

Tuning and hybrid parallelization of a genetic-based multi-point statistics simulation code



Oscar Peredo ^{a,b,*}, Julián M. Ortiz ^{b,c}, José R. Herrero ^d, Cristóbal Samaniego ^a

^a Barcelona Supercomputing Center (BSC-CNS), Department of Computer Applications in Science and Engineering, Edificio NEXUS I, Campus Nord UPC, Gran Capitán 2-4, 08034 Barcelona, Catalunya, Spain

^b Advanced Laboratory for Geostatistical Supercomputing, Advanced Mining Technology Center, University of Chile, Chile

^c Department of Mining Engineering, University of Chile, Av. Tupper 2069, Santiago 837-0451, 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 15 February 2013

Received in revised form 19 January 2014

Accepted 18 April 2014

Available online 29 April 2014

Keywords:

Geostatistics

Stochastic simulation

Multi-point statistics

Code optimization

Parallel computing

Genetic algorithms

ABSTRACT

One of the main difficulties using multi-point statistical (MPS) simulation based on annealing techniques or genetic algorithms concerns the excessive amount of time and memory that must be spent in order to achieve convergence. In this work we propose code optimizations and parallelization schemes over a genetic-based MPS code with the aim of speeding up the execution time. The code optimizations involve the reduction of cache misses in the array accesses, avoid branching instructions and increase the locality of the accessed data. The hybrid parallelization scheme involves a fine-grain parallelization of loops using a shared-memory programming model (OpenMP) and a coarse-grain distribution of load among several computational nodes using a distributed-memory programming model (MPI). Convergence, execution time and speed-up results are presented using 2D training images of sizes $100 \times 100 \times 1$ and $1000 \times 1000 \times 1$ on a distributed-shared memory supercomputing facility.

© 2014 Elsevier B.V. All rights reserved.

1. Multi-point statistics simulation

Numerical modeling with geostatistical techniques aims at characterizing natural phenomena by summarizing and using the spatial correlation of collected data in order to measure the uncertainty at unsampled locations in space. As explained by Deutsch [15], in simulation techniques, this spatial correlation is imposed into a model commonly constructed on a regular lattice. The models must reproduce the statistical (histogram) and spatial distribution (variogram or other spatial statistics) and their quality is often judged in terms of the reproduction of geological features.

Conventional techniques in geostatistics address the modeling using statistical measures of spatial correlation that quantify the expected dissimilarity (transition to a different category) between locations separated by a given vector distance, in reference to a given attribute, such as the facies, rock type, porosity, grade of an element of interest, among others. This is done using the variogram. Limitations of these techniques have been pointed out in that they only account for two locations at a time when defining the spatial structure [27]. Much richer features can be captured by using multi-point statistics (MPS)

* Corresponding author at: Barcelona Supercomputing Center (BSC-CNS), Department of Computer Applications in Science and Engineering, Edificio NEXUS I, Campus Nord UPC, Gran Capitán 2-4, 08034 Barcelona, Catalunya, Spain. Tel.: +34 934016817.

E-mail address: operedo@alges.cl (O. Peredo).




Fig. 1. Training image (left) and simulated realization (right).

that consider the simultaneous arrangement of the attribute of interest at several locations, providing the possibility to account for complex features, such as hierarchy between facies, delay effects, superposition or curvilinearity.

MPS simulation aims at generating realizations that reproduce pattern statistics inferred from some training source, usually a training image. For example, in Fig. 1, left, we can see a training image based on sinuous channels with a simulated realization. These training images are used as a pattern database to generate simulations of the underlying image, as shown in Fig. 1, right. The simulations use those patterns with the aim that the training and simulated images share the same pattern histogram.

There are several approaches to simulate accounting for MPS. Modifications of conventional methods to impose local directions of continuity using the variogram is a simple approach to impose some of the complex geological features [43,44]. Object based methods and methods inspired in the genetic rules and physics of the deposition of sediments in different environments also seek to overcome the limitations of conventional categorical simulation techniques, with significant progress [13,42,36]. Presently, the most popular method is a sequential approach based on Bayes' postulate to infer the conditional distribution from the frequencies of multi-point arrangements obtained from a training image. This method, originally proposed by Guardiano and Srivastava [21], and later efficiently implemented by Strebelle and Journel [41], is called single normal equation simulation (snesim) (see also [40]). This method has been the foundation for many variants such as simulating directly full patterns [1,17] and using filters to approximate the patterns [45]. The use of a Gibbs Sampling algorithm to account directly for patterns has also been proposed [2,29]. A sequential method using a fixed search pattern and a 'unilateral path' also provides good results [8,9,33]. Other approaches available consider the use of neural networks [3,4], updating conditional distributions with multi-point statistics as auxiliary information [30–32] or secondary variable [24]. Recently, a couple of new approaches focused on patching patterns directly to reduce computing time and impose larger scale structures, have been presented [37,18]. These methods have a significant potential for practical applications. Alternatively, the problem can be addressed as an optimization one, using simulated annealing [14] or genetic algorithms [35]. The genetic approach is still under development, but essentially follows the same stochastic strategy as the annealing scheme. This work focuses on code optimizations and parallelization of a genetic-based sequential code that simulates categorical variables to reproduce multi-point statistics. However, many of the techniques and ideas proposed here can be applied to other codes implementing similar simulation algorithms.

In Section 2 we explain the basic ideas about genetic algorithms, parallel architectures and programming models. After that, the main bottlenecks of the genetic-based simulation are detailed in Section 3. A brief explanation of the actual implementation is presented in Section 4, together with the proposed code optimizations and parallelization schemes, in Sections 5 and 6 respectively. Finally, in the last sections we include the results obtained and final conclusions.

2. Genetic algorithms and parallel computing

Genetic algorithms (GA) were developed in the 1970s with the work of Holland [23] and in subsequent decades with De Jong and Goldberg [12,19]. Initially used to find good feasible solutions for combinatorial optimization problems, today they are used in various industrial applications, and recent advances in parallel computing have allowed their development and continuing expansion.

In the canonical approach of GA, typically there is an initial population of individuals, where each individual is represented by a string of bits, as $indiv_k = 000110101$, and a fitness function $fitness(indiv_k)$ which represents the performance of each individual. The fitness function, or objective function, is the objective that must be minimized through the generations over all the individuals. A termination criteria must be defined in order to achieve the desired level of decrement in the

fitness function. The main steps and operations performed in a canonical GA can be viewed in Algorithm 1. *Selection* and *restart* are operations performed over the entire population. *Selection* extracts the best individuals and *restart* modifies part or the entire population in order to jump from local optimal values. *Crossover* and *mutation* are operations performed over particular individuals of the population. *Crossover* mixes the bits of two individuals according to a predefined set of cut points and *mutation* modifies specific random bits from one individual. Other operators can be found in the mentioned literature.

