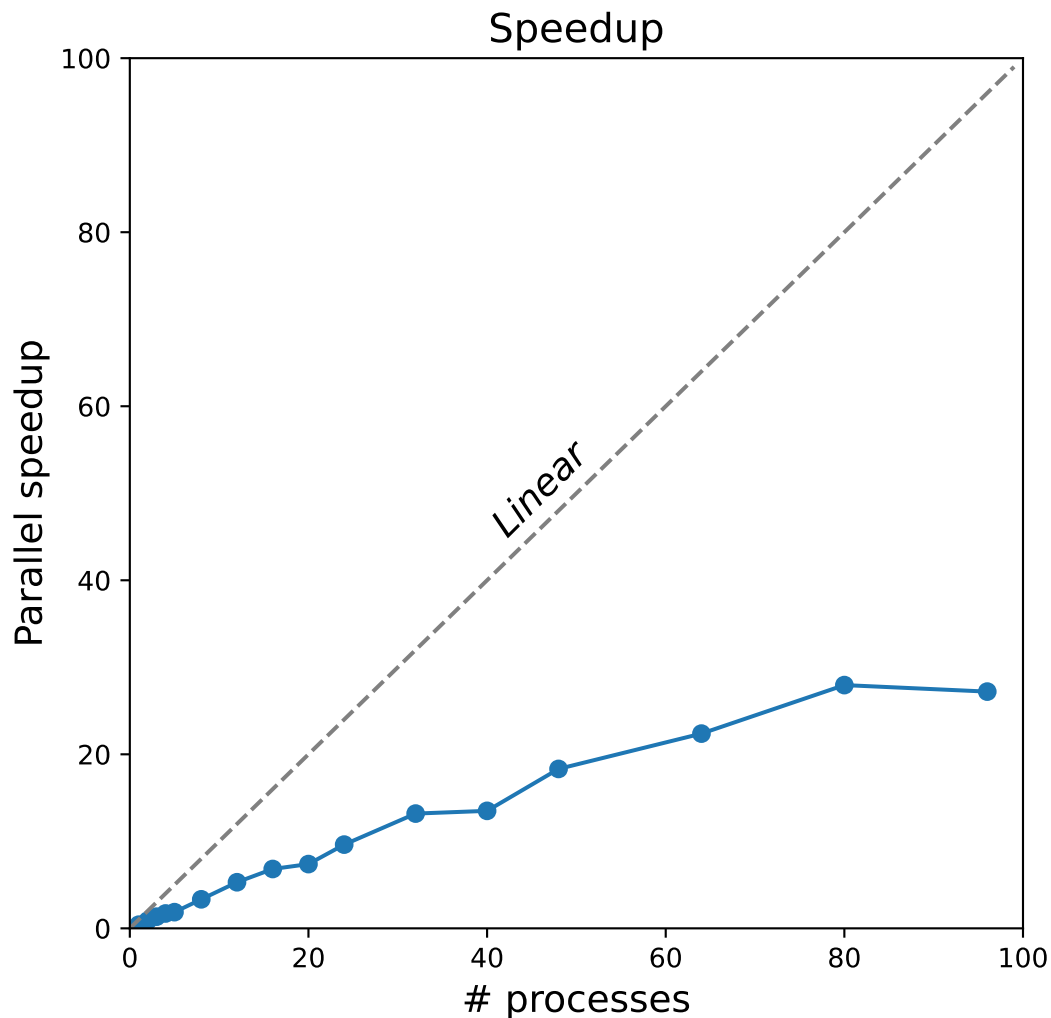
**Figure 15:** Thread and time exceed failures correlation

Furthermore, we have highlighted an elbow point in figure 15. This spot identifies the best visual tradeoff between number of processes and the execution time of the parallel solution. Indeed, for the problem configuration used for our tests, more than 16 processes do not bring an enough gain in order to motivate the expenses associated to the PBS request.

