

Received July 5, 2020, accepted July 27, 2020, date of publication August 7, 2020, date of current version August 19, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3015045

Auto-Tuning of Dynamic Scheduling Applied to 3D Reverse Time Migration on Multicore Systems

ÍTALO A. S. ASSIS^{1,2}, JOÃO B. FERNANDES¹, TIAGO BARROS¹,
AND SAMUEL XAVIER-DE-SOUZA¹, (Senior Member, IEEE)

¹Departamento de Engenharia de Computação e Automação, Universidade Federal do Rio Grande do Norte, Natal 59078-970, Brazil

²Departamento de Informática e Matemática Aplicada, Universidade Federal do Rio Grande do Norte, Natal 59078-970, Brazil

Corresponding author: Ítalo A. S. Assis (italoaug@ufrn.edu.br)

This work was supported by Shell Brasil Petróleo and the Agência Nacional do Petróleo, Gás Natural e Biocombustíveis (ANP) through the Project Novos Métodos de Exploração Sísmica por Inversão Completa das Formas de Onda.

ABSTRACT Reverse time migration (RTM) is an algorithm widely used in the oil and gas industry to process seismic data. It is a computationally intensive task that suits well in parallel computers. Methods such as RTM can be parallelized in shared memory systems through scheduling iterations of parallel loops to threads. However, several aspects, such as memory size and hierarchy, number of cores, and input size, make optimal scheduling very challenging. In this paper, we introduce a run-time strategy to automatically tune the dynamic scheduling of parallel loops iterations in iterative applications, such as the RTM, in multicore systems. The proposed method aims to reduce the execution time of such applications. To find the optimal granularity, we propose a coupled simulated annealing (CSA) based auto-tuning strategy that adjusts the chunk size of work that OpenMP parallel loops assign dynamically to worker threads during the initialization of a 3D RTM application. Experiments performed with different computational systems and input sizes show that the proposed method is consistently better than the default OpenMP schedulers, *static*, *auto*, and *guided*, causing the application to be up to 33% faster. We show that the possible reason for this performance is the reduction of cache misses, mainly level L3, and low overhead, inferior to 2%. Having shown to be robust and scalable for the 3D RTM, the proposed method could also improve the performance of similar wave-based algorithms, such as full-waveform inversion (FWI) and other iterative applications.

INDEX TERMS Auto-tuning, coupled simulated annealing, reverse time migration, openMP, dynamic scheduling.

I. INTRODUCTION

Seismic reflection surveying is the best known and used geophysical method for subsurface imaging. Oil and gas exploration is its main application [1]. Its main objective is to generate an image of a subsurface region to identify structures of interest.

Seismic data can go through several processing steps to improve the signal-to-noise ratio (SNR) and the seismic image's resolution. One of the most important of these steps is migration, which is responsible for positioning seismic reflection events in their correct place when imaging the subsurface. In this context, reverse time migration (RTM) [2], [3] has been widely used as a migration technique to more

accurately take into account the wave propagation effects resulting in subsurface images with higher definition.

Simulating wave propagation comprises the majority of an RTM and is computationally intensive, especially for the three-dimensional case. Therefore, the computational cost is the main factor limiting the application of RTM, as well as for several other geophysical algorithms [4], [5]. For this reason, parallel computing techniques have been widely applied to these methods (e.g., [6]).

Load balancing is one of the main aspects to be considered in parallel applications. It can be defined as the distribution of the computational load among the available processing resources (e.g., cores, computing nodes). A way to perform load balancing is by dividing the workload in chunks of computation to be distributed among the computational resources either statically or dynamically. In the context of parallel

The associate editor coordinating the review of this manuscript and approving it for publication was Ayaz Ahmad.

applications, auto-tuning techniques are becoming popular in the sense of obtaining portable near-optimal performance (e.g., when scheduling or load balancing tasks) in parallel programming [7].

Load balancing parallel seismic methods, such as RTM, is especially challenging with the rising interest for heterogeneous machines [8]. Nevertheless, as we show in this work, even for homogeneous architectures, a static load balancing may not be optimal.

This paper presents an execution time auto-tuning strategy to automatically find an optimal chunk size for OpenMP [9] dynamic scheduling of parallel loops. This method aims to reduce the run time of iterative applications. For that, we employ a global optimization method called coupled simulated annealing (CSA) [10]. This paper also provides numerical experiments where the proposed auto-tuning is applied to a 3D RTM algorithm.