Algorithm 1. Canonical genetic algorithm

```

1: INPUT:  $N$  individuals (population)
2: Evaluate a fitness function  $fitness$  in each individual
3: while termination criteria is not achieved do
4:   {Breeding a new generation}
5:   Sort the individuals by their fitness function value
6:   if no improvement is measured in the population then
7:     Restart: select some individuals and restart their bits
8:     Sort the individuals by their fitness function value
9:   end if
10:  Selection: select the best individuals based on their fitness function
11:  Crossover: breed new individuals crossing bits of individuals from the selection
12:  Mutation: breed new individuals mutating some bits of individuals from the selection
13:  Replace old individuals with new ones
14:  Evaluate a fitness function  $fitness$  in each individual
15: end while
16: OUTPUT: best individual in the population

```

In parallel computing architectures, as described by Culler et al. [7], the two main models are distributed-memory and shared-memory, with their respective best known programming models MPI [38] and OpenMP [6]. In the first model, each processor has its own private memory and the data interchanged between processors travels through a network in chunks of messages. The speed of this communication depends on the speed of the interconnection network. In the second model, each processor has access to a common memory through data coherence and data consistency methods.

In order to use efficiently all the resources of parallel architectures, we need to explore algorithms that can exploit the parallelism and be able to adapt to future trends.

Genetic algorithms receive the classification of *embarrassingly parallel* technique to solve problems. This classification comes from the fact that separating the workload of the problem into several parallel tasks is trivial. This property motivates its investigation and application in the field of geostatistics, and particularly in MPS simulation.

3. Bottlenecks of genetic-based MPS simulation

We can see in Algorithm 1 that the evaluation of the function $fitness$ is performed $\#generations \times \#individuals$, and strong evidence indicates that this function is the most time consuming routine (a profiling using the `gprof` tool, Graham et al. [20], tells us that for sufficiently large training images, more than 96% of the execution time is spent in this routine). Its calculation is based on an object called *template*. A *template* T consists in a set of coordinates that defines cell positions. Another interpretation is that a template is basically a pattern of memory accesses. $Patterns(T)$ represents all the possible patterns that can be generated from a template T given k possible categories. If a template is defined as $T = \{(1,1), (2,1), (1,2), (2,2)\}$ (4 nodes) and the number of categories is $k = 2$, the pattern database $Patterns(T)$ will have 2^4 elements. Complex geometries can be used to define the template and its corresponding pattern database, for example, in Fig. 2 we have a template defined as $T = \{(1,1), (5,1), (9,1), (4,4), (5,4), (6,4), (1,5), (4,5), (5,5), (6,5), (9,5), (4,6), (5,6), (6,6), (1,9), (5,9), (9,9)\}$. This template is disconnected and its memory-access pattern is very irregular. This irregularity induces a slowdown in the overall performance when we need to traverse all its nodes. In contrast, the template of Fig. 3 has a regular memory-access pattern, if we access to the first element $T(1,1)$, additionally the CPU will bring to the cache memory contiguous elements for free (each CPU has its own cache line size, for Example 64 bytes, which means that each cache line can store 16 consecutive integers of 4 bytes). In Fortran, the CPU will bring a column line, in C/C++, a row line (see [22] for more details about the CPU memory hierarchy).

Given the pattern database $Pattern(T)$, two main tasks must be performed:

- First we have to count the frequency of appearances of each pattern in the training image and store them in an appropriate structure.

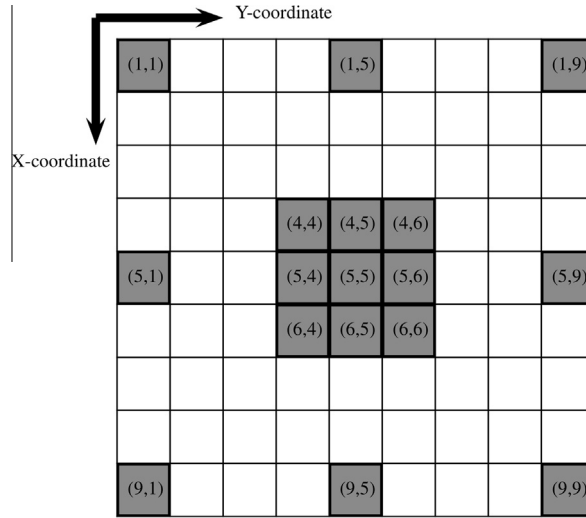


Fig. 2. Template of 17 nodes with a complex geometry (irregular memory accesses).




Fig. 3. Template of 18 nodes with a simple geometry (regular memory accesses).

- After that, for each individual in the population, and in each generation, we have to count the frequency of appearances of those patterns and calculate the following equation (other possible equations can be viewed in [34]):

$$fitness(indiv_k) = \sum_{p \in Patterns(T)} O_p \left(freq_T(p) - freq_{indiv_k}(p) \right)^2 \tag{1}$$

with O_p a weight factor for each pattern, $freq_T(p)$ and $freq_{indiv_k}(p)$ the number of appearances of pattern p in a training image and individual $indiv_k$ respectively.

In both tasks, we have to handle with patterns located at boundary nodes. A buffer band of halo nodes, with bandwidth equal to $h = \max\{width(T), height(T)\}$, is attached at the boundaries of the training image and realizations. For the training image, an extension of the original image is added keeping the geological continuity and statistical properties (alternatively, for sufficiently large images, a reduction of size h in each side of the images can be applied, keeping the removed space as buffer zone).

For the first task, following ideas from [39], we store the frequencies of the patterns that appear in the training image in a list \mathcal{L} . In this list each element is a pair (d, f) where $d = (s_1, \dots, s_{|T|})$ is the pattern (data event), stored as an array of integers of length $|T|$ with $s_i \in \{0, \dots, k - 1\}$ (k categories), and f is the frequency of appearance, stored as an integer. With this structures defined, the algorithm used in this task corresponds to Algorithm 2 using an empty frequency list \mathcal{L} .

For the second task, the accounting process is performed using Algorithm 2.

Algorithm 2. Fitness function calculation: $fitness(indiv_k)$

```

1: INPUT: individual  $indiv_k$ , training image frequency list  $\mathcal{L}$ , template  $T$ 
2:  $sum \leftarrow 0$ 
3:  $\mathcal{L}_{aux} \leftarrow \mathcal{L}$ 
4: for each node  $(i,j)$  from  $indiv_k$  do
5:   Extract pattern located in node  $(i,j)$  using template  $T$  and store it in an array  $localPattern$  of length  $|T|$ 
6:   Search  $localPattern$  in the list  $\mathcal{L}_{aux}$ 
7:   if  $localPattern$  exists in  $\mathcal{L}_{aux}$  then
8:      $\mathcal{L}_{aux}[localPattern] \leftarrow \mathcal{L}_{aux}[localPattern] - 1$ 
9:   end if
10: end for
11: for each pattern  $p$  in  $\mathcal{L}_{aux}$  do
12:    $sum \leftarrow sum + \mathcal{L}_{aux}[p] * \mathcal{L}_{aux}[p]$ 
13: end for
14: OUTPUT:  $sum$ 

```

In this algorithm, we compare the histograms of the individual and the training image. A considerable bottle-neck for this calculation is the access to the list \mathcal{L} which stores the training image histogram, because for each extracted pattern from an individual, a search must be performed over it in order to see if this pattern exists in the histogram of the training image or not. A proposed solution to this problem is to store the elements (d, f) of the list using a lexicographical order in the patterns d . This order allows to search the existence of a pattern in the list using a binary search with an average and worst case performance of order $O(\log_2 n)$ comparisons, with n the length of \mathcal{L} .

