




ELSEVIER

Contents lists available at ScienceDirect

Computers & Geosciences

journal homepage: www.elsevier.com/locate/cageo

Case study

Acceleration of the Geostatistical Software Library (GSLIB) by code optimization and hybrid parallel programming

Oscar Peredo ^{a,b,*}, Julián M. Ortiz ^{c,b}, José R. Herrero ^d^a Advanced Laboratory for Geostatistical Computing (ALGES), University of Chile. Avda. Tupper 2069, Santiago, Chile^b Advanced Mining Technology Center (AMTC), University of Chile. Avda. Tupper 2007, Edificio AMTC, Santiago, Chile^c Department of Mining Engineering, University of Chile. Avda. Tupper 2069, Santiago, Chile^d Computer Architecture Department, Universitat Politècnica de Catalunya (UPC-BarcelonaTech), Campus Nord, Mòdul C6, Desp. 206, C/Jordi Girona 1-3, 08034 Barcelona, Catalunya, Spain

ARTICLE INFO

Article history:

Received 22 May 2015

Received in revised form

20 September 2015

Accepted 23 September 2015

Available online 26 September 2015

Keywords:

Geostatistics

Stochastic simulation

Kriging

GSLIB

Code optimization

OpenMP

MPI

ABSTRACT

The Geostatistical Software Library (GSLIB) has been used in the geostatistical community for more than thirty years. It was designed as a bundle of sequential Fortran codes, and today it is still in use by many practitioners and researchers. Despite its widespread use, few attempts have been reported in order to bring this package to the multi-core era. Using all CPU resources, GSLIB algorithms can handle large datasets and grids, where tasks are compute- and memory-intensive applications. In this work, a methodology is presented to accelerate GSLIB applications using code optimization and hybrid parallel processing, specifically for compute-intensive applications. Minimal code modifications are added decreasing as much as possible the elapsed time of execution of the studied routines. If multi-core processing is available, the user can activate OpenMP directives to speed up the execution using all resources of the CPU. If multi-node processing is available, the execution is enhanced using MPI messages between the compute nodes. Four case studies are presented: experimental variogram calculation, kriging estimation, sequential gaussian and indicator simulation. For each application, three scenarios (small, large and extra large) are tested using a desktop environment with 4 CPU-cores and a multi-node server with 128 CPU-nodes. Elapsed times, speedup and efficiency results are shown.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

The Geostatistical Software Library (GSLIB), originally presented by Deutsch and Journel (1998), has been used in the geostatistical community for more than thirty years. It contains plotting utilities (histograms, probability plots, Q–Q/P–P plots, scatter plots, location maps), data transformation utilities, measures for spatial continuity (variograms), kriging estimation and stochastic simulation applications. Among these components, estimation and simulation are two of the most used components, and can be executed with large data sets and estimation/simulation grids. Large scenarios require several minutes/hours of elapsed time to finish, due to the heavy computations involved and their sequential implementation. Since their original development, these routines have helped many researchers and practitioners in their studies, mainly due to the accuracy and performance delivered by this package. Many efforts have been proposed to accelerate or enhance the scope of the original package, WinGslib (Statiols LLC, 2001),

SGeMS (Remy et al., 2009) and HPGL (High performance geostatistics library, 2010) being the most relevant efforts. SGeMS and HPGL move away from Fortran and implement Python and C/C++ code in conventional and new algorithms. Although there is a significant gain with this change, for many practitioners and researchers, the simplicity of Fortran code and the availability of an extensive pool of modified GSLIB-based programs make it hard to abandon this package.

According to the authors' knowledge, few efforts have been reported in order to accelerate the GSLIB package by itself: analyzing, optimizing and accelerating the original Fortran routines. In this work we present case studies of accelerations performed on original GSLIB routines (in their Fortran 90 versions), using code optimization and multi-core programming techniques. We explain our methodology, in which a performance profile is obtained from the original routine, with the aim of identifying overhead sources in the code. After that, incremental modifications are applied to the code in order to accelerate the execution. OpenMP (Chandra et al., 2001) directives and MPI (Snir et al., 1998) instructions are added in the most time consuming parts of the routines. Similar experiences in other geostatistical codes have been reported in Straubhaar et al. (2013) and Peredo et al. (2014).

* Corresponding author. Currently at Telefónica I+D Chile, Chile.

E-mail address: operedo@alges.cl (O. Peredo).URL: <http://www.tidchile.cl> (O. Peredo).

```

application
1  module geostat
2      integer,allocatable :: ...
3      ...
4  end module
5
6  program main
7      use geostat
8      include 'application.inc'
9      call readparm(...)
10     call application(...)
11     stop
12 end

```

Fig. 1. Original main program for `gamv`, `kt3d`, `sgsim` and `sisim` applications.

2. GSLIB structure

According to GSLIB documentation (Deutsch and Journel, 1998), the software package is composed of a set of utility routines, compiled and wrapped as a static library named `gslib.a`, and a set of applications that call some of the wrapped routines. We will refer to these two sets as *utilities* and *applications*. Typically, a main program and two subroutines compose an *application* (Fig. 1). The first subroutine is in charge of reading the parameters from the input files, and the second subroutine executes the main computation and writes out the results using predefined output formats. Additionally, two structures of static and dynamic variables are used by the main program and each subroutine: an include file and a `geostat` module. The include file contains static variable declarations, like constant parameters, fixed length arrays and common blocks of variables. The `geostat` module contains dynamic array declarations, which will be allocated in some of the subroutines with the `allocate` instruction. A utility is self-contained allowing sharing variables with other utilities and applications through common block variable declarations.

The above-mentioned structure is common to many applications and has advantages and disadvantages. The user/programmer can easily identify each part of the code and where the main computations are occurring. However, the use of implicit typing and module/include variable declarations in the applications and utilities makes it difficult to set up a data-flow analysis (Aho et al., 2006) of all the variables in

any state of the execution. With this kind of analysis, the user/programmer can estimate the state of the variables in different parts of the code, being able to know if a variable is alive or dead at some point of execution.

From a final user perspective, this information can be seen as irrelevant. However, from a programmer's perspective, who intends to re-design some parts of the code or accelerate the overall execution, the data-flow analysis is the first step into its journey.

With these concepts in mind, in the next sections we show how to apply a methodology to accelerate GSLIB applications and utilities.

3. Methodology

Re-design: First we have to re-design the application/utility code to identify the state of each variable, array or common block during the execution. This step is necessary to enable the user/programmer to identify the scope of each variable (data-flow analysis), in order to insert OpenMP directives into the code in a smooth and easy way.

Profiling and code optimization: After re-design, we have to study the run-time behavior of the application using a profiler tool. In our case we choose the Linux-based tools `gprof` (Graham et al., 1982), `strace/ltrace` (Johnson and Troan, 2004) and `Oprofile` (Oprofile community, 2015). These tools can deliver several statistics, among the most important are the following: elapsed time per routine or line of code, number of system/library calls, number of calls per routine, number of L1/L2 cache misses per line of code and number of misspredicted branches per line of code. Each of these statistics will be described in further sections. With this information, we can identify which lines of the application or used utilities are generating overhead. We can modify some parts of the code using the profiled information. For each modification, we must re-measure the elapsed time and statistics, in order to accept or reject the modification. In this work, only hardware-independent optimizations will be included, because of portability issues.

OpenMP parallelization: Once an optimized sequential version of the application is released, we can add OpenMP directives in the most time consuming parts of the code, for example in simulation-based loops or grid-based loops. Each directive defines a parallel region, which will be executed by several threads, with a maximum defined by the user.

For each directive the user/programmer must study the data-flow of the variables inside the parallel region, and specify if the

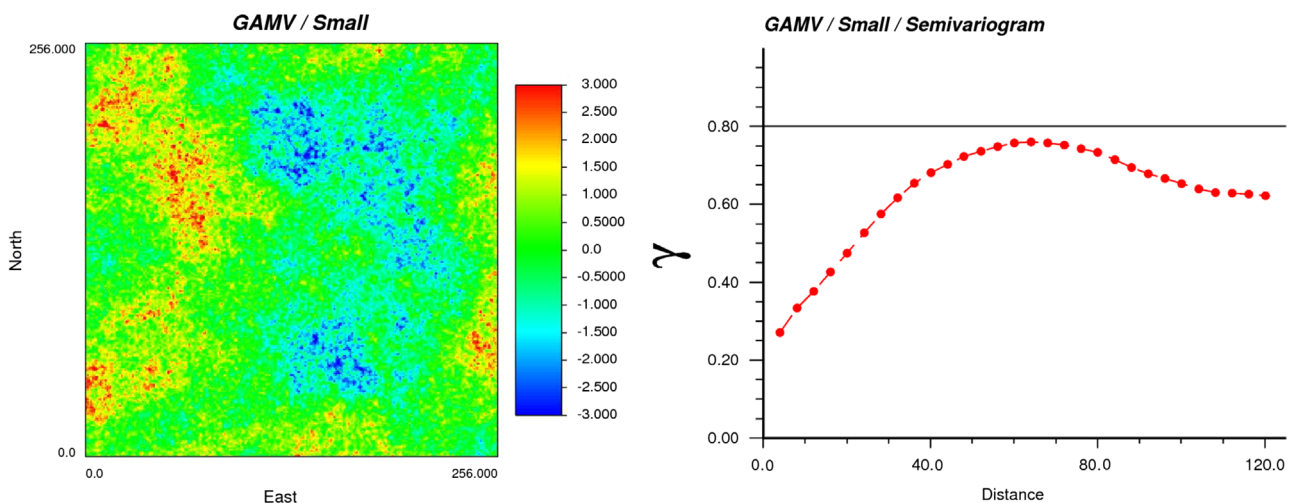


Fig. 2. Input dataset (left) and output (right) of `gamv` application using a small scenario.

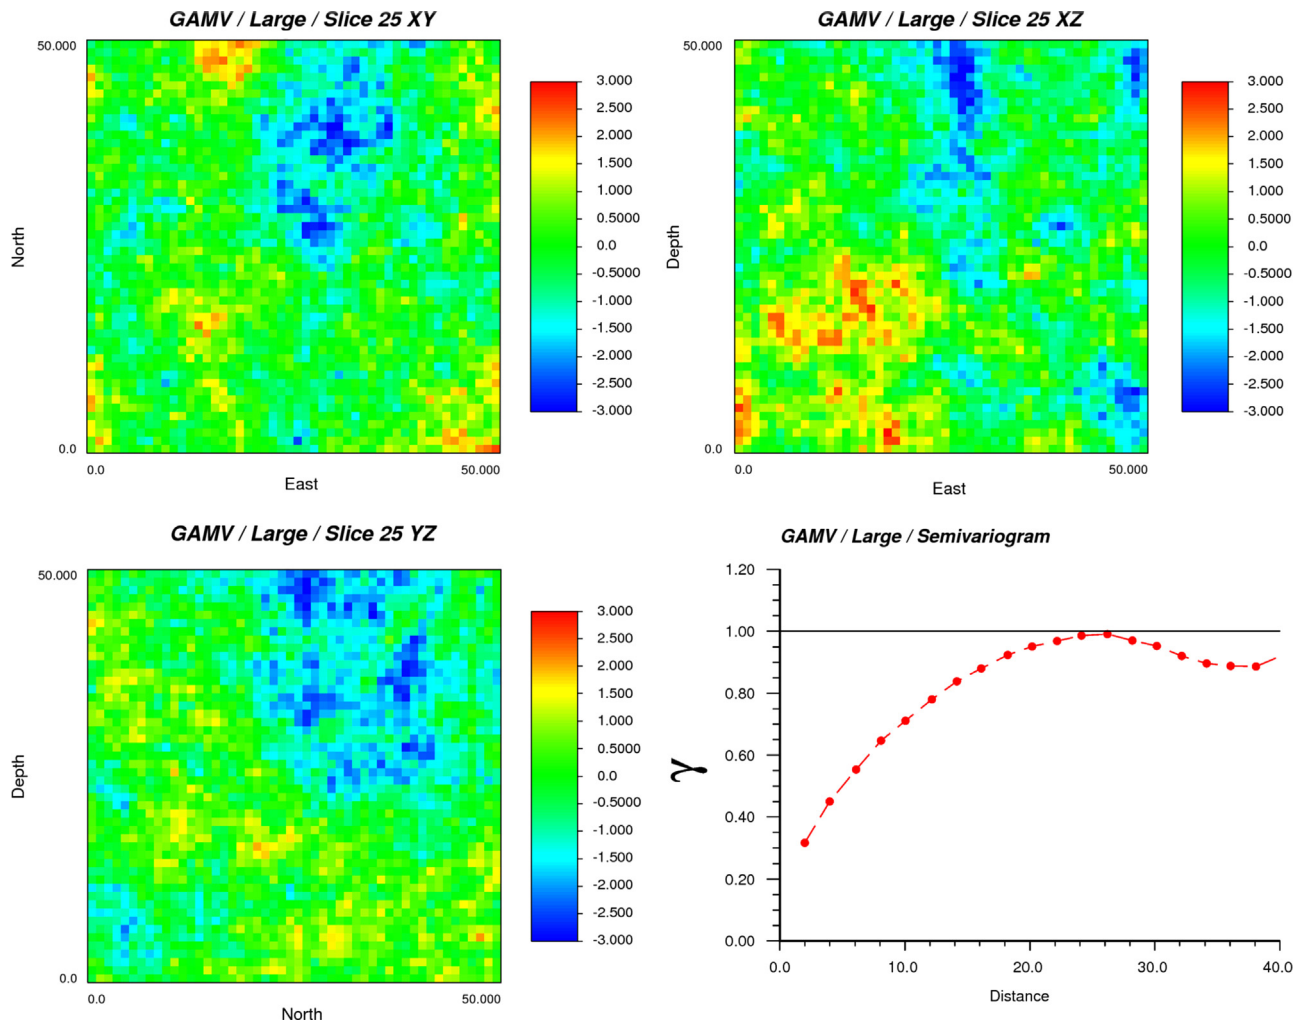


Fig. 3. Input dataset (top-left, top-right and bottom-left) and output (bottom-right) of `gamv` application using a large scenario.

Table 1

Output of `ltrace` tool applied to `gamv` application (small scenario) running in *Desktop* system. Large and extra large scenarios deliver the same information. Timing column (“seconds”) of `ltrace` may not show accurate measurements and must be taken only as an approximation reference.