The contribution of this paper to RTM and the scheduling of parallel loops are:

- 1) the introduction of a novel and robust auto-tuned scheduling strategy for dynamically scheduling parallel loops;
- 2) a methodology to perform the parametrization of the proposed method applied to the 3D RTM;
- 3) a comparison of the proposed method against the default OpenMP schedulers for the 3D RTM showing that the proposed auto-tuning reduces the number of cache misses;
- 4) an overhead analysis showing that the proposed method's overhead was inferior to 2% when employed to the 3D RTM in different computational environments and input sizes;
- 5) a consistent reduction in the run-time of the 3D RTM across five different and non-parametrized architectures.

The outline of the paper is as follows. We first present the basics of our target application: the RTM (Section II), the parallelization strategies made available by OpenMP (Section III) and the optimization method that comprises the proposed auto-tuning, the CSA (Section IV). Then, we provide a detailed description of our RTM implementation (Section V) as well as of the proposed auto-tuning approach (Section VI). Section VII displays the results of the proposed method in comparison with the standard OpenMP schedules. In Section VIII we present a literature review, highlighting the main aspects of our contribution. Section IX concludes this paper.

II. REVERSE TIME MIGRATION FORMULATION

The seismic reflection method consists of three main steps: acquisition, processing, and interpretation of seismic data. In the acquisition, seismic shots, reflected by subsurface interfaces, are recorded at surface level by receivers. The signal recorded by each detector, from each seismic shot, is called a seismic trace. A set of seismic traces is called a

seismogram. Seismograms can be converted to depth estimates of interfaces between different subsurface materials during processing.

After the acquisition step, several techniques can be used to process seismic data. In general, the purpose of processing reflection data is to increase the SNR and improve the vertical resolution of resulting seismic images. Migration is one of the main steps in the seismic data processing. It aims to 1) properly position seismic reflections at the coordinates of the reflector in the subsurface; 2) reduce diffraction effects in the images; 3) improve the spatial resolution.

Modern migration approaches use the seismic wave equation, a partial differential equation describing wave motion, generated by a source in a medium. The scalar equation for 3D acoustic waves is defined as

$$\frac{\partial^2 u(\mathbf{x})}{\partial x_1^2} + \frac{\partial^2 u(\mathbf{x})}{\partial x_2^2} + \frac{\partial^2 u(\mathbf{x})}{\partial x_3^2} = \frac{1}{c(\mathbf{x})^2} \frac{\partial^2 u(\mathbf{x})}{\partial t^2} + s(t), \quad (1)$$

where $\mathbf{x} = (x_1, x_2, x_3)$ are the spatial dimensions, $u(\mathbf{x})$ is the acoustic pressure, $c(\mathbf{x})$ is the propagation velocity and $s(t)$ is the source function at time t .

Spatial and time restrictions should be observed when solving finite differences by a numerical approach [11]. These restrictions are defined as:

$$\max(\Delta x_1, \Delta x_2, \Delta x_3) \leq \frac{c_{\min}}{W f_{\max}} \quad (2)$$

and

$$\Delta t \leq \frac{2 \min(\Delta x_1, \Delta x_2, \Delta x_3)}{\pi c_{\max} \sqrt{3}}, \quad (3)$$

where Δx_1 , Δx_2 and Δx_3 are the spatial sampling of dimensions x_1 , x_2 and x_3 , Δt is the time sampling; f_{\max} is the maximum frequency of $s(t)$; c_{\min} and c_{\max} are the minimum and the maximum values of $c(\mathbf{x})$; and W is the number of grid points per minimum wavelength. According to [11], W must be equal or greater than 4 for high order finite differences schemes. Non-compliance with (2) and (3) would result in numerical dispersion and instability.

Another important aspect is that the geological model encoded in $c(\mathbf{x})$ must be restricted to a finite number of points on a mesh, even though the Earth is heterogeneous and continuous. In order to represent real boundaries, it is common to apply artificial edges to the model limits, to absorb the energy reaching the borders [12].

There are several approaches to migrate seismic data. We use migration by finite differences (or wave equation migration), in which the wave equation is approximated by a finite-difference equation, suitable to be solved by a computer as explained above. One of the main migration methods by finite differences is RTM [2], [3]. In RTM, source and receiver wavefields are propagated forward and backward in time, respectively. RTM imaging relies on the physical property that those pressure waves must correlate at the reflective interfaces.