In the next section we will explain the implementation issues for this two tasks using code examples in Fortran 90 as programming language, in order to see the data structures that are used and the proposed optimizations.

4. Implementation

4.1. Storage of pattern frequencies from training image

The global variables and routines involved in the storage and management of the list \mathcal{L} are encapsulated in a Fortran module called `patternOperations`. The global variables are an integer `npatterns` and an array `patternList (:)`, which store the size and elements of the list \mathcal{L} respectively. Each entry of `patternList` is a structure which identifies a pattern. This structure is composed by an integer array `pattern (:)` (of length $|T|$ represented by the value `tem_nodes`) which stores the values of the pattern and an integer value `frequency` which stores its frequency of occurrence. The routines implemented in this module are `patternInsertion` which inserts a pattern into \mathcal{L} , `patternSearch` which returns the position of the queried pattern into \mathcal{L} (or -1 if it is not found), and `patternComparison` which compares two patterns.

Initially, when we store the pattern histogram of the training image, we use the global array `patternList` to keep track of the different patterns `patternList (i)%pattern (1:tem_nodes)` with their respective frequencies `patternList (i)%frequency`. A first scan to the training image must be performed in order to fill `patternList`. Using the routine `patternInsertion`, each time a new pattern is found, the memory space used by `patternList` is re-allocated (adding one new element with frequency equal to 1) and all the elements previously inserted together with the new pattern are re-ordered according to a lexicographical order. If an existing pattern is found, the frequencies are updated. The lexicographical order is as follows: given two arrays of integers of the same length $A = (a_1, \dots, a_n)$ and $B = (b_1, \dots, b_n)$ we will say that A is greater than B if and only if $\exists k \in \{1, \dots, n\}$ such that $a_k > b_k$ and $\forall i$ satisfying $i < k, a_i = b_i$ holds. For example, given this arrays $A = (0, 0, 0, 1, 0, 0)$ and $B = (0, 0, 0, 0, 0, 0)$, this order will indicate that A is greater than B .

After the filling step, `patternList` will be scanned each time we call the routine `patternSearch`. This routine gives the index position in the global array `patternList` where the searched pattern is located, if there is a match. If there is no match, it returns -1 . This routine is implemented through a simple iterative binary search, which uses the routine `patternComparison` to compare patterns.

The routine `patternComparison` basically traverses each pair of pattern's nodes until they have different values and keeps track of which pattern has the greater one, returning 0 if they are equal, 1 if the first array is *larger* than the second one, and -1 otherwise. The parameters of this routine are two arrays with the pattern values and the length of those arrays (must be equal on both).

As we explained before, `patternSearch` and `patternComparison` are intensively used by the routine that calculates the fitness function. In the worst case, if the training image pattern list \mathcal{L} has size n and the template has t nodes, a search in the list will perform $t \times O(\log_2 n)$ comparisons. For very large templates, this pattern search is very time consuming and different data structures must be used in order to get a reasonable execution time in the search task.

4.2. Calculate fitness($indiv_k$)

Following Algorithm 2, which explains all the steps involved in the calculation of the fitness function, in this subsection we explain its associated routine presented in code 1. In this routine, the input $indiv_k$ is a 2D integer array with an image loaded (an *individual* in the genetic algorithm terminology). The inputs tem_nodes , tem_rows , tem_cols , tem_coord_rows and tem_coord_cols correspond to the specific *template* that we are using. For example, in Fig. 2, $tem_nodes = 17$, $tem_rows = tem_cols = 9$ and tem_coord_rows and tem_coord_cols are arrays that store respectively the row and column coordinates of the nodes in template T .

```

subroutine fitnessFunction(
  rows,cols,           !! image dimesions
  indivk,             !! image values
                    !! (2D array of size rows x cols)
  tem_rows, tem_cols, tem_nodes,
  tem_coord_rows,tem_coord_cols,
                    !! template dimensions
                    !! template useful coordinates
                    !! (1D arrays of size tem_nodes)
  value)             !! returned value
use patternOperations !! contains global variables npatterns and patternList

... !! declaration of input/output parameters and local variables
integer(4)          :: freq_aux(npatterns), localPattern(tem_nodes)

do ii=1,npatterns
  freq_aux(ii)=patternList(ii)%frequency
end do
value=0
do icol = 0,cols-tem_cols
  do irow = 0,rows-tem_rows
    do inode = 1, tem_nodes
      localPattern(inode)=indivk(
        irow+tem_coord_rows(inode),
        icol+tem_coord_cols(inode)
      )

      end do
      call patternSearch(tem_nodes,localPattern, pos)
      if(pos/= -1) freq_aux(pos)=freq_aux(pos)-1
    end do
  end do
do ii=1,npatterns
  value=value+freq_aux(ii)*freq_aux(ii)
end do
end subroutine fitnessFunction

```

Code 1: Subroutine fitnessFunction

We can see that in this implementation, the values of O_p , the weight factors described in Eq. (1), are equal to 0 if $freq_{\pi}(p) = 0$ and equal to 1 otherwise (we only take into account patterns that are present in the training image).

5. Code optimization

We have applied several code optimization techniques to the previous described routines, in order to better exploit the CPU resources of the sequential execution. These optimizations can be grouped as: increase data locality, improve stack memory usage, code specialization of fitness routine, branch and load reductions.

5.1. Increasing data locality of the main data structures

The routines `patternSearch` and `patternComparison` are based on the global data structure `patternList`. The first modification consists in adapting this structure to the column-major order of the Fortran language in order to exploit the data and temporal locality.

The column-major order in Fortran is related to the way in which the CPU accesses the data stored in memory. In this order, the matrices are accessed using the address $row + (col - 1) * numrows$ with $numrows$ fixed. The cache lines that are moved from the main memory to the cache memory consist in contiguous memory addresses of fixed size. Leaving col fixed and traversing first all values of row , we can minimize the accesses to non-contiguous memory addresses [22]. Modifying the data structures in order to increase this kind of accesses reduces the cache data misses, reducing the overall execution time.

The new structure is simply a 2D array in which the row size is the number of nodes in the template, tem_nodes , and the column size is the number of patterns found in the training image.

```

module patternOperations
implicit none
integer(4) :: npatterns

type patternType
integer(4),pointer :: pattern(:)
integer(4) :: frequency
end type patternType

type(patternType),pointer :: patternList(:)

contains
...
end module patternOperations

```

Code 2: Original data structures

```

module patternOperations
implicit none
integer(4) :: npatterns

!! size = npatterns
integer(4), pointer :: frequency(:)

!! size = tem_nodes X npatterns
integer(4), pointer :: patternList(:, :)

contains
...
end module patternOperations

```

Code 3: Re-designed data structures

5.2. Using the stack memory to store local arrays

In Fortran, using the Intel's compiler `ifort`, we can use the option `-auto`, which causes all local, non-SAVED variables to be allocated on the run-time stack, including fixed-length arrays. The default is `-auto-scalar`, saving all the scalar variables in the run-time stack. The main advantage is that the access time to the run-time stack is faster than the time access to the run-time heap (which stores the dynamically allocated memory), which decreases the execution time [25]. A negative side of this option is related with the stack size. The stack has a maximum size fixed before the execution and if that size is exceeded, a stack overflow error can be obtained. Also, if we use several threads with OpenMP, each thread has its own stack memory space, so the amount of global stack space is increased proportionally to the number of threads. In order to avoid a stack overflow, we can calculate exactly how much data we need to allocate before compile time and see if it can fit in the stack. If it is too big, we can add more space to the stack, for example using the `ulimit` command in Linux operating systems or setting the environment variable `OMP_STACK_SIZE` with an appropriate value.