Time (<code>ltrace</code> , small) (%)	Seconds	Calls	Name
77.87	112.12	2,400,020	<code>_gfortran_transfer_real</code>
12.40	17.85	480,031	<code>_gfortran_st_read</code>
9.69	13.95	480,031	<code>_gfortran_st_read_done</code>
0.04	–	–	Other routines

variables will be shared or private. If original GSLIB unmodified routines are parallelized, this analysis can be very tedious and error prone. For this reason, the first step of the methodology is fundamental in order to facilitate this analysis. Further examples of common problems if the re-design is not applied will be detailed in the next sections.

MPI parallelization: After single-node multi-thread processing has been defined, multi-node MPI processing can be added in a straightforward way once the data-flow of variables has been studied. As for OpenMP directives, the most time consuming parts of the code can be even more parallelized among distributed parallel processes.

At last, we must re-optimize the application in order to exploit efficiently all resources (memory, I/O, CPU and network) running multiple threads of execution distributed among multiple

compute nodes. When using several threads in a multi-node system, new bottlenecks and sources of overhead can arise. For this reason we must repeat the second step of the methodology, using multi-node multi-thread profiling tools (Barcelona Supercomputing Center, 2015; ZIH - TU Dresden, 2015) whenever it is possible.

4. Case study

The proposed methodology was applied to accelerate four GSLIB applications: `gamv`, `kt3d`, `sgsim` and `sisim`. We tested the final versions of the applications in two Linux-based systems: the *Server*, running SUSE operating system with multiple nodes of 2×8 -cores Intel Xeon CPU E526702.60 GHz interconnected through a fast Infiniband FDR10 network, and the *Desktop*, running openSUSE operating system with a single node of 1×4 -cores Intel Xeon CPU E312253.10 GHz. All programs were compiled using GCC `gfortran` version 4.7 and Open MPI `mpif90` version 1.8.1, supporting OpenMP version 3.0, with the option `-O3` in all cases and `-fopenmp` in the multi-thread executions.

All results are the average value of 5 runs, in order to reduce external factors in the measurement.

In the first application, `gamv`, synthetic 2D and 3D datasets of normal scored Gaussian random fields are used as base dataset. In the other three applications a real mining 3D dataset of copper grades (2376 samples) is used as base dataset. In this 3D dataset,

the lithology of the samples corresponds to Granodiorite (15%), Diorite (69%) and Breccia (16%).

Through all this section, we use the notation defined in Deutsch and Journel (1998), where the basic entities are random functions or fields (RF) which are sets of continuous location-dependent random variables (RV), such as $\{Z(\mathbf{u}): \mathbf{u} \in \Omega\}$. For each application we include a brief description of the main mathematical formulas required by the underlying algorithm.

4.1. gamv

The measurement of the spatial variability/continuity of a variable in the geographic region of study is a key tool in any geostatistical analysis. The gamv application calculates several of these measures, in an experimental way, using the available dataset as a source. Available measures to be calculated are the following: semivariogram, cross-semivariogram, covariance, correlogram, general relative semivariogram, pairwise relative semivariogram, semivariogram of logarithms, semimadogram and indicator semivariogram. The description of each measure can be found in Deutsch and Journel (1998) (III.1). Among the most used, we can mention the semivariogram, which is defined as

$$\gamma(\mathbf{h}) = \frac{1}{2N(\mathbf{h})} \sum_{i=1}^{N(\mathbf{h})} (Z(\mathbf{u}_i) - Z(\mathbf{u}_i + \mathbf{h}))^2 \tag{1}$$

where \mathbf{h} is the separation vector, $N(\mathbf{h})$ is the number of pairs separated by \mathbf{h} (with certain tolerance), $Z(\mathbf{u}_i)$ is the value at the start of the vector

(tail) and $Z(\mathbf{u}_i + \mathbf{h})$ is the corresponding end (head). In Algorithm 1 we can see the main steps of the algorithm implemented in gamv application. We can observe that the steps in this algorithm are essentially the same regardless of the measure to be calculated. For example, to calculate the semivariogram (Eq. (1)) using just one variable, one direction \mathbf{h} and 10 lags with separation $h=1.0$, first we iterate through all pairs of points in the domain Ω (loops of lines 2 and 3 of Algorithm 1), then we have ten iterations in the next loops (line 4 with $ndir=1$, $nvarg=1$ and $nlag=10$). In each iteration, we must check if some geometrical tolerances are satisfied by the current pair of points (first condition of line 6) and then we must check if the separation vector between the points, $(\mathbf{p}_i - \mathbf{p}_j)$, is similar to the separation vector \mathbf{h} multiplied by the current lag $ilag$ and the separation lag ($h=1.0$). If both conditions are fulfilled, the pseudo-routine save_statistics saves the values of the variables in study into array β . In this case, only one variable is being queried ($h_{iv}==t_{iv}$). For each type of variogram, the variables $\mathbf{V}_{i,h_{iv}}$ and $\mathbf{V}_{j,h_{iv}}$ may be transformed using different algebraic expressions. In the case of the semivariogram, we must save $(\mathbf{V}_{i,h_{iv}} - \mathbf{V}_{j,h_{iv}})^2$. Finally, using the statistics stored in β , the pseudo-routine build_variogram saves the final variogram values in vector γ , which is stored in file output.txt.

The small test scenario consists in calculating semivariogram values of 30 lags with separation $h=4.0$ using the north direction in the XY plane ($\mathbf{h} = (0, 1, 0)$), with a regularly spaced 2D dataset

Algorithm 1. Pseudo-code of gamv, measurement of spatial variability/continuity (single-thread algorithm).

```

Input: ( $\mathbf{V}, \Omega$ ): sample data base values  $\mathbf{V} \in \mathbb{R}^{|\Omega| \times m}$  defined in a 3D domain  $\Omega$ ;  $nvar$ : number of
variables in study ( $nvar \leq m$ );  $nlag$ : number of lags;  $h$ : lag separation distance;  $ndir$ :
number of directions;  $\mathbf{h}_1, \dots, \mathbf{h}_{ndir}$ : directions;  $\tau_1, \dots, \tau_{ndir}$ : geometrical tolerance parameters
(azimuth and dip);  $nvarg$ : number of variograms;
( $type_1, t_1, h_1$ ),  $\dots$ , ( $type_{nvarg}, t_{nvarg}, h_{nvarg}$ ): variogram types, tail and head variables;

1  $\beta \leftarrow \text{zeros}(nvar \times nlag \times ndir \times nvarg)$ ;
2 for  $i \in \{1, \dots, |\Omega|\}$  do
3   for  $j \in \{i, \dots, |\Omega|\}$  do
4     for  $(id, iv, il) \in \{1, \dots, ndir\} \times \{1, \dots, nvarg\} \times \{1, \dots, nlag\}$  do
5        $(\mathbf{p}_i, \mathbf{p}_j) \leftarrow ((x_i, y_i, z_i), (x_j, y_j, z_j)) \in \Omega \times \Omega$ ;
6       if  $(\mathbf{p}_i, \mathbf{p}_j)$  satisfy tolerances  $\tau_{id} \wedge \|(\mathbf{p}_i - \mathbf{p}_j) - \mathbf{h}_{id} \times il \times h\| \approx 0$  then
7          $\beta \leftarrow \text{save\_statistics}(\mathbf{V}_{i,h_{iv}}, \mathbf{V}_{i,t_{iv}}, \mathbf{V}_{j,h_{iv}}, \mathbf{V}_{j,t_{iv}}, type_{iv})$ ;
8       end
9     end
10  end
11 end
12  $\gamma \leftarrow \text{build\_variogram}(\beta)$ ;
13 write(output.txt,  $\gamma$ );

Output: Output file with  $\gamma$  values

```

```

1 ...
2 do ilag=2,nlag+2
3   if(h.ge.(xlag*real(ilag-2)-xltol).and.
4 +   h.le.(xlag*real(ilag-2)+xltol)) then
5     if(lagbeg.lt.0) lagbeg = ilag
6     lagend = ilag
7   end if
8 end do
9 ...

1 ...
2 xlaginv=1.0/xlag
3 liminf=(ceiling((h-xltol)*xlaginv)+2)
4 limsup=(floor((h+xltol)*xlaginv)+2)
5 if(lagbeg.lt.0) then
6   do ilag=liminf,limsup
7     lagbeg = ilag
8     lagend = ilag
9   end do
10 else
11   do ilag=liminf,limsup
12     lagend = ilag
13   end do
14 end if
15 ...

```

Fig. 4. Original (left) and optimized (right) `gamv` code. In the original code, we can see a source of miss-predicted branches in lines 3–4. Pre-computing the limits of the loop in line 2 and hoisting out a static condition in line 5 of original code, we can obtain an optimized version with 50% less miss-predicted branches.

Table 2

Summary of code optimizations in `gamv` application running in *Desktop* system.

Optimization	Small	Large	Extra large
Baseline	1.00 ×	1.00 ×	1.00 ×
Re-design	1.09 ×	1.11 ×	1.12 ×
Reduction of miss-predicted branches	6.01 ×	2.86 ×	2.76 ×

of $256 \times 256 \times 1$ (65,536) nodes. The large and extra large test scenarios use 20 lags with separation $h=2.0$, the east direction in the XY plane ($\mathbf{h} = (1, 0, 0)$) and a regularly spaced 3D dataset of $50 \times 50 \times 50$ (125,000) and $100 \times 100 \times 100$ (1,000,000) nodes respectively. All datasets correspond to a Gaussian random field

with an exponential covariance model. Even though a special version of `gamv` is available for regularly spaced datasets, called `gam`, we have decided to test the scalability of the optimization and parallelization using the general application `gamv`, which can also be applied without loss of generality in regular or irregular




Fig. 5. Trace views of number of instructions per cycle (IPC) generated with `extrae/paraver` applied on multi-thread version of `gamv`, small test scenario. Left column: using default schedule. Right column: using "static,32" schedule. Green color indicates low IPC values, blue color indicates high IPC values. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

Table 3

Single-node time/speedup results for *gamv*. Small scenario: 2D, 30 lags, east direction, 65,536 data points, schedule "static,32". Large scenario: 3D, 20 lags, north direction, 125,000 data points, schedule "static,1024". Extra large scenario: 3D, 20 lags, north direction, 1,000,000 data points, schedule "static,1024"

#Threads	Time (s) (Speedup)					
	Server Small 1 Node	Server Large 1 Node	Server Extra large 1 Node	Desktop Small 1 Node	Desktop Large 1 Node	Desktop Extra large 1 Node
	1 (base gslib)	121.04 (1.0 ×)	535.52 (1.0 ×)	9956.0 (1.0 ×)	166.28 (1.0 ×)	508.66 (1.0 ×)
1 (optimized)	18.96 (6.38 ×)	119.95 (4.46 ×)	3670.0 (2.71 ×)	27.63 (6.01 ×)	177.35 (2.86 ×)	3318.0 (2.76 ×)
2 (optimized)	9.60 (12.60 ×)	60.12 (8.9 ×)	1888.35 (5.27 ×)	13.86 (11.99 ×)	87.90 (5.78 ×)	1901.2 (4.81 ×)
4 (optimized)	5.04 (24.01 ×)	30.61 (17.49 ×)	960.96 (10.36 ×)	7.51 (22.14 ×)	45.28 (11.23 ×)	954.88 (9.59 ×)
8 (optimized)	3.06 (39.55 ×)	16.03 (33.40 ×)	497.60 (20.00 ×)	–	–	–
16 (optimized)	1.74 (69.56 ×)	8.78 (60.99 ×)	265.72 (37.46 ×)	–	–	–

Table 4

Multi-node time/speedup results for *gamv* with extra large scenario.

#Threads	Time (s) (Speedup)		
	Server Extra large 2 Nodes	Server Extra large 4 Nodes	Server Extra large 8 Nodes
1 (optimized)	1854.73 (5.36 ×)	947.14 (10.51 ×)	478.54 (20.80 ×)
2 (optimized)	948.30 (10.49 ×)	475.70 (20.92 ×)	243.37 (40.90 ×)
4 (optimized)	479.91 (20.74 ×)	244.62 (40.69 ×)	127.18 (78.28 ×)
8 (optimized)	251.57 (39.57 ×)	130.27 (76.42 ×)	70.05 (142.12 ×)
16 (optimized)	140.52 (70.85 ×)	73.82 (134.86 ×)	42.90 (232.07 ×)

grids. In Figs. 2 and 3 we can see the small and large datasets with their corresponding semivariogram plots.

Re-design: Applying the first step of the methodology to the original *gamv* code of Fig. 1, a slight acceleration is observed, with a speedup of $1.09 \times / 1.11 \times / 1.12 \times$ for the small/large/extra large scenarios. All variable declarations inside of *geostat* module were relocated at the beginning of the main program. After that, the routines *readparm*, *gamv* and *writeout* were manually inlined. Finally all internal routine calls were redesigned allowing only argument-based communication between the caller and the callee. No global variables which are modified at run-time are allowed through the usage of include files or module variables. Additionally, whenever is possible, the usage of the statement *implicit none* is mandatory.

Profiling and code optimization: In the second step of the methodology we are interested in observing the outputs of all tools that can extract information during execution. In Table 1 we can see the output of *ltrace* for the small scenario (the large and extra large scenarios deliver the same information). The tools *gprof* and *strace* do not give us any useful information about possible bottlenecks or overhead sources during execution in any scenario. *gprof* accounts for subroutines execution timing, but in this case all the code is self contained in *gamv*. *strace* accounts for system calls during the execution and since only small scale I/O operations are performed at the beginning and end of the execution, no significant information is collected. The output of *ltrace* is reporting a large usage of the internal library function *_gfortran_transfer_real*. This function belongs to the native GCC Fortran library *libgfortran*, and performs casting functions among single and double precision floating-point numbers. Although a major source of overhead is added by this native function, optimizing this library function is not straightforward and requires critical modifications of internal library routines. Examining *gamv* code we infer that most of this type casting is used

in critical parts where numerical accuracy is mandatory, like double-precision square roots and trigonometric functions. For this reason we decided to skip it and focus our efforts in other high-level overhead sources. If we inspect the implemented algorithm, we can observe that almost all work is done in loops of lines 2 and 3 of Algorithm 1. Instrumenting the re-designed code with *system_clock* calls before and after these loops, approximately 99.767% of the elapsed time is spent on them.