The core of an RTM can be divided into three stages. The first stage is the simulation of the propagation of a wavefield

resulting from the excitation of a seismic source. The second stage is the backpropagation of wavefields registered in a seismogram. Finally, the third stage is the imaging condition, which is a correlation between the forward and backward propagated wavefields and produces an image of the subsurface. This process is repeated for all the shots of seismic data available.

The propagation and backpropagation steps use the same velocity model shown in (1) as $c(\mathbf{x})$. This model specifies the wave velocity for each mesh point and represents the different properties of the materials and boundaries, in the volume being imaged.

In the imaging condition stage, the wavefields generated by the propagation of the source and the backpropagation of the observed data are correlated pointwise, at each time interval, to generate an image. Mathematically, it is defined as

$$I(\mathbf{x}) = \int_{t=0}^T u_i(\mathbf{x}, t) \cdot u_r(\mathbf{x}, t) dt, \quad (4)$$

where $I(\mathbf{x})$ is the resulting image, $u_i(\mathbf{x}, t)$ is the wavefield propagated with the source excitation, $u_r(\mathbf{x}, t)$ is the wavefield of backpropagated data and T is the total simulation time. The migration of each seismic shot generates an image. These images are stacked to build the total migrated volume.

Each cycle of a seismic survey ends with the interpretation phase. Since both coverage and resolution are better with 3D data, these surveys lead to improved interpretation compared with 2D surveys and are standard today [13].

III. PARALLELIZATION STRATEGY

In this work, we used an RTM algorithm implemented with two degrees of parallelization. The first is the migration of different common-shot (CS) gathers, i.e., seismic data with the same shot coordinates, which is implemented with the message passing interface (MPI) [14], for distributed memory environments. The second is the migration of a single CS gather and is performed in shared memory environments, with OpenMP [9].

The proposed work is applied in the second degree of parallelization, where different loops of the RTM operation of each CS gather are parallelized among the cores of a multicore system, with OpenMP. This parallelization is performed by dividing each loop into loops of smaller sizes, which are computed in the different cores of the multicore system. The size of these smaller loops is usually referred to as the *chunk size*. Our work's primary goal is to balance the computation of the smaller loops by the different cores by choosing the proper chunk size, which is known as workload balancing. The proposed load balancing approach is discussed in more detail in Section VI.

For the parallelization with OpenMP, the `parallel for` construction was employed. This construction automatically distributes the workload (N_{loop}) among all threads ($N_{threads}$), in the loop where it is applied. The workload distribution within the threads can be changed by using the

OpenMP clause `schedule` and variable *chunk size*. The different OpenMP workload distributions used in this work are explained next.

Static: the load is distributed for each thread in fixed data blocks of roughly $N_{loop}/N_{threads}$. It is possible to choose the size of these blocks by changing the *chunk size* variable.

Dynamic: similar to the `static` one, with the main difference that, when a thread finishes to compute the work allocated to it and becomes idle, the system automatically assigns more work to this thread, until all the work finishes.

Guided: similar to the `dynamic` distribution, in the sense that, when a thread becomes idle, the system also allocates more load to this thread. The difference is that the size of the assigned loop subset starts with the value of $N_{loop}/N_{threads}$ and is decreased by the system until it reaches the chosen chunk size value. If the chunk size is omitted, the size of the final loop subset is 1.

Auto: is the automatic distribution provided by OpenMP. It delegates all the scheduling decisions to the compiler or run-time system.

IV. COUPLED SIMULATED ANNEALING

Coupled Simulated Annealing (CSA) [10] is a global optimization algorithm, based on the well-known simulated annealing (SA) algorithm [15]. The SA algorithm, also a global optimization method, is inspired by the thermodynamic annealing process, which consists of a heat treatment that alters a given material's physical properties. The SA algorithm is employed in minimization (or maximization) problems, where the goal is to obtain the minimum (or maximum) of a specific cost function, namely the *energy* of the annealing process. This work poses a minimization problem.

Briefly, the SA algorithm is divided into the generation of new solutions and the acceptance of these solutions. Algorithm parameters, known as generation and acceptance *temperatures*, control both these stages. New possible solutions, also known as probe solutions, are generated by a generation temperature function. If a probe solution yields a smaller value of the cost function, this solution is accepted as the new one with probability one; otherwise, this solution is only accepted as the new one with the probability given by a function of the acceptance temperature.