5.3. Specialization of fitness function to an input template

Given an input template, we can specialize our fitness function routine in order to exploit the data accesses provided by the template. For example, in the fitness function of code 1 using the template of Fig. 2, the variables `icol + 1, ..., icol + 9` are calculated several times, but in reality we only need to calculate them one time per row-iteration and keep their values stored in an auxiliary variable. This allows for avoiding the accesses to the coordinate arrays `tem_coord_rows` and `tem_coord_cols` (see code 5) and keeping the cache memory *clean* for other data.

If we denote by $t = \text{tem_nodes}$, $n = \text{cols} - \text{tem_cols}$ and $m = \text{rows} - \text{tem_rows}$, the total number of memory accesses performed by the fitness routine (only taking into account the arrays `tem_coord_rows`, `tem_coord_cols` and `indivk`) is $3 \times t \times n \times m$. Using this optimization the total number of memory accesses for the same arrays is $t \times n \times m$ with a reduction of $3 \times$ less memory accesses than the original scenario. In the modified code depicted in code 5, using the template described in Fig. 2, loop unrolling and common subexpression eliminations are included in order to eliminate the accesses of the arrays `tem_coord_rows` and `tem_coord_cols` and re-utilize the values of `icol + 1, ..., icol + 9`.

```

subroutine fitnessFunction(...)
...
...
do icol = 0, cols-tem_cols

do irow = 0, rows-tem_rows

do inode = 1, tem_nodes
localPattern(inode)=indivk(
irow+tem_coord_rows(inode),
icol+tem_coord_cols(inode)
)
end do

call patternSearch(tem_nodes, localPattern, pos)
if (pos/= -1) then
freq_aux(pos)=freq_aux(pos)-1
end if
end do
end do
...
end subroutine fitnessFunction

```

Code 4: Original fitnessFunction routine supporting a general template

```

subroutine fitnessFunction(...)
...
integer(4) :: rowplus1, rowplus4, rowplus5, rowplus6, rowplus9
integer(4) :: colplus1, colplus4, colplus5, colplus6, colplus9
...
do icol = 0, cols-tem_cols
colplus1=icol+1
colplus4=icol+4
colplus5=icol+5
colplus6=icol+6
colplus9=icol+9
do irow = 0, rows-tem_rows
rowplus1=irow+1
rowplus4=irow+4
rowplus5=irow+5
rowplus6=irow+6
rowplus9=irow+9
localPattern(1)=indivk(rowplus1, colplus1)
localPattern(2)=indivk(rowplus5, colplus1)
localPattern(3)=indivk(rowplus9, colplus1)
...
localPattern(15)=indivk(rowplus1, colplus9)
localPattern(16)=indivk(rowplus5, colplus9)
localPattern(17)=indivk(rowplus9, colplus9)
call patternSearch(tem_nodes, localPattern, pos)
if (pos/= -1) then
freq_aux(pos)=freq_aux(pos)-1
end if
end do
end do
...
end subroutine fitnessFunction

```

Code 5: Specialized fitnessFunction routine addressed to the template of figure 2

5.4. Branch reduction

The routine `patternComparison`, described in code 6, which is the most intensive in terms of execution time, can be re-designed in order to reduce the number of branch instructions (control flow, for example `if-then-else` or `while` loops) executed inside a loop.

When a branch instruction is processed by the CPU, some cycles may be lost due to an incorrect branch prediction or an expensive condition evaluation [22]. For that reason, we know that reducing the number of branches and relaxing their

boolean conditions are good practices in order to reduce the overall execution time. A first version of the modified routine consists in the replacement of the `if-elseif` statement by a simple arithmetic operation `value = onePattern1(ii) - onePattern2(ii)`. In this new version, the output `value` is equal to 0 if the patterns are equal, `value > 0` if `onePattern1` is greater and `value < 0` if `onePattern2` is greater.

In this re-design, almost all the branch instructions were eliminated from the original routine. Only the ones injected in the boolean conditions of the `while` loop remain. One possible way to get rid of the evaluation of those conditions is to unroll the loop and use explicit evaluations of each statement incrementing the value of the variable `ii`. Using this technique we are avoiding conditional evaluations and reducing the number of branch instructions performed by the CPU. The re-designed routine can be viewed in code 7. This modified routine is specialized and depends on the number of nodes of the template. In this case we are using the template described in Fig. 2.

```
subroutine patternComparison(length,onePattern1,onePattern2,value)
!
! Compare two patterns.
! If value = 0, both patterns are equals
! If value = 1, the first pattern is bigger
! If value = -1, the second pattern is bigger
!
integer(4), intent(in) :: length
integer(4), intent(in) :: onePattern1(length)
integer(4), intent(in) :: onePattern2(length)
integer(4), intent(out) :: value
integer(4) :: ii
value = 0
ii = 0
do while ( value == 0 .and. ii < length )
  ii = ii + 1
  if ( onePattern1(ii) > onePattern2(ii) ) then
    value = 1
  elseif ( onePattern1(ii) < onePattern2(ii) ) then
    value = -1
  end if
end do
end subroutine patternComparison
```

Code 6: Original `patternComparison` routine supporting general template geometry

```
subroutine patternComparison(length,onePattern1,onePattern2,value)
!
! Compare two patterns.
! If value = 0, both patterns are equals
! If value > 0, the first pattern is bigger
! If value < 0, the second pattern is bigger
!
integer(4), intent(in) :: length
integer(4), intent(in) :: onePattern1(length)
integer(4), intent(in) :: onePattern2(length)

value=onePattern1(1) - onePattern2(1)
if(value/=0)return
value=onePattern1(2) - onePattern2(2)
if(value/=0)return
...
value=onePattern1(16) - onePattern2(16)
if(value/=0)return
value=onePattern1(17) - onePattern2(17)
if(value/=0)return

end subroutine patternComparison
```

Code 7: Specialized `patternComparison` routine with reduced number of branch instructions and minimal boolean conditionals

5.5. Load reduction

In the previous optimized routine `patternComparison` (code 7) the value of the array `onePattern2` is invariant through the entire execution of the caller routine `patternSearch`, implemented as depicted in code 8.