Other profiling tool at our disposal is *Oprofile*. Different hardware events can be queried with this tool, such as the number of branch instructions, number of miss-predicted branches, number of unhalted CPU cycles and number of L1/L2 data cache events (loads, stores, cache misses and evictions). For a complete description of each hardware event and architecture details about performance-monitoring counters in Intel based architectures, see Intel Corporation (2014). After extracting several line-by-line profiles with this tool, the most influential event in the overall execution time is the number of miss-predicted branch instructions. More that 50% of this event is measured in a few lines of code. In Fig. 4 we can see the original and optimized lines of code. After the proposed optimization, a speedup of $6.01 \times / 2.86 \times / 2.76 \times$ was reached (Table 2). In addition, the results delivered by the optimized application match exactly with the results of the original application.

OpenMP parallelization: For the third step of the methodology, we add the OpenMP *parallel for* pragma before the loop of line 2 in Algorithm 1, using a run-time defined schedule for loop splitting.

MPI parallelization: The fourth step consists in a blocked partition of the same loop of line 2, using a block size defined by the user. The main steps of the multi-thread multi-node algorithm are depicted in Algorithm 2. The parallelization of the optimized application still delivers similar results to the original *gamv* application, allowing some tolerance up to the 5th or 6th decimal

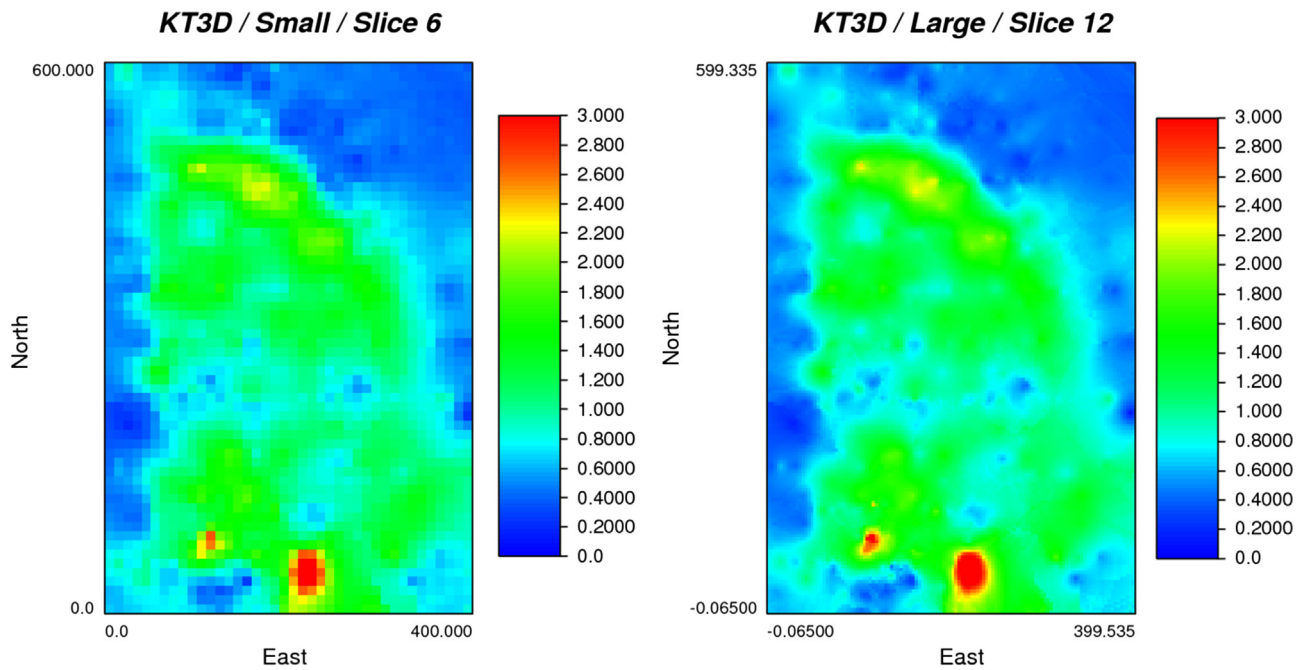


Fig. 6. Output of `kt3d` application using a small (left) and large (right) scenarios.

number (non-commutativity of floating-point operations).

In line 8 of this algorithm, we define a thread-local statistics container β_{ithread} , with the same dimensions as the global container β . This local container stores the statistics measured by each thread. After finishing the process of each local chunk of iterations from i_{start} to i_{end} , we merge the statistics into the global container, as shown in line 21. Finally, in lines 26 and 27 the master thread builds the variogram values with the global statistics collected by all threads. As mentioned before, the fourth step of the methodology is applied splitting the same loop of line 2 from Algorithm 1 using a block size B defined by the user. In line 3 of Algorithm 2 the value K indicates the number of blocks in which the total number of iterations $|Q|$ can be divided. Additionally in Algorithm 2 an outer block-based loop is defined in line 4, where each MPI task handles a block of iterations that will be divided among the local threads, as shown in lines 5 and 6. In order to collect all results among different nodes, the MPI instruction `MPI_Reduce` is launched by all MPI tasks in line 24 of Algorithm 2. With this instruction, all tasks send their local statistics β to the master task (`itask=0`) which adds all of them in the local container β^{global} .

A key step is the multi-thread loop splitting. As we can observe, if we split the loop of line 2 of the single-thread algorithm (Algorithm 1) using the default “static” schedule, an unbalanced workload will be performed by all threads (Fig. 5-left), since the outer and inner loops in lines 2 and 3 form a triangular iteration space (Navarro and Hitschfeld, 2013). For this reason, we allow the user to set the runtime schedule, depending on the size of the input dataset. In these particular test scenarios, small and large, we have inferred that the optimal schedules are “static,32” and “static,1024” respectively, evaluating all possible schedules in the form “{static,dynamic,guided},{8,16,...,512,1024}”.

Results: In Table 3 we can see the single-node average elapsed time of five executions using the optimal schedules and the corresponding speedup compared against the original GSLIB routine. Using the *Server* system, we have achieved a maximum

speedup of $69.56 \times$, $60.99 \times$ and $37.46 \times$ for the small, large and extra large respectively, using 16 CPU-cores. Using the *Desktop* system, the maximum achieved speedups with 4 CPU-cores are $22.14 \times$, $11.23 \times$ and $9.59 \times$ respectively. Using the optimized sequential time as baseline, the maximum achieved speedups are $10.89 \times$, $13.66 \times$ and $13.81 \times$ for the small, large and extra large scenarios in the *Server* system and $3.67 \times$, $3.91 \times$ and $3.47 \times$ respectively in the *Desktop* system. In terms of efficiency, defined as $\text{speedup}/\#\text{processes}$, these values represent 68%, 85% and 86% in the *Server* system and 91%, 97% and 86% in *Desktop*, for the small, large and extra large scenarios respectively.

The multi-node results can be viewed in Table 4, where the maximum achieved speedup is $232.07 \times$ for the extra large scenario with 16 CPU-cores and 8 compute nodes, a total of 128 distributed CPU-cores. Using the optimized sequential time with one node as baseline, the maximum achieved speedup is $85.54 \times$ using the same distributed 128 CPU-cores. This speedup accomplishes a 66% of efficiency using 128 CPU-cores. Comparatively, using 64 and 32 distributed CPU-cores and the optimized baseline, the average efficiency obtained is 79% and 86% respectively. The efficiencies obtained using 64 CPU-cores are 77% with 4 nodes and 16 threads each, and 81% with 8 nodes and 8 threads each. Using 32 CPU-cores, the efficiencies are 81% with 2 nodes and 16 threads each, 88% with 4 nodes and 8 threads each, and 90% with 8 nodes, 4 threads each.

The decay in efficiency observed in these results is related with the overhead added by thread synchronization and contention when the threads are writing to the local-node global container β . This container is implemented as a two dimensional shared array with size $(nvar \times nlag \times ndir \times nvar, T)$ where each thread uses the space $\beta(:, \text{ithread})$ to save their local results. With this implementation, contention may appear if $nvar \times nlag \times ndir \times nvar$ is smaller than the cache line size divided by the size in bytes of β data type (double-precision float).

Algorithm 2. Pseudo-code of gamv , measurement of spatial variability/continuity (multi-thread multi-node algorithm).

```

Input: Same inputs as Algorithm 1;  $P$ : number of MPI tasks;  $T$ : number of execution threads per
MPI task;  $S$ : runtime schedule of thread synchronization;  $B$ : block size for each MPI task;

1  $\beta \leftarrow \text{zeros}(nvar \times nlag \times ndir \times nvarg)$ ;
2  $itask \leftarrow \text{MPI\_COMM\_RANK}(\dots)$ ;
3  $K \leftarrow \left\lceil \frac{|\Omega| - B + 1}{B} \right\rceil$ ; // Count the number of blocks of size  $B$  that fit in  $|\Omega|$ 
4 for  $k \in \{1, \dots, K\}$  do
5   if  $itask == (k \bmod P)$  then
6      $(k_{start}, k_{end}) \leftarrow (k * B, \min\{(k + 1) * B, |\Omega|\})$ ;
7     /* omp parallel default(firstprivate) shared( $\mathbf{V}, \Omega$ ) */
8     foreach  $ithread \in \{1, \dots, T\}$  do
9        $\beta_{ithread} \leftarrow \text{zeros}(nvar \times nlag \times ndir \times nvarg)$ ;
10       $(i_{start}, i_{end}) \leftarrow \text{loop\_split}(ithread, k_{start}, k_{end}, S, T)$ ;
11      for  $i \in \{i_{start}, \dots, i_{end}\}$  do
12        for  $j \in \{i, \dots, |\Omega|\}$  do
13          for  $(id, iv, il) \in \{1, \dots, ndir\} \times \{1, \dots, nvarg\} \times \{1, \dots, nlag\}$  do
14             $(\mathbf{p}_i, \mathbf{p}_j) \leftarrow ((x_i, y_i, z_i), (x_j, y_j, z_j)) \in \Omega \times \Omega$ ;
15            if  $(\mathbf{p}_i, \mathbf{p}_j)$  satisfy tolerances  $\tau_{id} \wedge \|(\mathbf{p}_i - \mathbf{p}_j) - \mathbf{h}_{id} \times il \times h\| \approx 0$  then
16               $\beta_{ithread} \leftarrow \text{save\_statistics}(\mathbf{V}_{i, h_{iv}}, \mathbf{V}_{i, t_{iv}}, \mathbf{V}_{j, h_{iv}}, \mathbf{V}_{j, t_{iv}}, type_{iv})$ ;
17            end
18          end
19        end
20      end
21     $\beta \leftarrow \text{merge\_statistics}(\beta_1, \dots, \beta_T)$ ;
22  end
23 end
24  $\beta^{global} \leftarrow \text{MPI\_REDUCE}(\beta, \text{target}=0)$ ;
25 if  $itask == 0$  then
26    $\gamma \leftarrow \text{build\_variogram}(\beta^{global})$ ;
27    $\text{write}(\text{output.txt}, \gamma)$ ;
28 end

```

Output: Output file with γ values

Table 5

Output of `gprof` tool applied to `kt3d` application (small, large and extra large scenarios) running in `Desktop` system. Timing column (“seconds”) may not show accurate measurements and must be taken only as an approximation reference.

% Time (<code>gprof</code> , small)	Seconds	Calls	Name
35.12	1.77	36,946,936	cova3
30.95	1.56	28,798	ktcsol
19.25	0.97	85,615,247	sqdist
5.75	0.29	1	MAIN
4.56	0.23	28,800	srchsupr
4.17	0.21	28,801	sortem
0.02	–	–	Other routines
% Time (<code>gprof</code> , large)	Seconds	Calls	Name
30.82	32.00	665,354,374	cova3
26.79	27.82	518,358	ktcsol
16.43	17.06	1,485,946,159	sqdist
16.07	16.69	518,400	srchsupr
5.21	5.41	518,401	sortem
4.61	4.79	1	MAIN
0.06	–	–	Other routines
% Time (<code>gprof</code> , extra large)	Seconds	Calls	Name
27.07	63.96	1,330,982,767	cova3
26.05	61.54	1,036,800	srchsupr
23.84	56.32	1,036,728	ktcsol
14.18	33.50	2,968,320,694	sqdist
4.52	10.69	1,036,801	sortem
4.27	10.08	1	MAIN
0.07	–	–	Other routines

4.2. `kt3d`

The kriging algorithm was introduced to provide minimum error-variance estimates of unsampled locations using available data. The traditional application of kriging is to provide a regular grid of estimates that acts as a low-pass filter that tends to smooth out details and extreme values of the original dataset. It is extensively explained in [Deutsch and Journel \(1998\)](#) (IV.1). Different versions of the algorithm are available in the `kt3d` application. Simple Kriging (SK), in its stationary version, aims to obtain a linear regression estimator $Z_{SK}^*(\mathbf{u})$ defined as

$$Z_{SK}^*(\mathbf{u}) = \sum_{\alpha=1}^n \lambda_{\alpha}^{(SK)}(\mathbf{u})Z(\mathbf{u}_{\alpha}) + \left[1 - \sum_{\alpha=1}^n \lambda_{\alpha}^{(SK)}(\mathbf{u}) \right] m \quad (2)$$

where $m = E\{Z(\mathbf{u})\}$, $\forall \mathbf{u}$ is the location-independent (constant) expected value of $Z(\mathbf{u})$ and $\lambda_{\alpha}^{(SK)}(\mathbf{u})$ are weights given by the solution of the system of normal equations:

$$\sum_{\beta=1}^n \lambda_{\beta}^{(SK)}(\mathbf{u})C(\mathbf{u}_{\beta}, \mathbf{u}_{\alpha}) = C(\mathbf{u}, \mathbf{u}_{\alpha}), \quad \alpha = 1, \dots, n \quad (3)$$

with $\{C(\mathbf{u}_{\alpha}, \mathbf{u}_{\beta})\}_{\alpha, \beta=0, 1, \dots, n}$ being the covariance matrix adding the sampled data $\mathbf{u}_0 = \mathbf{u}$. With the stationary assumption, the covariance can be expressed as a function $C(\mathbf{h}) = C(\mathbf{u}, \mathbf{u} + \mathbf{h})$. Ordinary Kriging (OK) filters the mean m from the SK estimator of Eq. (2), imposing that the sum of weights is equal to one. The resulting OK estimator is as

$$Z_{OK}^*(\mathbf{u}) = \sum_{\alpha=1}^n \lambda_{\alpha}^{(OK)}(\mathbf{u})Z(\mathbf{u}_{\alpha}) \quad (4)$$

with the weights $\lambda_{\alpha}^{(OK)}$ solution of the extended system:

$$\sum_{\beta=1}^n \lambda_{\beta}^{(OK)}(\mathbf{u})C(\mathbf{u}_{\beta} - \mathbf{u}_{\alpha}) + \mu(\mathbf{u}) = C(\mathbf{u} - \mathbf{u}_{\alpha}), \quad \alpha = 1, \dots, n \quad (5)$$

$$\sum_{\beta=1}^n \lambda_{\beta}^{(OK)} = 1 \quad (6)$$

with $\mu(\mathbf{u})$ being a Lagrange parameter associated with the second constraint (6). It can be shown that OK amounts to re-estimating, at each new location \mathbf{u} , the mean m as used in the SK expression. Thus the OK estimator $Z_{OK}^*(\mathbf{u})$ is, in fact, a simple kriging estimator where the constant mean value m is replaced by the location-dependent estimate $m^*(\mathbf{u})$ (non-stationary algorithm with varying mean and constant covariance). Other kriging algorithms available add different trend models in 1, 2 or 3D. These last algorithms are not supported yet by the multi-thread proposed version of `kt3d`. The main steps of the kriging algorithms can be viewed in [Algorithm 3](#). The super block spatial search strategy ([Deutsch and Journel, 1998](#), II.4) is implemented in order to accelerate the search of neighbor values in cases where large datasets are available (lines 1 and 3). Once the neighbor values \mathbf{N} are available, the

Table 6

Output of `ltrace` tool applied to `kt3d` application (small, large and extra large scenarios) running in `Desktop` system. Timing column (“seconds”) may not show accurate measurements and must be taken only as an approximation reference. In this case it shows the measurements running the application approximately by two minutes.