In its turn, the CSA algorithm consists of a set of parallel SA algorithms, known as SA optimizers. Each SA optimizer generates and evaluates a probing solution, updating its current state. The generation and acceptance temperatures are equal for all the different SA instances. The main differences between CSA and SA are that 1) for accepting solutions with higher cost function values, the CSA considers all current solutions, and 2) the acceptance criterion is based on the current solutions and a *coupling* term between these solutions. The coupling approach has shown to be capable of reducing the algorithm's sensitivity to initialization parameters and providing information that might steer the overall optimization process toward the global optimum.

The CSA algorithm employed in this work is the one described by Gonçalves-e-Silva *et al.* [16], which is implemented as follows. Let $a_i \in \Theta$ and $b_i \in \Omega$ be, respectively, the current and probe solutions of the i -th SA optimizer; with Θ and Ω being the set of current and probe solutions, respectively, and $i = 1, \dots, m$, where m is the number of elements in both Θ and Ω . At the k -th iteration of the CSA algorithm, the probe solutions are given by

$$b_i = a_i + \epsilon_i T_k^{\text{gen}}, \quad (5)$$

where T_k^{gen} is the generation temperature and ϵ_i is a random variable sampled from the Cauchy distribution

$$g(\epsilon, T) = \frac{T}{(\epsilon^2 + T^2)^{(D+1)/2}}, \quad (6)$$

where $T = T_k^{\text{gen}}$ and D is the dimension of the problem. The rule for updating T_k^{gen} is also a free choice of the specific CSA implementation. We followed the guidelines from [16] and used as update $T_{k+1}^{\text{gen}} = 0.99999T_k^{\text{gen}}$, with T_k^{gen} being updated to 99.999% of its previous value.

Each solution, a_i and b_i , has an associated energy (or cost) value, $E(a_i)$ and $E(b_i)$. The acceptance probability function is defined as:

$$A_\Theta = \frac{\exp\left(\frac{E(a_i) - \max(E(a_i))_{a_i \in \Theta}}{T_k^{\text{ac}}}\right)}{\gamma}, \quad (7)$$

where T_k^{ac} is the acceptance temperature and γ is the coupling term, given by:

$$\gamma = \sum_{\forall a \in \Theta} \exp\left(\frac{E(a) - \max(E(a_i))_{a_i \in \Theta}}{T_k^{\text{ac}}}\right). \quad (8)$$

If $E(b_i) > E(a_i)$, a_i assumes the value of b_i only if $A_\Theta < r$, where r is a random variable sampled from an uniform distribution in the interval $[0, 1]$. Otherwise, if $E(b_i) < E(a_i)$, a_i assumes the value of b_i with probability one.

As shown by Xavier-de-Souza *et al.* [10], the CSA performance is improved if the variance of A_Θ is kept close to its maximum value. This variance might be written as

$$\sigma^2 = \frac{1}{m} \sum_{\forall a \in \Theta} A_\Theta^2 - \frac{1}{m^2} \quad (9)$$

and lays in the interval

$$0 \leq \sigma^2 \leq \frac{m-1}{m^2}. \quad (10)$$

The controlling of this variance value can be accomplished by using the following rule to update the acceptance temperature:

$$T_{k+1}^{\text{ac}} = \begin{cases} T_k^{\text{ac}}(1 - \alpha), & \text{if } \sigma^2 < \sigma_D^2 \\ T_k^{\text{ac}}(1 + \alpha), & \text{if } \sigma^2 \geq \sigma_D^2 \end{cases} \quad (11)$$

where σ_D^2 is the desired variance, which should be kept as close as possible to $\frac{m-1}{m^2}$, and α is the acceptance temperature modification rate, usually a value within the interval $(0, 0.1]$.

The CSA algorithm is parameterized by setting the initial temperature values T_0^{gen} and T_0^{ac} ; the total number of iterations, N ; and the number of optimizers, m . Setting the number of CSA optimizers, m , is heavily discussed in [16], where it is shown that, for minimizing several functions, good choices for the number of optimizers lie in the range between $m = 4$ and $m = 10$. In this work, we chose to use $m = 4$ to minimize the auto-tuning overhead without compromising the CSA quality performance. Moreover, as shown in [10], the CSA algorithm is very robust to the initial values of the acceptance temperature, T_0^{ac} , presenting satisfactory results for a large range of T_0^{ac} values. Regarding the initialization of the generation temperature, T_0^{gen} , and the number of iterations, N , section VII-A provides a guideline for setting these parameters when auto-tuning RTM applications.