Its values correspond to the pattern that is going to be searched in the pattern database stored in `patternList`. Storing each of the array cells in register variables (up to the maximum number of integer registers available in the CPU) is an alternative that reduces the number of loads performed by the CPU. The variables `op1` to `op17` store the values of `onePattern` (1) to `onePattern` (17), using the template of Fig. 2. Using an inlined version of the routine `patternComparison` together with the registers `op1` to `op17` allows us to reduce by a half the number of loads executed and also to reduce to zero the number of routine calls to `patternComparison`. Code 9 depicts this modifications.

```
subroutine patternSearch(length,onePattern,pos)
!
! Search a pattern in the pattern
! database of training image
! If pos=-1, found
! If pos=-1, not found
!
use patternOperations
integer(4), intent(in) :: length,onePattern(length)
integer(4), intent(out) :: pos
integer(4) :: minn,maxx,value,ii,jj
integer(4) :: isFound

isFound=0
minn = 1
maxx = npatterns
pos=-1
do while ( minn <= maxx .and. isFound == 0 )
  pos = int(real(minn+maxx)* 0.5)
  call patternComparison(
    length,
    patternList(pos)%pattern,
    onePattern,
    value
  )
  if (value == 0) then
    isFound = 1
  elseif (value == 1) then
    maxx = pos - 1
  else
    minn = pos + 1
  end if
end do
if (isFound == 0) pos = -1
end subroutine patternSearch
```

Code 8: Original `patternSearch` routine supporting general template geometry

```
subroutine patternSearch(length,onePattern,pos)
!
! Search a pattern in the pattern
! database of training image
! If pos=-1, found
! If pos=-1, not found
!
use patternOperations
integer(4), intent(in) :: length,onePattern(length)
integer(4), intent(out) :: pos
integer(4) :: minn,maxx,value,ii,jj
integer(4) :: op1,...,op17

op1=onePattern(1)
...
op17=onePattern(17)
value=-1
minn = 1
maxx = npatterns
pos=-1
do while ( minn <= maxx )
  pos = shiftr(minn+maxx,1)
  do
    value=patternList(1,pos) - op1
    if(value/=0)exit
    ...
    value=patternList(17,pos) - op17
    if(value/=0)exit
    return !! value==0
  end do
  if (value > 0) then
    maxx = pos - 1
  else
    minn = pos + 1
  end if
end do
pos = -1
end subroutine patternSearch
```

Code 9: Specialized `patternSearch` routine with reduced number of loads

An additional optimization was added to the final version of the `patterSearch` routine. The expression

```
pos = int(real(minn + maxx) * 0.5)
```

was replaced by

```
pos = shiftr(minn + maxx, 1)
```

in order to replace the floating-point operation $\frac{\text{minn} + \text{maxx}}{2}$ by one right-shift in the bits of the value `minn + maxx`.

6. Parallelization

6.1. Fine grained parallelization with OpenMP

A parallelization scheme using OpenMP over a genetic-based MPS code was described in [35]. In that work, the proposed parallelization was based on parallel do-loops over the fitness function routine. The parallelization presented in this section is essentially the same, based on parallel do-loops, but taking into consideration the previously described code optimizations.

```
subroutine fitnessFunction(...)
...
do ii=1,npatterns
  freq_aux(ii)=frequency(ii)
end do
value=0

do icol = 0,cols-tem_cols
  ...
  do irow = 0,rows-tem_rows
    ...
    if(pos/= -1) then
      freq_aux(pos)=freq_aux(pos)-1
    end if
  end do
end do

do ii=1,npatterns
  value=value+freq_aux(ii)+freq_aux(ii)
end do
end subroutine fitnessFunction
```

Code 10: Sequential version of `fitnessFunction` routine

```
subroutine fitnessFunction(...)
...
idthread=omp_get_thread_num()+1
numthreads=omp_get_num_threads()
...
do ii=1,npatterns
  freq_aux(ii)=frequency(ii)      !! global
  freq_aux_local(ii,idthread)=0  !! local
end do
value=0
!$OMP PARALLEL PRIVATE( localPattern,pos,icol,irow, &
!$OMP colplus1,...,rowplus9,freq_aux_local )
!$OMP DO
do icol = 0,cols-tem_cols
  ...
  do irow = 0,rows-tem_rows
    ...
    if(pos/= -1) then
      freq_aux_local(pos,idthread)=freq_aux_local(pos,idthread)+1
    end if
  end do
end do
!$OMP END DO
!$OMP END PARALLEL
do jj=1,numthreads
  do ii=1,npatterns
    freq_aux(ii)=freq_aux(ii)-freq_aux_local(ii,jj)
  end do
end do
do ii=1,npatterns
  value=value+freq_aux(ii)+freq_aux(ii)
end do
end subroutine fitnessFunction
```

Code 11: Parallelization of `fitnessFunction` routine using OpenMP

The main difference with the sequential code 1, besides the code optimizations, is the utilization of a local array `freq_aux_local` which stores the frequencies of the patterns calculated locally by all threads. After each thread finishes their corresponding loops (an implicit wait is placed at the end of the loop for thread synchronization) all those frequencies are gathered into the global array `freq_aux` and this array is used to calculate the final value of the fitness function.

Several schedules (*static*, *dynamic* and *guided* with different chunk sizes) were tested but none of them was considerably faster than the others, so we choose to stay using the *static* schedule in all of our tests. No profit was obtained after tuning the chunk size, because no false sharing (a review of this topic can be found in [7]) was introduced in the calculations since no writing is done in the matrix `indivk`. The default value of chunk size was used in the execution of the tests.

6.2. Coarse grained parallelization with MPI + OpenMP

If the initial population is distributed in several processes, the amount of work that each process does decreases with a reduction in the execution time. Following ideas from [5], we implement an *island-based* parallelization scheme in which each process waits for the best individuals calculated by the others (other *islands*), and asynchronously each process sends to everybody the best individual calculated by itself. In this way, all processes share the same best individuals and each one can combine them in order to breed a new generation. The steps are depicted in Algorithm 3. Using this distribution of load, each process can use several threads with OpenMP as in the previous subsection and the code optimizations explained in the previous section. The result of this integration is a hybrid distributed-shared memory parallelization based on the MPI and OpenMP programming models.

Algorithm 3. Parallel canonical genetic algorithm (based on island model)

```

1: INPUT:  $P$  processes,  $\frac{N}{P}$  individuals per process (population)
2: for each process  $p$  (in parallel) do
3:   Evaluate a fitness function  $fitness$  in each individual
4:   while termination criteria is not achieved do
5:     {Breed a new generation}
6:     Sort the individuals by their fitness function value
7:     if no improvement is measured in the local population then
8:       Restart: select some individuals and restart their bits
9:       Sort the individuals by their fitness function value
10:    end if
11:    Selection: select the best individuals based on its fitness functions
12:    Crossover: breed new individuals crossing bits of individuals from the selection
13:    Mutation: breed new individuals mutating some bits of individuals from the selection
14:    Replace old individuals by new ones
15:    Send (asynchronously) the best individual to the other processes
16:    Receive (asynchronously) the best individuals from other processes and copy them in the population (using the
    memory space of the worst individuals)
17:    Evaluate a fitness function  $fitness$  in each individual
18:  end while
19: end for
20: OUTPUT: best individual in population of process master.

```

The only drawback of this parallel implementation is the reduction of the population size in each island. Each process will have a population of $\frac{N}{P}$ individuals, and if N is not sufficiently large, local minimum will be achieved earlier in the convergence process, restricting the search for better solutions. A solution to this problem is to set large population sizes according to the number of processes involved in the executions and the size of the search space, in our case, the number of template nodes and training image nodes. Several tests must be done in order to get a good estimation of the optimal population size. In this work, however, our focus was to get reasonable values of performance and convergence, leaving this topic for future research.

The key part of the parallel strategy is the communication between the processes, which is equivalent to an all-to-all broadcast of the best individuals stored in each island. This communication can be implemented with MPI using asynchronous `MPI_Isend/MPI_Irecv/MPI_Wait` or with the intrinsic routine `MPI_Allgather`.