% Time (<code>ltrace</code> , small)	Seconds	Calls	Name
97.54	125.14	1,992,512	expf
1.44	1.84	19,046	_gfortran_transfer_real
1.02	–	–	Other routines
% Time (<code>ltrace</code> , large)	Seconds	Calls	Name
97.93	125.81	2,912,965	expf
1.79	2.30	19,046	_gfortran_transfer_real
0.28	–	–	Other routines
% Time (<code>ltrace</code> , extra large)	Seconds	Calls	Name
99.65	125.51	2,675,566	expf
0.14	0.17	19,046	_gfortran_transfer_real
0.21	–	–	Other routines

covariance matrix values $C(\mathbf{u}_\beta - \mathbf{u}_\alpha)$ must be assembled into the matrix \mathbf{A} (line 4). Additionally, the right-hand side is composed of the covariance values $C(\mathbf{u} - \mathbf{u}_\alpha)$ and is stored in the array \mathbf{b} (line 5). With the left and right-hand sides \mathbf{A} and \mathbf{b} already computed, the weight vector λ can be calculated solving the SK linear system (2) or the OK linear system (5)–(6) (line 6). Finally the estimated value is computed using the weights λ (line 7) and stored in a file (line 8).

`sqdist` calculates a squared anisotropic distance between two points. `sortem` implements a sorting algorithm for floating-point single precision arrays in an ascending order. `srchsupr` implements a spatial search using a super block strategy. `ktsol` solves a floating-point double precision linear system by gaussian elimination with partial pivoting.

In terms of elapsed time, a predominance of `cova3` is observed in the three scenarios, with 35%, 30% and 27% in each case.

The routine `srchsupr` shows a steady increase of participation

Algorithm 3. Pseudo-code of `kt3d`, kriging 3D estimation (single-thread algorithm).

```

Input: ( $\mathbf{V}, \Omega$ ): sample data base values defined in a 3D domain  $\Omega$ ;  $r$ : radius of search (neighbour of
node to estimate);  $\beta$ : super block strategy parameters;  $\gamma$ : structural variographic models;

1  $\mathbf{H} \leftarrow \text{set\_spatial\_hash}(\mathbf{V}, \Omega, \beta, r)$ ;
2 for  $i \in \{1, \dots, |\Omega|\}$  do
3    $\mathbf{N} \leftarrow \text{search\_spatial\_hash}(i, \mathbf{H}, \mathbf{V})$ ;
   /*  $\mathbf{N} \subset \mathbf{V}$ , neighbourhood values of node  $i$  */
4    $\mathbf{A} \leftarrow \text{assemble\_covariance\_matrix}(\mathbf{N}, \gamma)$ ;
5    $\mathbf{b} \leftarrow \text{assemble\_rhs\_vector}(\mathbf{N}, \gamma)$ ;
6    $\lambda \leftarrow \text{solve\_system}(\mathbf{A}, \mathbf{b})$ ;
7    $y \leftarrow \text{compute\_estimation}(\lambda, \mathbf{N})$ ;
8   write(output.txt, y);
9 end

Output: Output file output.txt with estimated values for all nodes

```

The small test scenario consists in performing an Ordinary Kriging estimate in a grid of $40 \times 60 \times 12$ nodes (28,800), representing a 3D volume of $400 \times 600 \times 120$ (m^3) with a search radius of 100 (m) and a maximum of 50 neighbors to include in the local estimation. The large and extra large test scenarios have the same parameters as the small scenario but using finer grids of $120 \times 180 \times 24$ and $120 \times 180 \times 48$ nodes (518,400 and 1,036,800 respectively) in the same 3D volume. All estimation grids use the same variographic inputs, which are a spherical plus an exponential structure, and the same input dataset (2376 copper grade 3D samples). In Fig. 6 we can see slices of the small and large kriging outputs.

Re-design: Applying the first step of the methodology to the original `kt3d` code as in the `gamv` application, a slight acceleration is observed, with a speedup of $1.03 \times / 1.03 \times / 1.02 \times$ for the small/large/extra large scenarios.

Profiling and code optimization: For the second step of the methodology, `gprof` tool delivers information depicted in Table 5 for the small (top), large (medium) and extra large (bottom) scenarios. All these routines are utilities from the `gslib.a` library. In the three scenarios, the most called routines are `cova3` and `sqdist`, where `cova3` calculates the covariance value $C(\mathbf{u}, \mathbf{u} + \mathbf{h})$ according to the variographic model passed as an input, and

from 4% to 26% of the elapsed time, which is proportional to the scenario size.

In Table 6 we can observe the output of `ltrace` for the three scenarios. In all of them, the library function `expf` is executed by the application several times, however its low-level optimization is out of the scope of this work. The tool `strace` does not report any useful information about possible bottlenecks or overhead sources during execution in all scenarios.

Analyzing `Oprofile` outputs, we can observe a large amount of branch instructions executed in `cova3` routine. A basic optimization applied to the code is to specialize this routine for the input parameters. In this case, the variographic structures passed as inputs define which branch instructions will be executed. The first structure is a spherical model and the second is an exponential model. In Fig. 7 we can observe the original and specialized routine `cova3`, using these two variographic structures as an input. A speedup of $1.06 \times / 1.05 \times / 1.05 \times$ is obtained after this optimization for the small/large/extra large scenarios. Additionally, removing dead code not used by the multi-thread version of the application can speed up the execution time. Using preprocessor directives as `#ifdef DEBUG` or `#ifdef _OPENMP` allows us to remove all unnecessary code at compilation time and reduce the

<pre> 1 ... 2 cova = 0.0 3 do is=1,nst(ivarg) 4 ist = 1start + is - 1 5 if(ist.ne.1) then 6 ir = min((irot+is-1),MAXROT) 7 hsqd=sqdist(x1,y1,z1,x2,y2,z2, 8 + ir,MAXROT,rmat) 9 end if 10 h = real(dsqrt(hsqd)) 11 if(it(ist).eq.1) then 12 hr = h/aa(ist) 13 if(hr.lt.1.) then 14 cova=cova + 15 + cc(ist)*(1.-hr*(1.5-.5*hr*hr)) 16 end if 17 else if(it(ist).eq.2) then 18 cova = cova + 19 + cc(ist)*exp(-3.0*h/aa(ist)) 20 else if(it(ist).eq.3) then 21 ... 22 else if(it(ist).eq.4) then 23 ... 24 else if(it(ist).eq.5) then 25 ... 26 endif 27 end do 28 return 29 end 30 ... </pre>	<pre> 1 ... 2 cova = 0.0 3 is=1 4 ist = 1start + is - 1 5 h = real(dsqrt(hsqd)) 6 hr = h/aa(ist) 7 if(hr.lt.1.) then 8 cova=cova+cc(ist)*(1.-hr*(1.5-.5*hr*hr)) 9 end if 10 is=2 11 ist = 1start + is - 1 12 ir = min((irot+is-1),MAXROT) 13 hsqd=sqdist(x1,y1,z1,x2,y2,z2,ir,MAXROT,rmat) 14 h = real(dsqrt(hsqd)) 15 cova = cova + 16 + cc(ist)*exp(-3.0*h/aa(ist)) 17 return 18 end 19 ... </pre>
--	---

Fig. 7. Original (left) and specialized (right) *cova3* utility. The specialized code works for two additive variographic structures ($nst(ivarg) = 2$), spherical ($it(1) = 1$) and exponential ($it(2) = 1$), and performs a full unroll of the main loop.

Table 7

Summary of code optimizations in *kt3d* application running in *Desktop* system.

Optimization	Small	Large	Extra large
Baseline	1.00 ×	1.00 ×	1.00 ×
Re-design	1.03 ×	1.03 ×	1.02 ×
Specialization of <i>cova3</i>	1.06 ×	1.05 ×	1.05 ×
Remove dead code	1.11 ×	1.08 ×	1.07 ×
Unbuffered binary writing	1.12 ×	1.10 ×	1.06 ×

Table 8

Single-node time/speedup results for *kt3d*. Small scenario: $40 \times 60 \times 12$ grid nodes. Large scenario: $120 \times 180 \times 24$ grid nodes. Extra large scenario: $120 \times 180 \times 48$ grid nodes.

#Threads	Time (s) (Speedup)					
	Server Small 1 Node	Server Large 1 Node	Server Extra large 1 Node	Desktop Small 1 Node	Desktop Large 1 Node	Desktop Extra large 1 Node
1 (base gslib)	5.69 (1.0 ×)	108.62 (1.0 ×)	244.74 (1.0 ×)	5.66 (1.0 ×)	112.27 (1.0 ×)	246.76 (1.0 ×)
1 (optimized)	5.10 (1.11 ×)	102.94 (1.05 ×)	227.78 (1.07 ×)	5.08 (1.12 ×)	102.06 (1.10 ×)	232.79 (1.06 ×)
2 (optimized)	2.59 (2.19 ×)	52.18 (2.08 ×)	118.03 (2.07 ×)	2.76 (2.05 ×)	52.85 (2.12 ×)	124.03 (1.98 ×)
4 (optimized)	1.33 (4.27 ×)	26.64 (4.07 ×)	60.67 (4.03 ×)	1.43 (3.95 ×)	28.30 (3.96 ×)	65.06 (3.79 ×)
8 (optimized)	0.71 (8.01 ×)	13.90 (7.81 ×)	32.10 (7.62 ×)	–	–	–
16 (optimized)	0.41 (13.87 ×)	7.50 (14.48 ×)	17.62 (13.88 ×)	–	–	–

Table 9
Multi-node time/speedup results for `kt3d`.

#Threads	Time (s) (Speedup)		
	Server Extra large 2 Nodes	Server Extra large 4 Nodes	Server Extra large 8 Nodes
1 (optimized)	123.10 (1.98 ×)	65.49 (3.73 ×)	33.55 (7.29 ×)
2 (optimized)	61.21 (3.99 ×)	31.99 (7.65 ×)	17.24 (14.19 ×)
4 (optimized)	31.75 (7.71 ×)	17.24 (14.19 ×)	9.80 (24.97 ×)
8 (optimized)	17.49 (13.99 ×)	9.99 (24.49 ×)	6.15 (39.79 ×)
16 (optimized)	10.28 (23.80 ×)	6.47 (37.82 ×)	4.31 (56.78 ×)

number of instruction cache misses. A speedup of $1.11 \times / 1.08 \times / 1.07 \times$ is obtained after this optimization (Table 7). Another easy modification involves the use of a buffer array to store the estimated results in each grid node, followed by one single `write` instruction. In line 8 of Algorithm 3 we can observe that a `write` call is performed for each grid node. A slight acceleration is obtained using the buffering technique ($1.11 \times / 1.10 \times / 1.09 \times$), however this feature slows down the execution using multiple threads. A possible reason for this behavior might be contention generated by false sharing (Culler et al., 1998).

False sharing can happen because the buffer array must be shared among all threads, even if they write in different memory locations of it. The usage of the shared array adds contention in the CPU-memory bus, in order to synchronize and manage all threads, for this reason we decide to skip this optimization and leave the code as it is. In order to accelerate the I/O operations, a binary format in the output file is adopted, adding the option `form= 'UNFORMATTED'` in the `open` instruction. A speedup of $1.12 \times / 1.10 \times / 1.06 \times$ is observed using this modification.

Regarding the number of cache misses and evictions in the L1 and L2 cache levels, the routine `srchsupr` shows the largest number of these events, which is in concordance with the size of scenario used in the tests. As shown in the previous paragraph with the tool `gprof`, the elapsed time of the routine `srchsupr` becomes larger if the size of the scenario increases. If the estimation grid size increases, the number of cache misses and eviction events increases proportionally. In order to optimize these features the access pattern can be improved by using a single container for the contiguous super block offsets stored in the arrays `ixsbtostr(1:N)`, `iysbtostr(1:N)` and `izsbtostr(1:N)`. The container can be denoted `ixyzsbtostr(1:3N)` and it can store the three values for a particular index i contiguously, following the access pattern implemented in `srchsupr`. Although this modification reduces the number of cache misses and evictions approximately 10–15% in the large and extra large scenarios, its effect in the total elapsed time is negligible. For this reason we decided to skip this optimization and leave the code of the original routine. Significant effects of this optimization probably can be observed using even larger scenarios, where the number of cache misses and evictions dominates the bottleneck factors of the execution. As for the `gamv` application, the optimized `kt3d` application delivers the same results as the original `kt3d` application.

OpenMP parallelization: For the third step of the methodology, the OpenMP `parallel for` pragma is added in the loop of line 2 from Algorithm 3. The default schedule of loop splitting is used.