In the proposed auto-tuning approach, the CSA is employed to minimize the execution time of different loops in the RTM algorithm, parallelized with OpenMP, by properly choosing the optimal chunk size for OpenMP parallel loops. Therefore, the cost function $E(a_i)$ is related to the execution time of an `OpenMP parallel for` construction and the variable a_i is related to the chunk size, in the dynamic OpenMP distribution.

V. IMPLEMENTATION ASPECTS OF RTM

The RTM program developed to test the proposed auto-tuning is introduced by Assis *et al.* [17]. It is implemented in C, using a hybrid parallel approach. MPI distributes shots among the nodes of a distributed memory system while OpenMP schedules chunks of the 3D mesh representing the spatial domain to cores of a shared memory system.

The wave propagator of our RTM implementation solves the wave equation by the finite difference method (FDM), using an eighth order in space and second order in time stencil. We used non-reflecting boundary condition to absorb the energy at the boundaries as described in [12]. The absorbing boundary coefficients are computed by

$$\phi(i) = \begin{cases} \pi f_{\text{peak}} \Delta t \left(\frac{w_i}{w_b}\right)^2, & \text{on the borders,} \\ 0, & \text{otherwise,} \end{cases} \quad (12)$$

$$\phi(\mathbf{x}) = \phi(x_1) + \phi(x_2) + \phi(x_3), \quad (13)$$

$$\phi_1(\mathbf{x}) = \frac{1}{1 + \phi(\mathbf{x})}, \quad (14)$$

$$\phi_2(\mathbf{x}) = 1 - \phi(\mathbf{x}), \quad (15)$$

where f_{peak} is the peak frequency of the source, w_b is the thickness of the absorbing boundary, in number of grid points, and w_i ranges from 0 to w_b , indicating the shortest distance from a point (x_1, x_2, x_3) to the border's interior edge. Note that, away from the borders, $\phi_1(\mathbf{x}) = \phi_2(\mathbf{x}) = 1$ and we recover the usual FDM solution.

Using a finite difference second order scheme in time and applying the coefficients of (14) and (15) to (1) leads to

$$\begin{aligned}
 &u(\mathbf{x}, t + \Delta t) \\
 &= \phi_1(\mathbf{x}) \cdot \left\{ 2u(\mathbf{x}, t) - \phi_2(\mathbf{x}) \cdot u(\mathbf{x}, t - \Delta t) \right. \\
 &\quad \left. + (c(\mathbf{x})\Delta t)^2 \cdot \left[\frac{\partial^2 u(\mathbf{x})}{\partial x_1^2} + \frac{\partial^2 u(\mathbf{x})}{\partial x_2^2} + \frac{\partial^2 u(\mathbf{x})}{\partial x_3^2} - s(t) \right] \right\}, \tag{16}
 \end{aligned}$$

where the source $s(t)$ is modeled as a Ricker wavelet [18].

Our RTM code implements the optimal checkpointing strategy, described in [19], [20], to avoid the use of secondary storage and memory. Details of our RTM implementation are shown in Algorithm 1, which is further discussed in Section VI.

VI. CSA-BASED AUTO-TUNING

Katagiri *et al.* [21] defines three types of auto-tuning: i) install time, when the estimation procedure is affected by machine environments, ii) before execution invocation, when the estimation procedure is affected by user’s knowledge, input parameters or the number of processors, for example, and iii) run time, when the estimation is affected by other parameters generated in run time. In this paper, we propose a run-time auto-tuning for adequately determining the size of parallel loops subsets to be dynamically distributed among OpenMP threads.

The modeling of parallel applications is not trivial, especially if this application runs in different architectures. Variations of aspects such as memory size and hierarchy, number of cores, and input size may have very diverse effects in parallel software performance [22]. Since the relation between the chunk size of the parallel loops and the total execution time of a program is unknown, using a stochastic optimization method to find an optimal chunk size is an alternative. Estimating this relation is particularly challenging for the FDM because of the memory access pattern generated by the stencil computations. Using a multidimensional stencil means that the access to memory is non-linear at each wave propagation time step, making it more complex to avoid cache misses. For this reason, the proposed auto-tuning employs CSA to find the chunk size that minimizes the execution time.