A performance analysis tool called Paraver, described in [28], was used to generate trace views associated to an execution of the proposed implementation using 16 processes with 1 and 12 threads each. In these traces we can see the different states of execution described in Algorithm 3. In the X-axis we have the execution time in microseconds and in the Y-axis we have the states of processes and threads. In Fig. 4 we can see the overall execution time with both processors sets using the same time scale.




Fig. 4. Execution trace (using the same time scale) with two processor sets, 16×1 (top) and 16×12 (bottom) processes \times threads, using a training image of size 1000×1000 running 10 generations of the Algorithm 3.

7. Results

7.1. Timing

In order to study the effectiveness of the optimizations and parallelizations, we use a test routine which loads two 2D images, using one as training image and the other one as individual. After that, it calculates the fitness function of the individual using the previous routines. All measurements are average of 100 executions of this test routine. Two image sizes were used, 100×100 and 1000×1000 . They are shown in Fig. 5. Additionally, full iterations were measured running 30 generations of the genetic algorithm.

The code optimizations were tested incrementally: if an optimization X reduces the execution time, all subsequent optimizations include that optimization. The code compilation was done using `ifort` (Intel Fortran Compiler) version 12.0.4 (20110427) [26], without any added options or flags. The default level of compiler optimization is `-O2`. The tests were executed in a cluster of compute nodes running Linux operating system where every node has two processors Intel Xeon E5649 6-Core (12 CPUs per node) at 2.53 GHz with 24 GB of RAM memory, 12 MB of cache memory and 250 GB local disk storage. The execution times and speed-up values obtained is detailed in Tables 1–5.

Based on the results from Table 1, the code optimization accelerates considerably the sequential execution of the fitness function calculation. A speedup of $2.03\times$ and $1.87\times$ was obtained for images of 100×100 and 1000×1000 respectively. The best results were obtained after re-designing the data structures that handle the pattern list, allowing it to reduce the execution time by improving the locality of the memory accesses.




Fig. 5. Training images of size 100×100 (left) and 1000×1000 (right).

Table 1

Code optimization: fitness function calculation of a 100×100 and 1000×1000 images using the template in Fig. 2.

Optimization	100×100		1000×1000	
	Time (s)	Speed up	Time (s)	Speed up
Baseline	0.00319	1	0.402347	1
Increase data locality	0.00213	1.49 \times	0.26368	1.53 \times
Use Stack memory	0.00208	1.53 \times	0.257903	1.56 \times
Specialization of fitness	0.00181	1.76 \times	0.243201	1.65 \times
Branch reduction	0.00160	1.99 \times	0.222511	1.80 \times
Load reduction	0.00157	2.03 \times	0.214222	1.87 \times

Table 2

Fine-grained parallelization without code-optimizations: fitness function calculation of a 100×100 and 1000×1000 images using the template in Fig. 2.

Threads	100×100		1000×1000	
	Time (s)	Speed up	Time (s)	Speed up
1	0.003190	1 \times	0.402347	1 \times
2	0.001800	1.77 \times	0.186401	2.15 \times
4	0.031770	0.10 \times	0.094410	4.26 \times
6	0.020684	0.15 \times	0.065002	6.18 \times
8	0.027191	0.11 \times	0.055411	7.26 \times
10	0.046754	0.06 \times	0.049403	8.14 \times
12	0.027156	0.11 \times	0.037000	10.87 \times

Table 3Fine-grained parallelization with code-optimizations: fitness function calculation of a 100×100 and 1000×1000 images using the template in Fig. 2.

Threads	100×100		1000×1000	
	Time (s)	Speed up	Time (s)	Speed up
1	0.003190	1×	0.402347	1×
1 + code-opt	0.001573	2.03×	0.222202	1.80×
2 + code-opt	0.002028	1.57×	0.109205	3.68×
4 + code-opt	0.012321	0.25×	0.056436	7.12×
6 + code-opt	0.015411	0.20×	0.041220	9.76×
8 + code-opt	0.027145	0.11×	0.036421	11.04×
10 + code-opt	0.016694	0.19×	0.030886	13.02×
12 + code-opt	0.017766	0.17×	0.023531	17.09×

Table 4Coarse-grained parallelization without code optimizations: average generation step (30 generations) of the genetic algorithm using a population of 1000 individuals, each one a 100×100 and 1000×1000 images respectively, using the template in Fig. 2.

Processes × Threads	100×100		Processes × Threads	1000×1000	
	Time (s)	Speed up		Time (s)	Speed up
1 × 1	3.466	1×	1 × 1	312.133	1×
1 × 2	1.852	1.87×	1 × 12	55.62	5.61×
2 × 2	0.934	3.71×	2 × 12	27.53	11.33×
4 × 2	0.556	6.23×	4 × 12	14.03	22.24×
8 × 2	0.295	11.74×	8 × 12	7.26	42.99×
16 × 2	0.213	16.27×	16 × 12	4.02	77.64×

Table 5Coarse-grained parallelization with code optimizations: average generation step (30 generations) of the genetic algorithm using a population of 1000 individuals, each one a 100×100 and 1000×1000 images respectively, using the template in Fig. 2.

Processes × Threads	100×100		Processes × Threads	1000×1000	
	Time (s)	Speed up		Time (s)	Speed up
1 × 1	3.466	1×	1 × 1	312.133	1×
1 × 1 + code-opt	2.032	1.70×	1 × 12 + code-opt	42.911	7.10×
2 × 1 + code-opt	1.027	3.37×	2 × 12 + code-opt	22.102	14.12×
4 × 1 + code-opt	0.561	6.17×	4 × 12 + code-opt	11.354	27.49×
8 × 1 + code-opt	0.310	11.18×	8 × 12 + code-opt	5.916	52.76×
16 × 1 + code-opt	0.221	15.68×	16 × 12 + code-opt	3.085	101.17×

According to the results from Tables 2 and 3, the fine-grain parallelization allows to accelerate the execution considerably only when the size of the training image and realizations are large, in our case a size of 1000×1000 reaching a speedup of 17.09× with 12 threads in the best case. With a size of 100×100 no profit from the parallelization was observed (except in the non-optimized code with 2 threads), due to the small amount of work that each thread performs compared with the overhead introduced by thread management. The best speedup result, 2.03×, was obtained using the optimized code with only one thread.

The coarse-grain parallelization accelerates the execution proportionally to the number of processes involved in it. This feature is obtained after distributing the workload evenly among the processes. In our case, the workload is represented by the population, which is divided among the processes. According to Tables 4 and 5, the best results achieved in terms of speedup were obtained using the hybrid parallelization scheme with optimized code, MPI and OpenMP on large training and realization images of size 1000×1000 . The thread selection policy used to define how many threads will run in each test was based in the best results obtained in the corresponding fine-grain tests from Tables 2 and 3, namely, 2 and 12 threads for small and large images respectively, using non-optimized code, and 1 and 12 threads for small and large images respectively, using optimized-code.

7.2. Convergence

In order to test the convergence of the method using the parallel optimized code, we choose a fixed set of parameters for the genetic algorithm and with those parameters we observe the behavior of the fitness function and the realization obtained, starting from a randomly generated population of individuals. Those parameters are:

- population size: 640 individuals,
- mutation rate: described in the next paragraph,
- crossover percentage: 50%, corresponding to the number of individuals selected to perform a crossover with other individual,
- restart percentage: 10%, corresponding to the percentage of individuals that will be restarted, in each restart step,
- number of cut-points: 10%, corresponding to the number of cut-points in the crossover operation,
- number of mutation nodes: described in the next paragraph.