As mentioned in section 2.4, the analysis of the available similar works has required to directly manage others' code. Unfortunately, since the intricate problem was resolved in a matter of seconds, we claim that the findings are deceptive. We believe that the main reasons are the way some PSO instances perform the iterations since some of them stop whenever the solution is below an error threshold, hence performing fewer iterations than those required; as well as our difficulty in comprehending the actual behavior of other people's code, which has resulted in the development of an inconsistent algorithm.


Based on the prior results, we have chosen to take into account only the multi-process solution, and we have deepened our investigation by examining the parallel performance improvement using speedup and efficiency graphs. To begin with, the notion of scalability cannot be directly analyzed considering the problem we are optimizing. The reason for that regards the notion of *problem size*, which cannot be trivially defined. Naively, one would conclude that the problem size is doubled when the problem dimension is doubled, however, the parallelization influence is limited only in the time for the position and velocity update. On the other hand, we cannot argue that the problem size is doubled when the number of particles is doubled, since the problem persists, but holds in the opposite direction, namely only some portions of the code benefit from the parallelization. Therefore, we claim that the concept of problem size is represented by a tight coupling between the problem dimension and the number of particles. Due to this non-trivial correlation, we have decided to focus only on one

hard problem configuration and support our results with hundreds of runs.



**Figure 16:** Speedup

From figure 16 we can see that the speedup is very limited. This can be seen as a consequence of the time needed for communication between multiple processors, as at the end of each iteration all the processes must be synchronized and the number of exchanged messages is considerably high. Therefore, we claim that the overhead time we pay for the parallelization plays a relevant role, however, parallelization is still capable of providing a massive improvement in terms of time. This inevitably implies that the perfect parallelization and the ideal speedup cannot be achieved, as highlighted in the plot.



**Figure 17:** Efficiency

Likewise figure 16, figure 17 shows the efficiency curve considering only the jobs run with MPI, which is the reason why the efficiency does not start with 1. As a first consideration, the single process MPI job requires additional time compared to the serial version of the program to complete the execution, showing that the MPI, if not used properly introduces a non-negligible overhead. The efficiency plot, on the other hand, shows that the best trade-off between the number of processes employed and the speedup gained has two peaks, respectively around 3 and 12 processes. Furthermore, we can see that the more the number of processes the less the efficiency is, which means that despite taking less time to execute, it is not convenient to employ a huge number of processes. Moreover, figure 17 highlights an irregular curve, we believe that this is a consequence of an inhomogeneous cluster, hence, there are nodes which are slower compared to others.

To conclude, table 2 provides a complete overview of the program analysis.

| # Pr. | Seconds | Diff  | Speedup | Efficiency | Type    |
|-------|---------|-------|---------|------------|---------|
| 1     | 2099    | 0     | 1       | 1          | Serial  |
| 1     | 4997    | -2898 | 0.41    | 0.41       | OpenMPI |
| 2     | 2392    | -293  | 0.87    | 0.43       | OpenMPI |
| 3     | 1564    | 535   | 1.34    | 0.44       | OpenMPI |
| 4     | 1227    | 872   | 1.70    | 0.42       | OpenMPI |
| 5     | 1128    | 971   | 1.86    | 0.37       | OpenMPI |
| 8     | 629     | 1470  | 3.33    | 0.41       | OpenMPI |
| 12    | 396     | 1703  | 5.29    | 0.44       | OpenMPI |
| 16    | 307     | 1792  | 6.82    | 0.42       | OpenMPI |
| 20    | 284     | 1815  | 7.38    | 0.36       | OpenMPI |
| 24    | 218     | 1881  | 9.62    | 0.40       | OpenMPI |
| 32    | 159     | 1940  | 13.18   | 0.41       | OpenMPI |
| 40    | 155     | 1944  | 13.50   | 0.33       | OpenMPI |
| 48    | 114     | 1985  | 18.32   | 0.38       | OpenMPI |
| 64    | 93      | 2006  | 22.37   | 0.34       | OpenMPI |
| 80    | 75      | 2024  | 27.95   | 0.34       | OpenMPI |
| 96    | 77      | 2022  | 27.20   | 0.28       | OpenMPI |

**Table 2:** Speedup and efficiency table.

## 5 DevOps

In order to automate the build process of documentation, executable, and report a *Continuous Integration* workflow was designed.