For all the tests performed in this work, we adopted the following parameterization rules. For the initial acceptance temperature, we adopted one of the values suggested in [10], $T_0^{ac} = 0.9$. The number of iterations, N , and the number of SA optimizers, m , influence the convergence and exploration of the solution variables space. Large values of N and m might result in chunk size estimates closer to the global optimum, with the drawback of greater execution times. For setting the number of optimizers, we chose one of the values tested in [16], $m = 4$. The parameters σ_D^2 and α do not need to be configured, they can be kept fixed, in all simulations, according to [10], with $\sigma_D^2 = 0.99 \left(\frac{m-1}{m^2} \right)$ and $\alpha = 0.005$.

Algorithm 1 Reverse Time Migration With Auto-Tuning. n_s Is the Number of Time Steps. t_i Is the i -th Time Step in the RTM Algorithm

```

1: begin time measurement
2: distribute shots among nodes using MPI
3: read RTM parameters
4: initialize checkpointing variables
5: compute absorbing boundaries coefficients
6: initialize auto-tuning parameters
7: #OpenMP parallel section begin
8: for all shots location do
9:   read shot seismogram
10:  if it is the first shot then
11:    autotuning() (See Algorithm 2)
12:  end if
13:  for ( $t_i = 0$  to  $n_s - 1$ ) do
14:    #OpenMP parallel loop using the auto-tuned chunk
    size in a dynamic distribution
15:    for all grid points do
16:      compute the wavefield
17:    end for
18:    add the source wavelet
19:    if ( $t_i$  is a checkpoint) then
20:      save Checkpoint
21:    end if
22:  end for
23:  for ( $t_i = n_s - 1$  to  $0$ ) do
24:    #OpenMP parallel loop using the auto-tuned chunk
    size in a dynamic distribution
25:    for all grid points do
26:      compute the wavefield
27:    end for
28:    #OpenMP parallel loop using static distribution
29:    for all receivers location do
30:      inject observed data samples at time  $t_i$ 
31:    end for
32:    get forward wavefield at  $t_i$  from the checkpoints
    using the auto-tuned chunk size in a dynamic dis-
    tribution
33:    #OpenMP parallel loop using static distribution
34:    for all main grid points do
35:      perform image condition
36:    end for
37:  end for
38: end for
39: #OpenMP parallel section end
40: reduce all nodes migrated sections
41: finish time measurement

```

For the CSA-based auto-tuning algorithm, the only parameters that needed to be configured were the generation temperature, T_0^{gen} , and the number of iterations, N . In section VII-A we provide insights on how to configure these parameters. For all the tests that we performed, with different data sets, problem sizes, and machines, the CSA

algorithm has shown to be quite robust to the initialization of the parameters. We were able to achieve consistent results by using the same parameter values in all the tests.

Algorithm 1 presents the pseudo-code of the proposed implementation of the RTM. Regarding the application of CSA for auto-tuning the RTM, four parallel loops of the RTM have been enabled to use the dynamic scheduler with a chunk size that could be defined by the proposed auto-tuning method: i) the forward propagation of the source (Line 15); ii) the backward propagation of the observed data (Line 25); iii) the insertion of the receivers data (Line 29); and iv) the image condition (Line 34). The optimal checkpointing strategy (Line 32) recomputes some of the forward propagation time steps and can also have its chunk size defined by an auto-tuning method.

Since the propagation loops are essentially the same, we only apply the proposed auto-tuning for the forward propagation. The chunk size obtained is then used to the forward propagation, the backward propagation, and the checkpointing. On the other hand, the receivers' insertion and image condition loops are not auto-tuned. These loops have a significantly smaller dimension in comparison with the propagation loops. In our tests, these loops together spent less than 2% of the total execution time. Furthermore, they mostly perform linear access to memory, which is ideal for a static distribution. Its overhead may overcome the benefit of auto-tuning the receivers' data insertion and image condition loops. As shown in Algorithm 1 (Line 10), the proposed auto-tuning is performed only for the first shot. All the following shots use the same chunk size computed for the first shot.