In order to accelerate the convergence, a multi-point mutation strategy was implemented. In this strategy, a random node and a random category were selected (in our tests, we have only 2 categories). Using an influence radius r calculated a priori (in our case $r = 0.02 \times \min(\#rows, \#columns)$) all nodes that fall within the circle centered in the random node with radius r are *mutated* (changed) into a new selected category (leaving the conditionant nodes without modification). Additionally, a cyclic-cooling scheme was implemented in order to control the variability of the population, by means of reducing the mutation ratio by a factor $\lambda \in (0, 1)$ each 1000 generations. Starting from a mutation rate $m = 1$, if λm goes below a threshold, the ratio is restarted from $m = 1$ and a new mutation cycle begins.

Samples of realization images and convergence plots of function 1 are included in Figs. 6, 7. The convergence plots show the relative decrease in percentage of the fitness function with respect to the initial value (the best fitness value obtained in the first population of individuals). As we can see in these figures, the simulated images are not equal to the training image, which is not bad, because one of the objectives of the simulation process using MPS methods is to obtain simulated images that fits the underlying statistics of the training image but not being necessarily equal. If conditional data is used (nodes with information from the training image which are not modified in the simulation process) the resulting simulated image will be more similar to the training image with the corresponding match of the underlying statistics. The tests that we are considering do not use any conditional data. Therefore they must be considered as a worst case scenario in terms of convergence.




Fig. 6. Simulated realization of size 100×100 (left) and convergence plot showing 17000 generations.

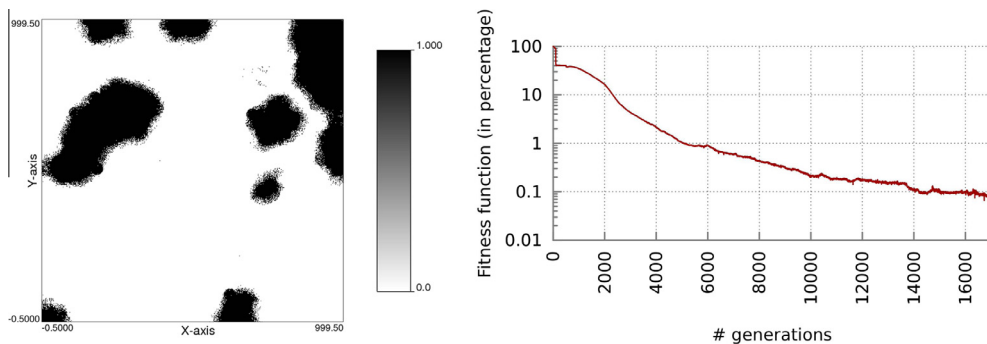


Fig. 7. Simulated realization of size 1000×1000 (left) and convergence plot showing 17000 generations.

8. Conclusions and future work

The proposed hybrid parallelization using an optimized code has shown reasonable speedup results, according to the time measurements reported in Section 7. The differences between non-optimized and optimized code execution are considerable and justify our research in this recent topic. A negative aspect has to do with the loss of generality of some of the proposed

optimizations. Applying routine specialization or branch/load reduction introduces several modifications in the corresponding code, were specific values of the size of the template are used. If the size or geometry of the template changes, we need to modify those routines in order to adapt them to the new template's values. If the complete code has to be used in an agile framework allowing the utilization of user-defined templates, this behavior can generate inconveniences because each new template needs its own optimized routines, with their corresponding creation, compilation and inclusion in the main application. A possible solution to this problem can be the adoption of auto-tuning techniques to automatically deploy new kernels according to the geometry of the template in use. Examples of auto-tuning techniques applied to stencil optimization for Finite Differences PDE solvers can be reviewed in [10,11].

The strategies proposed to accelerate the convergence allow us to get fast realizations, using reasonable small populations (40 individuals per process) and a small restart percentage (10% of the population is restarted). However, a further research on the acceleration of convergence by tuning the genetic parameters (population size, mutation and crossover rates, percentage of selection, percentage of restarted population, number of cross points and mutation points, among others) is left open for future research. Also, several other research topics can be explored, among them we can mention: dynamic mutation rate (using the full annealing scheme), non-linear crossovers (using external information to mix two realizations with some physical or geological interpretation), or selective restart (allowing to reset only individuals that meet certain properties).

This work focuses on 2D training images and realizations. However 3D models are used in real geostatistical scenarios. In order to adapt our code to the 3D scenario, several further optimizations and modifications must be done, but these are left as future work. Among the most relevant ones we can mention the specialization of the fitness evaluation to 3D templates, use of efficient data structures to manage large 3D images and 1D individuals, new methods to perform crossover and mutations in these individuals, and modifications of the fine and coarse grain strategies to the new 3D scenario. Another related future work could be the application of the fitness calculation optimized routines into the simulated annealing scheme of simulation as described in [14], using the MPI implementation described in [34] or the standard routines implemented in the GSLIB library from [16].

Finally, another possible research area is related to explore new computer architectures with this algorithm. Among the possible alternatives are NVIDIA's GPUs (using two programming models, CUDA and OpenACC), Intel's MICs (several cores in one chip and accelerators working together in Intel architectures) or energy-efficient new supercomputers that will be available in the next years.