MPI parallelization: The fourth step of the methodology consists in using MPI to split even more the same loop. With MPI task-based combined with OpenMP thread-based loop splittings, multi-node processing is available using all CPU-cores in each compute node. The parallelization of the optimized application still delivers

similar results to the original GSLIB application, allowing some tolerance up to the 5th or 6th decimal number (non-commutativity of floating-point operations).

The main steps of the multi-thread multi-node algorithm are depicted in Algorithm 4. In line 3, each MPI task selects a balanced share of the iterations of loop 2 from Algorithm 3. After that, in line 6 of Algorithm 4 each local thread opens a file `output-itag.txt` where it will store the estimated values of each local chunk of iterations from j_{start} to j_{end} in the loop of lines 8–15. The value of `itag` is a unique identifier of each thread in each compute node. The rest of the algorithm is the same as the single-thread version.

Results: Analyzing the single-node results of Table 8, the maximum achieved speedup is $13.87 \times$ for the small scenario, $14.48 \times$ for the large scenario and $13.88 \times$ for the extra large scenario using 16 CPU-cores in the *Server* system. Using the *Desktop* system, the maximum achieved speedups with 4 CPU-cores are $3.95 \times$, $3.96 \times$ and $3.79 \times$ respectively. Using the optimized sequential time as baseline, the maximum achieved speedups are $12.43 \times$, $13.72 \times$ and $12.92 \times$ for the small, large and extra large scenarios in the *Server* system and $3.55 \times$, $3.60 \times$ and $3.57 \times$ respectively in the *Desktop* system. In terms of efficiency (speedup/#CPU-cores), these values represent 77%, 85% and 81% in the *Server* system and 88%, 90% and 89% in *Desktop*, for the small, large and extra large scenarios respectively.

Regarding the multi-node results of Table 9, the maximum achieved speedup is $56.78 \times$ for the extra large scenario with 16 CPU-cores and 8 compute nodes, a total of 128 distributed CPU-cores. Using the optimized sequential single-node time as baseline, the maximum achieved speedup is $52.84 \times$ using the same distributed 128 CPU-cores. This speedup accomplishes a 41% of efficiency using 128 CPU-cores. Comparatively, using 64 and 32 distributed CPU-cores and the optimized baseline, the average efficiency obtained is 56% and 71% respectively. The efficiencies obtained using 64 CPU-cores are 55% with 4 nodes and 16 threads each, and 57% with 8 nodes and 8 threads each. Using 32 CPU-cores, the efficiencies are 69% with 2 nodes and 16 threads each, 71% with 4 nodes and 8 threads each, and 72% with 8 nodes, 4 threads each.

The decay in efficiency observed in these results is related with overhead generated by thread synchronization and bus contention when the threads are writing to disk.

No further overhead is added by network contention since no MPI messages are shared between tasks during the main computations. Further tests are needed in order to check if multiple threads can use the I/O library buffer in a contention-free way, and also to explore if external OS factors can affect the behavior of the threads when writing to disk.

Algorithm 4. Pseudo-code of `kt3d`, kriging 3D estimation (multi-thread multi-node algorithm).

```

Input: Same inputs as Algorithm 3;  $P$ : number of MPI tasks;  $T$ : number of execution threads per
MPI task;  $S$ : runtime schedule of thread synchronization;

1  $itask \leftarrow \text{MPI\_COMM\_RANK}(\dots)$ ;
2  $\mathbf{H} \leftarrow \text{set\_spatial\_hash}(\mathbf{V}, \Omega, \beta, r)$ ;
3  $(i_{start}, i_{end}) \leftarrow \text{loop\_split}(itask, 1, |\Omega|, \text{"static"}, P)$ ;
   /* omp parallel default(firstprivate) shared( $\mathbf{H}, \mathbf{V}, \Omega$ ) */
   /* omp for schedule( $S$ ) ,  $S = \text{"static"}$  */
4 foreach  $ithread \in \{1, \dots, T\}$  do
5    $itag \leftarrow ithread + itask \times T$ ;
6    $\text{open}(\text{output-}itag.\text{txt})$ ;
7    $(j_{start}, j_{end}) \leftarrow \text{loop\_split}(ithread, i_{start}, i_{end}, S, T)$ ;
8   for  $j \in \{j_{start}, \dots, j_{end}\}$  do
9      $\mathbf{N} \leftarrow \text{search\_spatial\_hash}(j, \mathbf{H}, \mathbf{V})$ ;
       /*  $\mathbf{N} \subset \mathbf{V}$ , neighbourhood values of node  $j$  */
10     $\mathbf{A} \leftarrow \text{assemble\_covariance\_matrix}(\mathbf{N}, \gamma)$ ;
11     $\mathbf{b} \leftarrow \text{assemble\_rhs\_vector}(\mathbf{N}, \gamma)$ ;
12     $\boldsymbol{\lambda} \leftarrow \text{solve\_system}(\mathbf{A}, \mathbf{b})$ ;
13     $y \leftarrow \text{compute\_estimation}(\boldsymbol{\lambda}, \mathbf{N})$ ;
14     $\text{write}(\text{output-}itag.\text{txt}, y)$ ;
15  end
16 end

```

Output: Output files `output- $itag$.txt` with estimated values for all nodes

4.3. `sgsim`

The application `sgsim` implements the sequential gaussian simulation algorithm, as described in [Deutsch and Journel \(1998\)](#) (V.2.3). It is considered as the most straightforward algorithm for generating realizations of a multivariate Gaussian field. Its main steps can be viewed in [Algorithm 5](#). The first step is to transform the original dataset into a standard normally distributed dataset (line 1). Then a random path must be generated, visiting all nodes

of the simulation grid, for each simulation to be calculated (line 3). At each location $ixyz$, a conditional cumulative distribution function (ccdf) must be estimated by simple or ordinary kriging (Eqs. (2) and (5)–(6)), and then a random variable must be drawn from the generated ccdf (lines 7 and 8). The next step is to translate the simulated value in the normal distribution to the original distribution of the sample data (line 9). Finally, the back-transformed scalar result is stored in a file `output.txt` using a system call (`write`) for each nodal value (line 10).

Algorithm 5. Pseudo-code of `sgsim`, sequential gaussian simulation program (single-thread algorithm).

```

Input: ( $\mathbf{V}, \Omega$ ): sample data base values defined in a 3D domain;  $\gamma$ : structural variographic models;  $\kappa$ :
kriging parameters (radius, max number of neighbours and others);  $\tau$ : seed for pseudo-random
number generator;  $N$ : number of generated simulations; output.txt: output file

1  $\mathbf{Y} \leftarrow \text{normal\_score}(\mathbf{V})$ ;
2 for  $i \text{ sim} \in \{1, \dots, N\}$  do
3    $\mathcal{P} \leftarrow \text{create\_random\_path}(\Omega, \tau)$ ;
4    $\mathbf{Y}^{tmp} \leftarrow \text{zeros}(\mathbf{Y})$ ;
5   for  $ixyz \in \{1, \dots, |\Omega|\}$  do
6      $\text{index} \leftarrow \mathcal{P}_{ixyz}$ ;
7      $p \leftarrow \text{kriging}(\text{index}, \gamma, \kappa)$ ;
8      $\mathbf{Y}_{\text{index}}^{tmp} \leftarrow \text{simulate}(\text{index}, p, \tau)$ ;
9      $\mathbf{V}^{tmp} \leftarrow \text{back\_transform}(\mathbf{Y}_{\text{index}}^{tmp})$ ;
10     $\text{write}(\text{output.txt}, \mathbf{V}_{\text{index}}^{tmp})$ ;
11  end
12 end

Output:  $N$  stochastic simulations stored in file output.txt

```

The small test scenario consists in generating 96 realizations using a regular grid of $40 \times 60 \times 12$ (28,800) nodes, simulating a single variable with three variographic structures (one spherical plus two different exponentials). The large and extra large test scenarios consist in 96 and 128 realizations respectively, using a regular grid of $400 \times 600 \times 12$ (2,880,000) nodes with the same variographic features. In Fig. 8 we can see a slice in the plane XY at level $z=6$ of a simulated realization using the small and large grids. Regarding the spatial search strategy, in the three scenarios the parameter `sstrat` is set to 1, which means that the data are relocated to grid nodes and a spiral search is used, as described in Deutsch and Journel (1998) (II.4).

Re-design: Applying the first step of the methodologies in the `gamv` application, no significant acceleration is observed, only a speedup of $1.01 \times / 1.01 \times / 1.02 \times$ for the small/large/extra large scenarios is obtained.

Profiling and code optimization: Regarding the second step, the output report of `gprof` can be viewed in Table 10. We can observe that the routines `krige` (implemented as subroutine in `sgsim`, inside the pseudo-routine `kriging` from Algorithm 5), `ksol` (utility contained in `gslib.a` library, launched inside of `krige`) and `srchnd` (utility contained in `gslib.a`) together consist in approximately 89%, 90% and 91% of the total elapsed time for the small, large and extra large scenarios. `krige` assembles the covariance matrix and right-hand side of the kriging linear system, `ksol` solves it using a standard gaussian elimination algorithm without pivot search and `srchnd` search for nearby simulated grid nodes using the spiral search strategy. Although not considered in the three test scenarios, the super block search can be activated by setting `sstrat=0`. This parameter makes the hard data and previously simulated nodes been searched separately (the hard

data are searched with the super block strategy and previously simulated nodes are searched with the spiral search). In this case, the average profile in the three scenarios is dominated by the routine `srchsupr`, as described in the `kt3d` application, reaching levels of 27% of the total elapsed time, followed by `cova3`, `sqdist`, `ksol` and `krige`. The optimization of `sgsim` using the parameter `sstrat=0`, which implies the optimization of the routine `srchsupr`, is left as a future work in further versions of the code. In Table 11, the output of `ltrace` is depicted. We can see a large number of library write calls (`_gfortran_st_write`, `_gfortran_st_write_done` and `_gfortran_transfer_real_write`), executed inside the grid node loop. Additionally, there is a large number of library calls to mathematical (`expf` and `log`). The tool `strace` does not report any useful information about possible bottlenecks or overhead sources during execution in all scenarios.

A first optimization is related with code hoisting and dead code removal. A simple inspection in the code allow us to perform this simple optimization, achieving a speedup of $1.03 \times / 1.02 / 1.03 \times$. Profiles generated with `Oprofile` tool indicate that a large number of branch instructions are executed in `ksol`. For this reason we specialize this routine in order to reduce the number of branches executed. In Fig. 9 we can observe the original and specialized `ksol` routine. This optimization achieves a speedup of $1.05 \times / 1.03 \times / 1.05 \times$. In order to reduce the number of write calls, we use a buffer array to store each simulation results (\mathbf{V}^{tmp} , Algorithm 5), and at the end of each grid node loop only one write call is executed. In contrast to `kt3d` application, in this case there is no contention by false sharing when each thread stores the simulated values in the buffer, since each thread has its own local buffer array. Additionally, in order to accelerate the write call, we use a binary format in the output file, adding the option

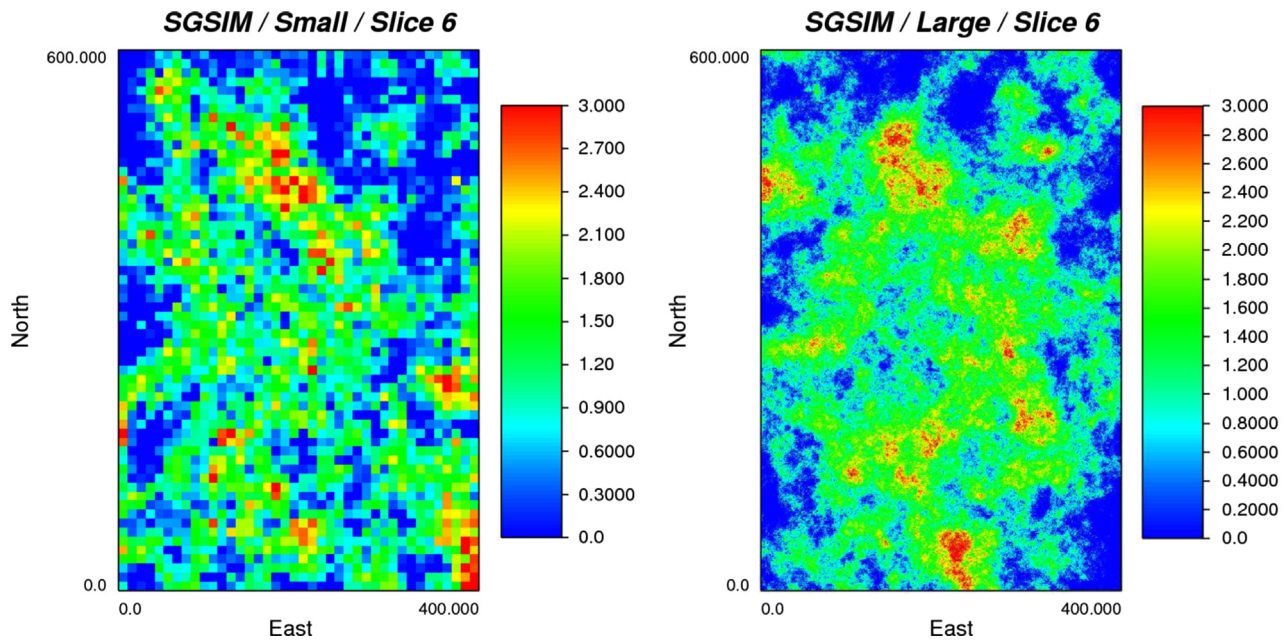


Fig. 8. Output of `sgsim` application using a small (left) and large (right) scenarios.

`form='UNFORMATTED'` in the `open` instruction. With this modification, a considerable speedup of $1.49 \times / 1.32 \times / 1.36 \times$ is achieved. Two further optimizations are added to accelerate the multi-thread execution. The first is a specialization of the GSLIB utility `sortem`, as described previously for `kt3d` application. The second is a local definition of an array, denoted `ixv`, which is stored in a common block used by the GSLIB utility routine `acorni`. This utility calculates pseudo-random numbers using a seed as an input and storing it in the first element of the array `ixv` at the beginning of the execution (Wikramaratna, 1989). In the single-thread version, these optimizations deliver a speedup of $1.53 \times / 1.32 \times / 1.38 \times$, matching the numerical results of the original `sgsim` application. A summary of all the optimizations performed can be viewed in Table 12.