### 5.1 Nix

The first step of the workflow is composed of a package manager called *Nix*. It is designed to create environments ensuring reproducibility across different machines.

The behavior of nix commands is described inside the `flake.nix` file, it also contains all the dependencies required for the execution of each possible scenario. To enforce the reproducibility in the future all required packages are fixed inside the `flake.lock` file.

The build entry points offered by nix are:

- `nix build .#report`: builds this report from Markdown files using pandoc and a latex template;
- `nix build .#particle-swarm-optimization`: builds the final binary;
- `nix build .:` same as `nix build .#particle-swarm-optimization`.



There exists also a development entrypoint that can be called using the command `nix shell`, it provides a new shell containing all dependencies required for build, testing, and report.

## 5.2 Docker

The reproducibility during execution of the final binary is provided by *containers*. A container is a sand box that contains all runtime dependencies required by the executable placed within it.

A manager for the life-cycle of containers is required, we opted for *Docker*, a state-of-the-art software used in many environments. It builds an image from the descriptor file `Dockerfile` placed in the root of the project, then from the image it is possible to create several containers that are independent of each other.

The main problem encountered was the root permission required by docker for the build operation. This prevented us to use it inside the cluster. To overcome this requirement we had to setup other technologies described in the following sections.


## 5.3 GitHub actions

A *GitHub action* is a list of commands executed when a specific event is triggered within a repository. It is possible to define multiple actions (also on different files) inside the `.github/workflows` directory.

We defined two workflows, one for the container creation and the other for the documentation compilation. They are triggered on each commit.

### 5.3.1 Container creation

The container creation workflow uses the `Dockerfile` to build up a docker image (as explained in section 5.2) that is then pushed to a *container registry* called *DockerHub*, the process is shown in figure 18. The action also includes an automation that executes unit tests, if only one them is marked as failed then the whole process is interrupted to avoid runtime errors in production.



**Figure 18:** Container creation workflow

### 5.3.2 Documentation compilation

The documentation is compiled in two formats:

- **html**: it includes all the content of this report and in addition it also provides code documentation extracted with *Doxygen* following rules defined in the *Doxyfile* file placed in the root of the project;
- **pdf**: two products have pdf format:
  - this report;
  - the sort version of this report.

Both formats are generated using a github action and the html website is hosted on *GitHub Pages* at this link.

## 5.4 Udocker

The final step of the workflow is represented by *Udocker*, it is a container manager like Docker but it does not require root permission for the execution phase.

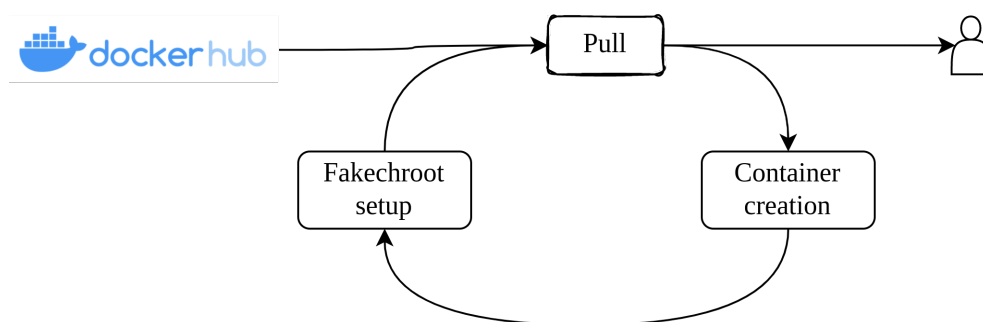
During the usage inside the cluster we encountered two main problems:

- the build phase;
- the OpenMPI communication between independent sandboxes.