References

- [1] B. Arpat, J. Caers, *Stochastic simulation with patterns*, *Math. Geol.* 39 (2007) 177–203.
- [2] J.B. Boisvert, S. Lyster, C.V. Deutsch, Constructing training images for veins and using them in multiple-point geostatistical simulation, in: E.J. Magri (Ed.), 33rd International Symposium on Application of Computers and Operations Research in the Mineral Industry, APCOM 2007, 2007, pp. 113–120.
- [3] J. Caers, A.G. Journel, Stochastic reservoir simulation using neural networks trained on outcrop data, in: SPE Annual Technical Conference and Exhibition, New Orleans, LA, September 1998. Society of Petroleum Engineers, SPE paper # 49026, 1998, pp. 321–336.
- [4] J. Caers, X. Ma, Modeling conditional distributions of facies from seismic using neural nets, *Math. Geol.* 34 (2002) 143–167.
- [5] E. Cantú-Paz, A survey of parallel genetic algorithms, *Calculateurs paralleles, Reseaux et Systems Repartis* 10, 1998.
- [6] R. Chandra, L. Dagum, D. Kohr, D. Maydan, J. McDonald, R. Menon, *Parallel Programming in OpenMP*, Morgan Kaufmann Pub. Inc., San Francisco, CA, USA, 2001.
- [7] D. Culler, J. Singh, A. Gupta, *Parallel computer architecture: a hardware/software approach*, The Morgan Kaufmann Series in Computer Architecture and Design, 1st ed., Morgan Kaufmann, 1998.
- [8] C. Daly, Higher order models using entropy, markov random fields and sequential simulation, in: O. Leuangthong, C.V. Deutsch (Eds.), *Geostatistics Banff 2004*, Springer, 2005, pp. 215–224.
- [9] C. Daly, C. Knudby, Multipoint statistics in reservoir modelling and in computer vision, *Petroleum Geostatistics A*, 2007.
- [10] K. Datta, M. Murphy, V. Volkov, S. Williams, J. Carter, L. Oliker, D. Patterson, J. Shalf, K. Yelick, Stencil computation optimization and auto-tuning on state-of-the-art multicore architectures, in: *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, IEEE Press, Piscataway, NJ, USA, 2008.
- [11] K. Datta, S. Williams, V. Volkov, J. Carter, L. Oliker, J. Shalf, K. Yelick, Auto-tuning the 27-point stencil for multicore, in: *Proc. iWAPT2009: The Fourth International Workshop on Automatic Performance Tuning*, 2009.
- [12] K.A. De Jong, Adaptive system design: a genetic approach, in: *IEEE Transactions on Systems, Man, and Cybernetics SMC-10*, 1980, pp. 566–574.
- [13] C. Deutsch, L. Wang, Hierarchical object-based stochastic modeling of fluvial reservoirs, *Math. Geol.* 28 (1996) 857–880.
- [14] C.V. Deutsch, Annealing techniques applied to reservoir modeling and the integration of geological and engineering (well test) data, Doctoral dissertation (Ph.D. thesis), Stanford University, 1992.
- [15] C.V. Deutsch, *Geostatistical Reservoir Modeling*, Oxford University Press, New York, 2002.
- [16] C.V. Deutsch, A.G. Journel, *GSLIB: Geostatistical Software Library and User's Guide*, Oxford University Press, New York, Oxford, 1992.
- [17] K. Eskandari, S. Srinivasan, Growthsim – a multiple point framework for pattern simulation. *Petroleum Geostatistics A*, 2007.
- [18] C. Faucher, A. Scahier, D. Marcotte, A new patchwork simulation method with control of the local-mean histogram, *Stochastic Environ. Res. Risk Assess.* 27 (2013) 1–21.
- [19] D.E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, 1989.
- [20] S.L. Graham, P.B. Kessler, M.K. McKusick, Gprof: a call graph execution profiler, *SIGPLAN Not.* 39 (2004) 49–57.
- [21] F. Guardiano, M. Srivastava, Multivariate geostatistics: beyond bivariate moments, *Geostatistics Troia 1*, 1993, pp. 133–144.
- [22] J.L. Hennessy, D.A. Patterson, *Computer Architecture: A Quantitative Approach*, Morgan Kaufmann, San Mateo, CA, 1990.
- [23] J. Holland, *Adaptation in Natural and Artificial Systems*, U. of Michigan Press, 1975.
- [24] S. Hong, J.M. Ortiz, C.V. Deutsch, Multivariate density estimation as an alternative to probabilistic combination schemes for data integration, in: J.M. Ortiz, X. Emery (Eds.), *Geostats 2008 – Proceedings of the Eighth International Geostatistics Congress*, vol. 1, Gecamin Ltda., Santiago, Chile, 2008, pp. 197–206.
- [25] Intel Corporation, 2006, Fortran Compiler Use of Temporaries. <<http://software.intel.com/file/20415>>, 2012.
- [26] Intel Corporation, 2011, Intel® Fortran Composer XE 2011. <<http://software.intel.com/en-us/articles/intel-composer-xe/>>, 2011.
- [27] S. Krishnan, A.G. Journel, Spatial connectivity: from variograms to multiple-point measures, *Math. Geol.* 35 (2003) 915–925.
- [28] J. Labarta, S. Girona, V. Pilet, T. Cortes, J.M. Cela, A parallel program development environment, in: *Proceedings of the 2th. Int. Euro-Par Conference*, Springer, 1995, pp. 665–674.

- [29] S. Lyster, C.V. Deutsch, Mps simulation in a gibbs sampler algorithm, in: J.M. Ortiz, X. Emery (Eds.), *Geostats 2008 – Proceedings of the Eighth International Geostatistics Congress*, vol. 1, Gecamin Ltda., Santiago, Chile, 2008, pp. 79–88.
- [30] J.M. Ortiz, Characterization of high order correlation for enhanced indicator simulation, Unpublished doctoral dissertation (Ph.D. thesis), University of Alberta, 2003.
- [31] J.M. Ortiz, C.V. Deutsch, Indicator simulation accounting for multiple-point statistics, *Math. Geol.* 36 (2004) 545–565.
- [32] J.M. Ortiz, X. Emery, Integrating multiple point statistics into sequential simulation algorithms, in: O. Leuangthong, C.V. Deutsch (Eds.), *Geostatistics Banff 2004*, Springer, 2005, pp. 969–978.
- [33] A. Parra, J.M. Ortiz, Conditional multiple-point simulation with a texture synthesis algorithm, in: *IAMG 09 Conference*, Stanford University, 2009.
- [34] O. Peredo, J.M. Ortiz, Parallel implementation of simulated annealing to reproduce multiple-point statistics, *Comput. Geosci.* 37 (2011) 1110–1121.
- [35] O. Peredo, J.M. Ortiz, Multiple-point geostatistical simulation based on genetic algorithms implemented in a shared-memory supercomputer, in: P. Abrahamsen, R. Hauge, O. Kolbjørnsen (Eds.), *Geostatistics Oslo 2012, Quantitative Geology and Geostatistics*, vol. 17, Springer, Netherlands, 2012, pp. 103–114.
- [36] M.J. Pyrcz, S. Strebelle, A showcase of event-based geostatistical models, in: J.M. Ortiz, X. Emery (Eds.), *Geostats Oslo 2012*, vol. 2, Gecamin Ltda., Santiago, Chile, 2008, pp. 1143–1148.
- [37] H. Rezaee, G. Mariethoz, M. Koneshloo, O. Omid Asghari, Multiple-point geostatistical simulation using the bunch-pasting direct sampling method, *Comput. Geosci.* 54 (2013) 293–308.
- [38] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, J. Dongarra, *MPI-The complete reference, The MPI Core, 2nd (revised) ed.*, vol. 1, MIT Press, Cambridge, MA, USA, 1998.
- [39] J. Straubhaar, P. Renard, G. Mariethoz, R. Froidevaux, O. Besson, An improved parallel multiple-point algorithm using a list approach, *Math. Geosci.* (2011) 1–24.
- [40] S. Strebelle, Conditional simulation of complex geological structures using multiple-point statistics, *Math. Geol.* 34 (2002) 1–21.
- [41] S. Strebelle, A.G. Journel, Sequential simulation drawing structures from training images, in: *6th International Geostatistics Congress*, Cape Town, South Africa. Geostatistical Association of Southern Africa, 2000.
- [42] H. Tjelmeland, Stochastic models in reservoir characterization and Markov random fields for compact objects, Unpublished doctoral dissertation (Ph.D. thesis), Norwegian University of Science and Technology, 1996.
- [43] W. Xu, Conditional curvilinear stochastic simulation using pixel-based algorithms, *Math. Geol.* 28 (1996) 937–949.
- [44] S. Zanon, *Advanced Aspects of Sequential Gaussian Simulation* (Master's thesis), University of Alberta, 2004.
- [45] T. Zhang, P. Switzer, A. Journel, Filter-based classification of training image patterns for spatial simulation, *Math. Geol.* 38 (2006) 63–80.