OpenMP parallelization: The next step is to add OpenMP pragmas into the source code. According to Mariethoz (2010), there are three levels of parallelization in sequential stochastic simulation algorithms: realization, path and node level. We choose to parallelize at realization level because it introduces minimal modifications in the code, which is one of our main objectives. Although the node level is the finest-grained scheme, we leave this task as a future work to explore with other technologies (accelerators and co-processors), nested-parallelism with threads or task-based parallelism (Rasera et al., 2015). The `parallel` pragma is added in the loop of line 2 from Algorithm 5. The default static block scheduling is used for loop splitting, since each thread will perform the same number of iterations, $iters(N, T) = \lceil (N - T + 1) / T \rceil$, being N the total number of simulations and T the number of threads.

MPI parallelization: The distribution of load using MPI is performed in a similar way, splitting the loop of line 2 from Algorithm 5, where each MPI task will perform the same number of iterations $iters(N, P)$, with P being the total number of MPI tasks. In this way, using multiple threads per task, each thread will perform $iters(iters(N, P), T)$ iterations. As in the previous applications, the parallelization of the optimized application still delivers similar results to the original GSLIB application, allowing some tolerance up to the 5th or 6th decimal number (non-commutativity of floating-point operations), and also allowing different random seeds for each set of realizations obtained by each thread in each node.

The main steps of the multi-thread multi-node algorithm are

depicted in Algorithm 6. The random seed τ can be modified differently in each thread and node, using a function f depending on the thread and task identifier (line 6). In terms of reproducibility of results, the function f can be defined using a lookup table, where the values defined for each thread and node correspond the historical values of `ixv(1)` extracted from the utility `acorni` just before the time where the first pseudo-random number was generated in the corresponding realization of the target reproducible set of simulations. Additionally, the rest of historical values of the array `ixv`, namely `ixv(2)` to `ixv(KORDEI+1)`, must be loaded in the current `ixv` array in order to maintain the reproducibility of pseudo-random numbers. In this work, the reproducibility features were not implemented, leaving this topic for future versions of the code. After the random seed generation step, each thread opens a binary local file `output-itag.bin` where it will store the simulated back-transformed values \mathbf{V}^{tmp} (line 8). The next steps are the same as in the single-thread version, with the only difference that each thread in each node uses its own random seed τ_{itag} and the result is saved to a file after all nodes have been simulated (line 19).

Results: In Table 13 we can see the maximum achieved speedup for the small, large and extra large scenarios with 16 CPU-cores in the *Server* system, $15.51 \times$, $15.01 \times$ and $14.17 \times$ respectively. Using the *Desktop* system, the maximum achieved speedups with 4 CPU-cores are $5.66 \times$, $4.70 \times$ and $5.02 \times$ respectively. Using the optimized sequential time as baseline, the maximum achieved speedups are $10.13 \times$, $11.16 \times$ and $11.38 \times$ for the small, large and extra large scenarios in the *Server* system and $3.70 \times$, $3.55 \times$ and $3.63 \times$ respectively in the *Desktop* system. In terms of efficiency, these values represent a 63%, 69% and 71% in the *Server* system and 92%, 88% and 90% in *Desktop*, for the small, large and extra large scenarios respectively.

Multi-node results are shown in Table 14, with a maximum achieved speedup of $100.57 \times$ for the extra large scenario with 16 CPU-cores per node and 8 compute nodes. Using the optimized sequential time as baseline, the maximum achieved speedup is $80.76 \times$ using the same distributed CPU-cores. This speedup accomplishes a 63% of efficiency using 128 CPU-cores. Using 64 and 32 distributed CPU-cores and the optimized baseline, the average efficiency obtained is 72% and 80% respectively. The efficiencies

obtained using 64 CPU-cores are 65% with 4 nodes and 16 threads each, and 79% with 8 nodes and 8 threads each. Using 32 CPU-cores, the efficiencies are 76% with 2 nodes and 16 threads each, 82% with 4 nodes and 8 threads each, and 83% with 8 nodes, 4 threads each.

As in `kt3d`, the decay in efficiency observed in these results is related with overhead added by bus contention when the threads are writing to disk. No additional overhead is added since no contention is generated by false sharing, thread synchronization or MPI message sharing.

(V.3.1). Its main steps can be viewed in [Algorithm 7](#).

The main steps are analogous to the sequential gaussian simulation algorithm, implemented by the `sgsim` application. The differences with [Algorithm 5](#) are the addition of a category-based loop (line 6) and the usage of the pseudo-routine `indicator_kriging` instead of the traditional `kriging` (line 7).

Indicator kriging (IK) is a variant of the traditional simple and ordinary kriging from Eqs. (2) and (5)–(6), which uses a nonlinear transformation of the data $Z(\mathbf{u})$, based in the indicator function

Algorithm 6. Pseudo-code of `sgsim`, sequential gaussian simulation program (multi-thread multi-node algorithm).

```

Input: Same inputs as Algorithm 5;  $P$ : number of MPI tasks;  $T$ : number of execution threads;

1   $itask \leftarrow \text{MPI\_COMM\_RANK}(\dots)$ ;
2   $N_{itask} \leftarrow \lceil \frac{N-P+1}{P} \rceil$ ;
3   $\mathbf{Y} \leftarrow \text{normal\_score}(\mathbf{V})$ ;
   /* omp parallel default(firstprivate) shared( $\mathbf{Y}, \Omega$ ) */
4  foreach  $ithread \in \{1, \dots, T\}$  do
5      $itag \leftarrow ithread + itask \times T$ ;
6      $\tau_{itag} \leftarrow \tau * f(ithread, itask)$ ;
7      $N_{itag} \leftarrow \lceil \frac{N_{itask} - T + 1}{T} \rceil$ ;
8     open(output- $itag$ .bin);
9      $\mathbf{V}^{tmp} \leftarrow \text{zeros}(|\Omega|)$ ;
10    for  $isim \in \{1, \dots, N_{itag}\}$  do
11        $\mathcal{P} \leftarrow \text{create\_random\_path}(\Omega, \tau_{itag})$ ;
12        $\mathbf{Y}^{tmp} \leftarrow \text{zeros}(\mathbf{Y})$ ;
13       for  $ixyz \in \{1, \dots, |\Omega|\}$  do
14           $\text{index} \leftarrow \mathcal{P}_{ixyz}$ ;
15           $p \leftarrow \text{kriging}(\text{index}, \gamma, \kappa)$ ;
16           $\mathbf{Y}_{\text{index}}^{tmp} \leftarrow \text{simulate}(\text{index}, p, \tau_{itag})$ ;
17           $\mathbf{V}_{\text{index}}^{tmp} \leftarrow \text{back\_transform}(\mathbf{Y}_{\text{index}}^{tmp})$ ;
18       end
19       write(output- $itag$ .bin,  $\mathbf{V}^{tmp}$ );
20    end
21 end

```

Output: N stochastic simulations stored in binary files `output-1.bin, ..., output- $T \times P$.bin`

4.4. `sisim`

The application `sisim` implements the sequential indicator simulation algorithm, as described in [Deutsch and Journel \(1998\)](#)

$I(Z(\mathbf{u})) = 1$ if $Z(\mathbf{u}) \leq z$ and $I(Z(\mathbf{u})) = 0$ if $Z(\mathbf{u}) > z$. The algorithm provides a least-squares estimate on the conditional cumulative density function (ccdf) at cutoff z_k , defined as

$$[I(\mathbf{u}; z_k)]^* = \mathbb{E}\{I(\mathbf{u}; z_k | (n))\}^* = \mathbb{P}\{Z(\mathbf{u}) \leq z_k | (n)\} \quad (7)$$

where (n) is the conditional information available in the

Table 10

Output of *gprof* tool applied to *sgsim* application (small, large and extra large scenario) running in *Desktop* system. Timing column (“seconds”) may not show accurate measurements and must be taken only as an approximation reference.

% Time (<i>gprof</i> , small)	Seconds	Calls	Name
36.05	2.43	2,558,429	krige
35.61	2.40	2,558,009	ksol
17.21	1.16	2,558,592	srchnd
6.23	0.42	5,324,392	acorni
2.37	0.16	2,558,592	gauinv
1.63	0.11	1	MAIN
0.90	-	-	Other routines
% Time (<i>gprof</i> , large)	Seconds	Calls	Name
34.33	319.07	276,255,462	krige
29.24	271.73	276,258,240	srchnd
26.58	247.02	276,250,404	ksol
3.12	29.00	97	sortem
3.08	28.61	552,739,240	acorni
2.43	22.56	1	MAIN
1.22	-	-	Other routines
% Time (<i>gprof</i> , extra large)	Seconds	Calls	Name
35.36	453.53	368,340,629	krige
29.37	376.79	368,344,320	srchnd
25.83	331.36	368,333,882	ksol
3.01	38.62	129	sortem
3.01	38.56	736,985,320	acorni
2.31	29.65	1	MAIN
1.11	-	-	Other routines

Table 11

Output of *ltrace* tool applied to *sgsim* application (small, large and extra large scenario) running in *Desktop* system. Timing column (“seconds”) may not show accurate measurements and must be taken only as an approximation reference. In large and extra large scenarios, it shows the measurements running the application approximately by two minutes.

% Time (<i>ltrace</i> , small)	Seconds	Calls	Name
24.02	21.36	374,902	_gfortran_st_write_done
20.05	17.83	374,903	_gfortran_st_write
18.60	16.54	374,868	_gfortran_transfer_real_write
17.30	15.39	373,128	log
17.11	15.21	367,032	expf
2.76	2.45	28,552	_gfortran_transfer_real
0.16	-	-	Other routines
% Time (<i>ltrace</i> , large)	Seconds	Calls	Name
46.42	128.82	2,877,690	log
19.57	54.31	1,056,004	_gfortran_transfer_real_write
17.87	49.58	1,056,019	_gfortran_st_write_done
15.93	44.19	1,056,020	_gfortran_st_write
0.17	0.46	28,538	_gfortran_transfer_real
0.04	-	-	Other routines
% Time (<i>ltrace</i> , extra large)	Seconds	Calls	Name
37.48	108.51	2,877,690	log
24.32	70.41	1,333,265	_gfortran_st_write
21.06	60.97	1,333,249	_gfortran_transfer_real_write
16.92	48.97	1,333,265	_gfortran_st_write_done
0.16	0.46	28,538	_gfortran_transfer_real
0.06	-	-	Other routines

neighbourhood of location \mathbf{u} . The Simple Kriging estimator of the indicator transform $i(\mathbf{u}; z)$ can be written as

$$\begin{aligned}
 [i(\mathbf{u}; z_k)]_{SK} &= \sum_{\alpha=1}^n \lambda_{\alpha}^{(ISK)}(\mathbf{u}; z) i(\mathbf{u}_{\alpha}; z) \\
 &+ \left[1 - \sum_{\alpha=1}^n \lambda_{\alpha}^{(ISK)}(\mathbf{u}; z) \right] P\{Z(\mathbf{u}) \leq z\}
 \end{aligned}
 \tag{8}$$

where $\lambda_{\alpha}^{(ISK)}(\mathbf{u}; z)$ are the SK weights corresponding to cutoff z . These weights are calculated as the solution of the following linear system:

$$\sum_{\beta=1}^n \lambda_{\beta}^{(ISK)}(\mathbf{u}; z) C_I(\mathbf{u}_{\beta} - \mathbf{u}_{\alpha}; z) = C_I(\mathbf{u} - \mathbf{u}_{\alpha}; z), \quad \alpha = 1, \dots, n
 \tag{9}$$

with $C_I(\mathbf{h}; z) = C\{I(\mathbf{u}; z), I(\mathbf{u} + \mathbf{h}; z)\}$ being the indicator covariance at cutoff z . Ordinary Kriging estimator aims at re-estimating locally

the prior cumulative density function $P\{Z(\mathbf{u}) \leq z\}$.

time, followed by `sqdist` and `krige`, all of them with comparatively balanced weights. The effect of the super block search

Algorithm 7. Pseudo-code of `sisim`, sequential indicator simulation program (single-thread algorithm).

```

Input: ( $\mathbf{V}, \Omega$ ): sample data base values defined in a 3D domain;  $C$ : number of categories to be
reproduced;  $\gamma_1, \dots, \gamma_C$ : structural variographic models;  $\kappa$ : local interpolation parameters;  $\tau$ :
seed for pseudo-random number generator;  $N$ : number of generated simulations; output.txt:
output file

1 for  $isim \in \{1, \dots, N\}$  do
2    $\mathcal{P} \leftarrow \text{create\_random\_path}(\Omega, \tau)$ ;
3    $\mathbf{V}^{tmp} \leftarrow \text{zeros}(\mathbf{V})$ ;
4   for  $ixyz \in \{1, \dots, |\Omega|\}$  do
5      $index \leftarrow \mathcal{P}_{ixyz}$ ;
6     for  $icut \in \{1, \dots, C\}$  do
7        $p_{icut} \leftarrow \text{indicator\_kriging}(index, \gamma_{icut}, \kappa)$ ;
8     end
9      $\mathbf{V}_{index}^{tmp} \leftarrow \text{simulate}(index, p_1, \dots, p_C, \tau)$ ;
10     $\text{write}(\text{output.txt}, \mathbf{V}_{index}^{tmp})$ ;
11  end
12 end

Output:  $N$  stochastic simulations stored in file output.txt

```

The small test scenario consists in generating 96 realizations using a regular grid of $40 \times 60 \times 12$ (28,800) nodes in a 3D volume of $400 \times 600 \times 120$ (m³), simulating three categories (lithologies) with one variographic structure each (spherical). The large and extra large test scenarios consist in generating 16 and 128 realizations respectively, using a regular grid of $400 \times 600 \times 12$ (2,880,000) nodes with the same 3D volume and the same variographic features. In Fig. 10 we can see a slice in the plane XY at level $z=6$ of a simulated realization using the small and large grid. Regarding the spatial search strategy, in the three scenarios the parameter `sstrat` is set to 1, the same parameter used by the `sgsim` application.

Re-design: Applying the first step of the methodology in the `gamv` application, we can achieve a speedup of $1.05 \times / 1.03 \times / 1.01 \times$ for the small/large/extra large scenarios.

Profiling and code optimization: In the second step, the `gprof` output report can be viewed in Table 15. We can see in this report that the routines `krige` and `ksol` together consist in approximately 93%, 89% and 88% of the total elapsed time in the small, large and extra large scenarios. As in the `sgsim` application, the three test scenarios can activate the super block strategy for hard data by setting `sstrat=0`. By changing this parameter, the average profile in the three scenarios is dominated by the routines `ksol` and `cova3`, together with more than 50% of the total elapsed

routine `srchsupr` is not dominant, reaching levels between 3% and 7% in the three scenarios. In order to optimize `sisim` using the parameter `sstrat=0`, a specialization of the routine `cova3` can be applied, as shown for the `kt3d` application. However, the analysis and performance measurement of this task is left as a future work in further versions of the code.