### 5.4.1 Build phase

Unfortunately *udocker* does not offer primitives for containers building, it only provides one intake operation that is a pull from a container registry. For this reason we created the workflow described in section 5.3.1.

The pulled image created by GitHub and hosted on DockerHub can be used to create a container within the cluster. The sandbox is executed using a special environment thanks to *Fakechroot*. The whole process is shown in figure 19



**Figure 19:** Container pull workflow

### 5.4.2 OpenMPI communication

Normally the OpenMPI communication between different processes is not compatible with sandbox systems like docker and udocker, consequentially some additional steps are required to do that:

- the OpenMPI version installed inside the container must match the version used by the host;

- the executed container must share environment variables with the host, in particular the one that points to the dynamic library of OpenMPI.
- the executed container must share the host authentication and network.

## 6 Conclusion

Up until this point, we produced a hybrid OpenMP-MPI algorithm to solve complex continuous optimization problems, equipped with an efficient and reproducible DevOps pipeline.

We have realized that thread parallelization does not fit well all the problems. Indeed, due to the high overhead implied by the thread generation, we have observed that using OpenMP worsen the result, not providing the much-wanted speed benefit.

Benchmarking in the case of thread parallelization is a task which is far from trivial. Every system may perform differently in the presence or absence of threads. Moreover, it is hard to decide whether to parallelize or not some piece of code based on general assumptions. As an effective parallelization, we started our project by parallelizing each for loop in the code. This has resulted in a waste of resources and a worsening of performances for small-size problems. Unfortunately, the same has happened even in the case when the threads acted on the most time-consuming region of the code.

To conclude, the program we provided is suitable for single-threaded process parallelization and, as shown in the efficiency and speedup plots, it provides the best result when the number of processes is limited, as even if the computational time decreases, the more the processes the more the overhead required for the MPI communication to take place is.

### 6.1 Future Work

As a further work, it would be interesting to complement the already present architecture with different type of neighborhood and analyze which configuration brought the best results in presence of parallelization, and in terms of quality of the provided solutions. However, the scope of our project was to implement the above described parallel algorithm, which already posed significant challenges, especially because we could not base our implementation on pre-existing works.

## References

- [1] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of ICNN'95 - international conference on neural networks*, 1995, vol. 4, pp. 1942–1948 vol.4.
- [2] toddgquant, "PSO library for c." <https://github.com/toddgaunt/cpso>, 2019.
- [3] sousouho, "Succing PSO." <https://github.com/sousouhou/succinctPSO>, 2019.
- [4] kkentzo, "Pso." <https://github.com/kkentzo/pso>, 2020.
- [5] fisherling, "Pso." <https://github.com/fisherling/pso>, 2020.
- [6] A. O. S. Moraes, J. F. Mitre, P. L. C. Lage, and A. R. Secchi, "A robust parallel algorithm of the particle swarm optimization method for large dimensional engineering problems," *Applied Mathematical Modelling*, vol. 39, no. 14, pp. 4223–4241, 2015.
- [7] N. Nedjah, R. de Moraes Calazan, and L. de Macedo Mourelle, "A fine-grained parallel particle swarm optimization on many-core and multi-core architectures," in *Parallel computing technologies*, 2017, pp. 215–224.
- [8] abhi4578, "Parallelization-of-PSO." <https://github.com/abhi4578/Parallelization-of-PSO>, 2019.
- [9] LaSEEB, "Openpso." <https://github.com/abhi4578/openpso>, 2020.
- [10] pg443, "Particle-swarm-optimization-OpenMP." <https://github.com/pg443/Particle-Swarm-Optimizati-on-OpenMP>, 2021.