Additionally, in Table 16 we can see a large number of library write calls (`_gfortran_st_write` and `_gfortran_transfer_real_write`) performed inside of the grid node loop.

Further information can be obtained with `Oprofile`. After exploring different profiles obtained with this tool and studying the source code, we inferred that the two main optimizations to be performed in this application are related with branch instructions reductions of the routines `krige` and `ksol`, and library call reduction using buffered and unformatted writes to disk. Inspecting the source code, we can see several loop-invariant expressions (if-then-else type) that can be hoisted out of the internal loops of both routines. If we modify first the `krige` routine, the optimized code can achieve a speedup of $1.09 \times / 1.04 \times / 1.03 \times$. After that, the routine `ksol` is specialized to use fixed arguments (in the original code those arguments never change), reducing the number of branch instructions executed, allowing us to achieve a speedup of $1.11 \times / 1.06 \times / 1.05 \times$. In order to reduce the number of library calls, we store each nodal value in a buffer array, and after all nodes have been simulated successfully, just one write call is

performed, writing the buffer array in an unformatted file previously opened with the option `form='UNFORMATTED'` (same optimization applied before in `kt3d` and `sgsim`). The resulting speedup on the *Desktop* system is $1.16 \times / 1.12 \times / 1.11 \times$ for each scenario, with numerical results matching the original `sisim` application.

A summary of the speedups obtained can be viewed in [Table 17](#).

OpenMP and MPI parallelization: Following the same strategy defined for `sgsim`, the proposed parallel implementation is depicted in [Algorithm 8](#). A single `parallelpragma` was added before the simulation loop. As for the multi-thread multi-node sequential gaussian simulation, each thread in each node initializes its own random seed τ_{itag} and writes the simulated values in a local binary file.

$14.34 \times$ for the large scenario and $15.07 \times$ for the extra large scenario using 16 CPU-cores in the *Server* system. Using the *Desktop* system, the maximum achieved speedups with 4 CPU-cores are $4.28 \times$, $4.07 \times$ and $4.04 \times$ respectively. Using the optimized sequential time as baseline, the maximum achieved speedups are $13.16 \times$, $13.17 \times$ and $13.81 \times$ for the small, large and extra large scenarios in the *Server* system and $3.66 \times$, $3.60 \times$ and $3.62 \times$ respectively in the *Desktop* system. In terms of efficiency, these values represent 82%, 82% and 86% in the *Server* system and 91%, 90% and 90% in *Desktop*, for the small, large and extra large scenarios respectively.

Multi-node results are shown in [Table 19](#), with a maximum achieved speedup of $109.40 \times$ for the extra large scenario with 16 CPU-cores per node and 8 compute nodes. Using the optimized sequential time as baseline, the maximum achieved speedup is

Algorithm 8. Pseudo-code of `sisim`, sequential indicator simulation program (multi-thread multi-node algorithm).

```

Input: Same inputs as Algorithm 7;  $P$ : number of MPI tasks;  $T$ : number of execution threads;

1   $itask \leftarrow \text{MPI\_COMM\_RANK}(\dots)$ ;
2   $N_{itask} \leftarrow \lceil \frac{N-P+1}{P} \rceil$ ;
   /* omp parallel default(firstprivate) shared(V,Ω) */
3  foreach  $ithread \in \{1, \dots, T\}$  do
4  |    $\tau_{itag} \leftarrow \tau * f(ithread, itask)$ ;
5  |    $N_{itag} \leftarrow \lceil \frac{N_{itask}-T+1}{T} \rceil$ ;
6  |   open(output- $itag$ .bin);
7  |   for  $isim \in \{1, \dots, N_{itag}\}$  do
8  | |    $\mathcal{P} \leftarrow \text{create\_random\_path}(\Omega, \tau_{itag})$ ;
9  | |    $\mathbf{V}^{tmp} \leftarrow \text{zeros}(\mathbf{V})$ ;
10 | |   for  $ixyz \in \{1, \dots, |\Omega|\}$  do
11 | | |    $\text{index} \leftarrow \mathcal{P}_{ixyz}$ ;
12 | | |   for  $icut \in \{1, \dots, C\}$  do
13 | | | |    $p_{icut} \leftarrow \text{indicator\_kriging}(\text{index}, \gamma_{icut}, \kappa)$ ;
14 | | |   end
15 | | |    $\mathbf{V}_{\text{index}}^{tmp} \leftarrow \text{simulate}(\text{index}, p_1, \dots, p_C, \tau_{itag})$ ;
16 | |   end
17 |   write(output- $itag$ .bin,  $\mathbf{V}^{tmp}$ );
18 | end
19 end

Output:  $N$  stochastic simulations stored in binary files output-1.bin, \dots, output- $T \times P$ .bin

```

Results: Analyzing the single-node results of [Table 18](#), the maximum achieved speedup is $15.13 \times$ for the small scenario,

$100.29 \times$ using the same distributed CPU-cores. This speedup accomplishes a 78% of efficiency using 128 CPU-cores. Using 64 and 32 distributed CPU-cores and the optimized baseline, the average

<pre> 1 subroutine ksol(nright,neq,nsb,a,r,s,ising) 2 ... 3 nm = nsb*neq 4 ... 5 do k=1,m1 6 ... 7 do iv=1,nright 8 nm1=nm*(iv-1) 9 ii=kk+nn*(iv-1) 10 ... 11 end do 12 end do 13 ijm=ij-nn*(nright-1) 14 ... 15 do iv=1,nright 16 nm1=nm*(iv-1) 17 ij=ijm+nn*(iv-1) 18 ... 19 end do 20 return 21 end </pre>	<pre> 1 subroutine ksol01(nright,neq,nsb,a,r,s,ising) 2 ... 3 nm = neq 4 ... 5 do k=1,m1 6 ... 7 nm1=0 8 ii=kk 9 ... 10 ... 11 end do 12 end do 13 ijm=ij 14 ... 15 nm1=0 17 ij=ijm 18 ... 19 return 21 end </pre>
---	--

Fig. 9. Original (left) and specialized (right) `ksol` utility. The original `ksol` routine accepts as arguments several right-hand side vectors ($nright > 1$ and $nsb > 1$). However, the `sgsim` application uses it only with one right-hand side vector ($nright = 1$ and $nsb = 1$). We can specialize this routine if each appearance of `nright` and `nsb` is replaced by 1.

Table 12
Summary of code optimizations in `sgsim` application running in *Desktop* system.

Optimization	Small	Large	Extra large
Baseline	1.00 ×	1.00 ×	1.00 ×
Re-design	1.01 ×	1.01 ×	1.02 ×
Code hoisting + remove dead code	1.03 ×	1.02 ×	1.03 ×
Specialization of <code>ksol</code>	1.05 ×	1.03 ×	1.05 ×
Buffered binary writing	1.49 ×	1.32 ×	1.36 ×
Specialization of <code>sortem</code>	1.51 ×	1.33 ×	1.39 ×
Local arrays of <code>acorni</code>	1.53 ×	1.32 ×	1.38 ×

Table 13
Single-node time/speedup results for `sgsim`. Small scenario: $40 \times 60 \times 12$ grid nodes, 96 realizations. Large scenario: $400 \times 600 \times 12$ grid nodes, 96 realizations. Extra large scenario: $400 \times 600 \times 12$ grid nodes, 128 realizations.

#Threads	Time (s) (Speedup)					
	Server Small 1 Node	Server Large 1 Node	Server Extra large 1 Node	Desktop Small 1 Node	Desktop Large 1 Node	Desktop Extra large 1 Node
1 (base gslib)	10.24 (1.0 ×)	1177.52 (1.0 ×)	1528.76 (1.0 ×)	10.25 (1.0 ×)	1268.98 (1.0 ×)	1807.52 (1.0 ×)
1 (optimized)	6.69 (1.53 ×)	875.17 (1.34 ×)	1227.67 (1.24 ×)	6.70 (1.53 ×)	959.73 (1.32 ×)	1307.87 (1.38 ×)
2 (optimized)	3.49 (2.93 ×)	456.48 (2.57 ×)	640.38 (2.38 ×)	3.51 (2.92 ×)	511.9 (2.47 ×)	685.41 (2.63 ×)
4 (optimized)	2.08 (4.92 ×)	245.01 (4.80 ×)	351.47 (4.34 ×)	1.81 (5.66 ×)	269.95 (4.70 ×)	359.78 (5.02 ×)
8 (optimized)	0.92 (11.13 ×)	134.5 (8.75 ×)	178.18 (8.57 ×)	–	–	–
16 (optimized)	0.66 (15.51 ×)	78.42 (15.01 ×)	107.82 (14.17 ×)	–	–	–

efficiency obtained is 85% and 89% respectively. The efficiencies obtained using 64 CPU-cores are 80% with 4 nodes and 16 threads each, and 89% with 8 nodes and 8 threads each. Using 32 CPU-cores, the efficiencies are 83% with 2 nodes and 16 threads each, 91% with 4 nodes and 8 threads each, and 95% with 8 nodes, 4 threads each.

Although this application is similar to `sgsim`, the main difference is the amount of computations performed. In `sgsim`, a small kriging system is solved for each node in the grid. In `sisim`, `C` kriging systems are solved for each node in the grid (in these

examples $C=3$). For this reason, the effect of contention due to disk writing has less impact in the parallel efficiency obtained.

5. Conclusions and future work

We have shown a methodology to accelerate GSLIB applications and utilities based on code optimizations and hybrid parallel programming using multi-core and multi-node execution with OpenMP directives and MPI task distribution. The methodology

Table 14
Multi-node time/speedup results for *sgsim*.

#Threads	Time (s) (Speedup)		
	Server Extra large 2 Nodes	Server Extra large 4 Nodes	Server Extra large 8 Nodes
1 (optimized)	611.77 (2.49 ×)	309.9 (4.93 ×)	156.54 (9.76 ×)
2 (optimized)	309.10 (4.94 ×)	161.48 (9.46 ×)	83.05 (18.40 ×)
4 (optimized)	170.65 (8.95 ×)	88.55 (17.26 ×)	45.87 (33.32 ×)
8 (optimized)	90.73 (16.84 ×)	46.57 (32.82 ×)	24.15 (63.30 ×)
16 (optimized)	50.09 (30.52 ×)	29.24 (52.28 ×)	15.20 (100.57 ×)

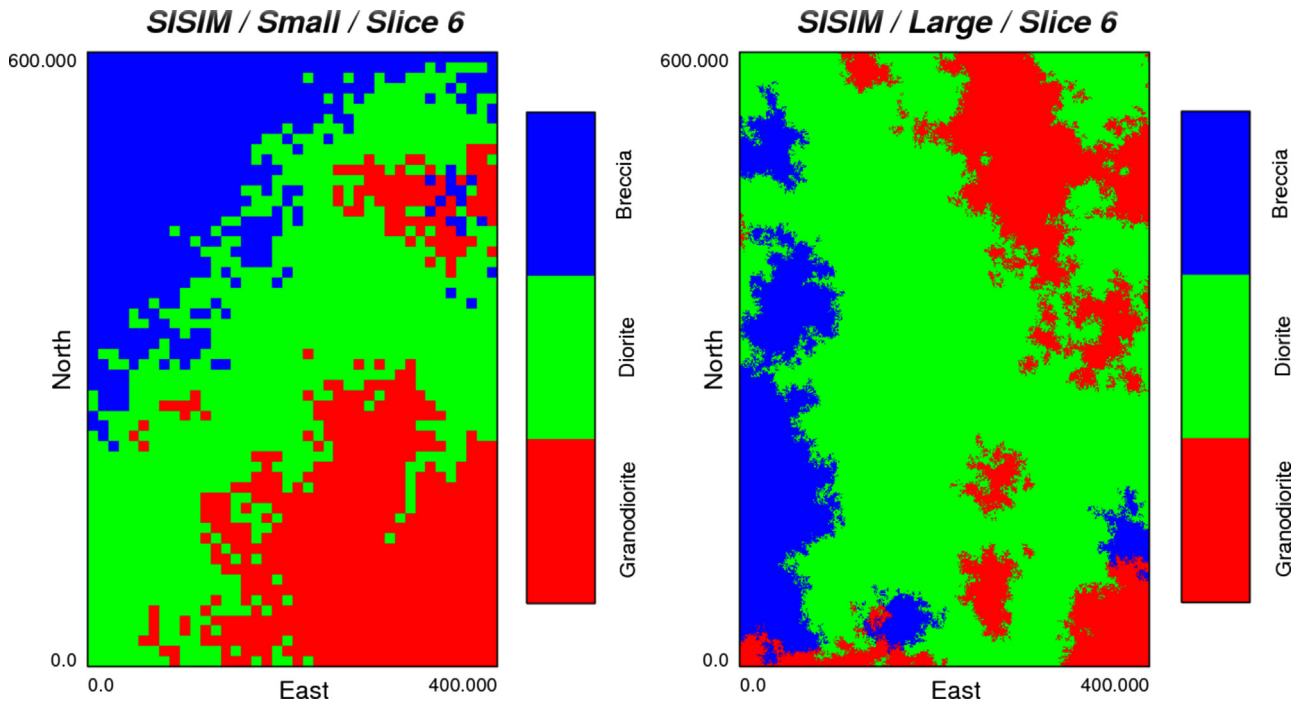


Fig. 10. Output of *sgsim* application using a small (left) and large (right) scenarios.

Table 15
Output of *gprof* tool applied to *sgsim* application (small, large and extra large scenarios) running in *Desktop* system. Timing column (“seconds”) may not show accurate measurements and must be taken only as an approximation reference.

% Time (<i>gprof</i> , small)	Seconds	Calls	Name
56.96	38.05	7,673,982	ksol
36.47	24.36	7,675,245	krige
4.25	2.84	2,558,592	srchnd
1.26	0.84	1	MAIN
1.06	–	–	Other routines
% Time (<i>gprof</i> , large)	Seconds	Calls	Name
55.24	711.83	138,125,112	ksol
33.73	434.68	138,127,767	krige
8.73	112.47	46,043,040	srchnd
1.20	15.50	1	MAIN
1.10	–	–	Other routines
% Time (<i>gprof</i> , extra large)	Seconds	Calls	Name
51.69	4896.94	1,105,001,571	ksol
36.64	3471.48	1,105,022,088	krige
9.12	863.74	368,344,320	srchnd
1.33	126.40	1	MAIN
1.211	–	–	Other routines

Table 16

Output of `ltrace` tool applied to `sisim` application (small, large and extra large scenarios) running in `Desktop` system. Timing column (“seconds”) may not show accurate measurements and must be taken only as an approximation reference. In this case it shows the measurements running the application approximately by two minutes.

% Time (<code>ltrace</code> , small)	Seconds	Calls	Name
36.97	114.90	1,931,004	<code>_gfortran_transfer_real_write</code>
32.15	99.90	1,931,126	<code>_gfortran_st_write</code>
30.02	93.28	1,931,125	<code>_gfortran_st_write_done</code>
0.86	–	–	Other routines
% Time (<code>ltrace</code> , large)	Seconds	Calls	Name
42.98	102.56	2,017,680	<code>_gfortran_transfer_real_write</code>
29.60	70.64	2,017,669	<code>_gfortran_st_write</code>
26.28	62.71	2,017,668	<code>_gfortran_st_write_done</code>
1.14	–	–	Other routines
% Time (<code>ltrace</code> , extra large)	Seconds	Calls	Name
35.42	96.16	2,024,550	<code>_gfortran_st_write</code>
33.26	90.31	2,024,562	<code>_gfortran_transfer_real_write</code>
31.07	84.36	2,024,550	<code>_gfortran_st_write_done</code>
0.25	–	–	Other routines

Table 17

Summary of code optimizations in `sisim` application running in `Desktop` system.

Optimization	Small	Large	Extra large
Baseline	1.00 ×	1.00 ×	1.00 ×
Re-design	1.05 ×	1.03 ×	1.01 ×
Code hoisting + remove dead code	1.09 ×	1.04 ×	1.03 ×
Specialization of <code>ksol</code>	1.11 ×	1.06 ×	1.05 ×
Buffered binary writing	1.16 ×	1.12 ×	1.11 ×

was tested in four well-known GSLIB applications: `gamv`, `kt3d`, `sgsim` and `sisim`. All tests were performed in Linux-based systems. However, no additional external libraries or intrinsic operating system routines were used, so the code could be compiled and tested in other distributed systems (Windows and Mac) without further modifications.

Hardware independent code optimizations are applied to the four applications, in order to maintain the basic portability of the original code. The most critical optimizations are related to reductions of branch instruction execution, avoidance of miss-predicted branch instructions, specialization of generic routines and buffering I/O calls. OpenMP parallel loops and regions, as well as MPI instructions, are added to the optimized codes, with minimal modifications, delivering reasonable speedup results compared to

the performance of original GSLIB applications.

The variogram application `gamv` reaches a speedup of 69.56 ×, 60.99 × and 37.46 × with 16 non-distributed CPU-cores in the small, large and extra large test scenarios, and 232.07 × with 128 distributed CPU-cores in the extra large scenario. This high speedup value was obtained after applying a small code optimization which reduces the number of miss-predicted branches during execution. Additionally, a static optimal chunk-size strategy was set for the run-time scheduling of parallelization. The kriging application `kt3d` reaches a speedup of 13.87 ×, 14.48 × and 13.88 × with 16 non-distributed CPU-cores in the small, large and extra large test scenarios, and 56.78 × with 128 distributed CPU-cores in the extra large scenario. Code optimizations were difficult to apply in this application, and some features of the single-thread

Table 18

Time/speedup results for `sisim`. Small scenario: 40 × 60 × 12 grid nodes, 96 realizations. Large scenario: 400 × 600 × 12 grid nodes, 16 realizations. Extra large scenario: 400 × 600 × 12 grid nodes, 128 realizations.

#Threads	Time (s) (Speedup)					
	Server Small 1 Node	Server Large 1 Node	Server Extra large 1 Node	Desktop Small 1 Node	Desktop Large 1 Node	Desktop Extra large 1 Node
1 (base gslib)	74.46 (1.0 ×)	1378.92 (1.0 ×)	11060 (1.0 ×)	73.68 (1.0 ×)	1369.97 (1.0 ×)	10995.0 (1.0 ×)
1 (optimized)	64.77 (1.14 ×)	1266.31 (1.08 ×)	10130 (1.09 ×)	63.14 (1.16 ×)	1213.48 (1.12 ×)	9854.0 (1.11 ×)
2 (optimized)	33.93 (2.19 ×)	641.79 (2.14 ×)	5100 (2.16 ×)	33.16 (2.22 ×)	649.46 (2.11 ×)	5235.0 (2.10 ×)
4 (optimized)	16.96 (4.39 ×)	327.95 (4.20 ×)	2597 (4.25 ×)	17.21 (4.28 ×)	336.76 (4.07 ×)	2716.72 (4.04 ×)
8 (optimized)	8.89 (8.37 ×)	173.03 (7.96 ×)	1358 (8.14 ×)	–	–	–
16 (optimized)	4.92 (15.13 ×)	96.15 (14.34 ×)	733 (15.07 ×)	–	–	–

Table 19
Multi-node time/speedup results for *sisim*.

#Threads	Time (s) (Speedup)		
	Server Extra large 2 Nodes	Server Extra large 4 Nodes	Server Extra large 8 Nodes
1 (optimized)	5242 (2.10 ×)	2560 (4.31 ×)	1273 (8.68 ×)
2 (optimized)	2556 (4.32 ×)	1288 (8.58 ×)	640 (17.27 ×)
4 (optimized)	1308 (8.45 ×)	653 (16.91 ×)	332 (33.29 ×)
8 (optimized)	697 (15.86 ×)	345 (31.98 ×)	176 (62.80 ×)
16 (optimized)	377 (29.33 ×)	196 (56.42 ×)	101 (109.40 ×)

version have to be eliminated in the multi-thread version, in order to accelerate the parallel execution. The default schedule was the best alternative for the run-time scheduling of parallelization. Both sequential simulation applications, *sgsim* and *sisim*, reach speedup levels of 15.51 × and 15.13 × in the small scenario, 15.01 × and 14.34 × in the large scenario, and 14.17 × and 15.07 × in the extra large scenario using a single node of execution. In the extra large scenario using 128 distributed CPU-cores, both applications reach a speedup of 100.57 × and 109.40 × respectively. The code optimizations that showed successful results were the specialization of routines and buffering I/O calls. As shown in Algorithms 6 and 8, a parallel region was defined for both applications without using a parallel loop to distribute the workload, reducing the thread synchronization overhead.

Using the current multi-thread multi-node versions of the applications, we can conclude that parallel *sisim* is the best candidate to scale and keep the efficiency of the parallel execution using a large number of CPU-cores. Results showed that using a small and large dataset/grid, *sisim* delivers a similar efficiency (82% both) using 16 CPU-cores in a single node. Using the extra large dataset/grid, it improves its efficiency (86%) using 16 CPU-cores in a single node, without degrading it considerably using up to 128 distributed CPU-cores (89%, 85% and 78% with 32, 64 and 128 CPU-cores). *kt3d* showed an increment in efficiency from the small to large scenarios (77–85%) using 16 CPU-cores in a single node. However, with the extra large scenario a considerable decay is observed (80%). This behavior is accentuated in the multi-node tests, reaching even lower levels of average efficiency up to 128 distributed CPU-cores (71%, 56% and 41% with 32, 64 and 128 CPU-cores). *gamv* showed an increment in efficiency proportional to the size of dataset, from small to extra large scenarios (68%, 85% and 86% respectively) using 16 CPU-cores in a single node. In the multi-node tests, the average efficiency decreases steadily using up to 128 distributed CPU-cores (86%, 79% and 66% with 32, 64 and 128 CPU-cores), without reaching lower levels as in the *kt3d* application. *sgsim* showed the lowest efficiency percentages in the small and large scenarios (63% and 69% respectively) using 16 CPU-cores in a single node. However, in the extra large scenario the efficiency improved using 16 CPU-cores in a single node (71%) and slightly decreased using up to 128 distributed CPU-cores (80%, 72% and 63% with 32, 64 and 128 CPU-cores).

According to the single-node results, the efficiency increases as the size of the input dataset and/or estimation/simulation grid gets larger, which is in concordance with Gustafson's (1988) law, a counterpart of the well-known Amdahl's (1967) law. Gustafson's law says that computations involving arbitrarily large data sets can be efficiently parallelized. If larger scenarios are considered in future versions of the applications, we can expect higher efficiency values. Additionally, extremely large scenarios may benefit from memory-aware optimizations, such as the improved memory access pattern proposed for the super block strategy *srchsupr* in

the context of *kt3d* application.

Our final goal is to apply this methodology to the most important applications and utilities from GSLIB, setting the compilation scripts as general as possible, so other scientists and practitioners could benefit from this effort. The addition of new acceleration technologies, like Graphics Processing Units (GPU) (Kirk and Hwu, 2010) and Many Integrated Cores (MIC) (Jeffers and Reinders, 2013), and advanced parallelization techniques, like task-based OpenMP pragmas (Ayguaé et al., 2009), are being explored using the present work as base.

6. Source code

The current version of the modified codes can be downloaded from <https://github.com/operedo/gslib-alges>.

Acknowledgements

The authors thankfully acknowledge the computer resources, technical expertise and assistance provided by the Barcelona Supercomputing Center – Centro Nacional de Supercomputación (Spain) which supports the Marenostrum supercomputer, and the National Laboratory for High Performance Computing (Chile), which supports the Leftraru supercomputer. Additional thanks are owed to industrial supporters of ALGES laboratory, in particular Yamana Gold, as well as the Advanced Mining Technology Center (AMTC) and the whole ALGES team.

References

- Aho, A.V., Lam, M.S., Sethi, R., Ullman, J.D., 2006. *Compilers: Principles, Techniques, and Tools*, 2nd edition. Addison Wesley.
- Amdahl, G.M., 1967. Validity of the single processor approach to achieving large scale computing capabilities. In: Proceedings of the April 18–20, 1967, Spring Joint Computer Conference, AFIPS '67 (Spring), ACM, New York, NY, USA, 1967, pp. 483–485.
- Ayguaé, E., Copty, N., Duran, A., Hoeflinger, J., Lin, Y., Massaioli, F., Teruel, X., Unnikrishnan, P., Zhang, G., 2009. The Design of OpenMP Tasks. *IEEE Trans. Parallel Distrib. Syst.* 20 (3), 404–418.
- Barcelona Supercomputing Center, 2015. Paraver/extrae performance analysis tools. Computer Science Department. Website (<http://www.bsc.es/computer-sciences/performance-tools>) (accessed 04.08.2015).
- Chandra, R., Dagum, L., Kohr, D., Maydan, D., McDonald, J., Menon, R., 2001. *Parallel Programming in OpenMP*. Morgan Kaufmann Pub. Inc., San Francisco, CA, USA.
- Culler, D., Singh, J., Gupta, A., 1998. Parallel computer architecture: a hardware/software approach. In: The Morgan Kaufmann Series in Computer Architecture and Design, 1st edition. Morgan Kaufmann.
- Deutsch, C.V., Journel, A.G., 1998. *GSLIB: Geostatistical Software Library and User's Guide*. Oxford University Press.
- Graham, S.L., Kessler, P.B., Mckusick, M.K., 1982. *gprof: A Call Graph Execution Profiler*. SIGPLAN Not 17 (6), 120–126.
- Gustafson, J.L., 1988. Reevaluating Amdahl's law. *Commun. ACM* 31, 532–533.
- High performance geostatistics library, 2010. version 0.9.9, (<http://hppl.mit-ufa.com>).
- Intel Corporation, September 2014. Intel 64 and IA-32 architectures software developer's manual. In: System Programming Guide, vol. 3 (3A, 3B & 3C).
- Jeffers, J., Reinders, J., 2013. *Intel Xeon. Phi Coprocessor High Performance Programming*, 1st Edition Morgan Kaufmann Publishers Inc, San Francisco, CA, USA.
- Johnson, M.K., Troan, E.W., 2004. *Linux Application Development*, 2nd edition. Addison-Wesley Professional.
- Kirk, D.B., Hwu, W.-m.W., 2010. *Programming Massively Parallel Processors: A Hands-on Approach*, 1st Edition Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Mariethoz, G., 2010. A general parallelization strategy for random path based geostatistical simulation methods. *Comput. Geosci.* 36 (7), 953–958.
- Navarro, C.A., Hirschfeld, N., Improving the GPU space of computation under triangular domain problems. *Comput. Res. Reposit. CoRR* abs/1308.1419.
- Oprofile community, 2015. Oprofile, system-wide profiler for Linux systems. Website. URL (<http://oprofile.sourceforge.net/about/>) (accessed 04.08.2015).
- Peredo, O., Ortiz, J.M., Herrero, J.R., Samaniego, C., 2014. Tuning and hybrid parallelization of a genetic-based multi-point statistics simulation code. *Parallel*

- Comput. 40 (5–6), 144–158.
- Rasera, L.G., Lopes Machado, P., Costa, J.F.C.L., 2015. A conflict-free, path-level parallelization approach for sequential simulation algorithms. *Comput. Geosci.* 80, 49–61.
- Remy, N., Boucher, A., Wu, J., 2009. *Applied Geostatistics with SGeMS: A User's Guide*. Cambridge University Press.
- Snir, M., Otto, S., Huss-Lederman, S., Walker, D., Dongarra, J., 1998. MPI-the complete reference. In: *The MPI Core*, vol. 1, 2nd edition. MIT Press, Cambridge, MA, USA.
- Staios LLC, 2001. WinGslib installation and getting started guide. Website (<http://www.staios.com/WinGslib/GettingStarted.pdf>) (accessed 04.08.2015).
- Straubhaar, J., Walgenwitz, A., Renard, P., 2013. Parallel multiple-point statistics algorithm based on list and tree structures. *Math. Geosci.* 45 (2), 131–147.
- Wikramaratna, R.S., 1989. ACORN: a new method for generating sequences of uniformly distributed pseudo-random numbers. *J. Comput. Phys.* 83 (1), 16–31.
- ZIH - TU Dresden, 2015. VampirTrace 5.14.4 user manual. Website URL (http://tu-dresden.de/die_tu_dresden/zentrale_einrichtungen/zih/forschung/projekte/vampirtrace) (accessed 04.08.2015).