Algorithm 2 details the implementation of the proposed auto-tuning. The initial set of solutions (chunk sizes) is randomly chosen in the interval $[50, N_{\text{loop}}/N_{\text{threads}}]$. We disregarded small chunk sizes because of the high overhead to schedule them dynamically. Chunk sizes greater than the chunk size of the standard `static` distribution ($N_{\text{loop}}/N_{\text{threads}}$) are also not taken into consideration because they would lead to the number of blocks to be less or equal than the number of threads and thus, the distribution would be forced to be static.

For each CSA iteration, each optimizer only measures the execution time of the first time step in the forward propagation, using its current chunk size (Lines 6 and 13). As shown in [23], the run time of the first time step can accurately represent the total propagation execution time. This first-time step is performed twice (Line 4) and only the elapsed time of the second repetition is registered (Lines 5 and 12) in order to avoid cache population effects. The CSA then uses those time measures as the cost function values and generates the next set of solutions (Line 17).

The overhead in the proposed auto-tuning algorithm is related to the number of iterations used in the CSA global optimization method, which repeats the first time step of the forward wave propagation for each optimization iteration. For instance, in a wave propagation of 1000 time

Algorithm 2 Proposed Auto-Tuning Method, Function `autotuning()` of Algorithm 1. t_i Is the i -th Time Step in the RTM Algorithm

```

1:  $t_i = 0$ 
2: for all (auto-tuning iterations) do
3:   for all optimizers do
4:     for ( $i = 1$  to 2) do
5:       if ( $i == 2$ ) then
6:         time measure begin
7:       end if
8:       #OpenMP parallel loop using the current chunk
       size in a dynamic distribution
9:       for all grid points do
10:        compute the wavefield
11:       end for
12:       if ( $i == 2$ ) then
13:        time measure end
14:       end if
15:       end for
16:     end for
17:     CSA generates a new solution for each optimizer from
       the time measures
18:   end for
19: return the solution with the lowest cost function

```

steps of both forward and backward wave equations and the use of checkpoints with a forward propagation recomputing rate of 2 to evaluate the image condition (4000 times steps in total), with 100 CSA iterations, the overhead would be 2.5% for one shot. Since the auto-tuning is performed only in the forward wave propagation of the first shot, for more shots, the overhead is shared among the time steps of all shots, being significantly reduced in these cases. We tested only instances of a small number of shots, but it is known that 3D seismic acquisitions performed in complex regions can present hundreds of thousands of shots.

VII. NUMERICAL EXPERIMENTS

Our experiments were conducted in five different computational environments, namely:

- Leuven: Single compute node hosting four sixteen-core AMD Opteron(TM) Processor 6376 at 2.3 GHz and 256 GB RAM. This equipment is located at the *Universidade Federal do Rio Grande do Norte (UFRN)*.
- NPAD: 68 compute nodes. Each compute node hosts two CPUs Intel Xeon Sixteen-Core E5-2698v3 at 2.3 GHz and 128 GB RAM DDR4 2133. It is equipped with a 60 TB Lustre parallel distributed file system. This equipment is located at the High-Performance Computing Center at UFRN (NPAD/UFRN).
- Yemoja: 856 compute nodes. Each compute node hosts two 10-core Intel Xeon E5-2690 Ivy Bridge v2 at 3 GHz. 200 nodes with 256 GB of RAM and 656 nodes with

128 GB RAM. It is equipped with an 850 TB Lustre parallel distributed file system. This equipment is located at the Manufacturing and Technology Integrated Campus of the National Service of Industrial Training (SENAI CIMATEC).

- Ogun: 48 compute nodes. Each compute node hosts two Intel Xeon Gold 6148 at 2.40 GHz and 192 GB RAM DDR4. It is equipped with a 340 TB Lustre parallel distributed file system. This equipment is also located at SENAI CIMATEC.
- SDumont: 504 compute nodes. Each compute node hosts two Intel Xeon E5-2695v2 Ivy Bridge at 2.40 GHz and 64 GB RAM DDR4. It is equipped with a 640 TB Lustre parallel distributed file system. This equipment is located at the National Laboratory of Scientific Computing (LNCC).

The compiler used was *gcc*, using its default run-time system (*libgomp*) and the optimization flag *-O3*. We used the most recent version of *gcc* available for each machine, namely 8.1 for Leuven, 7.3 for NPAD, 8.2 for Yemoja, 7.3 for Ogun and 8.3 for SDumont. No controls were used to bind threads to processors. Neither the OpenMP [24] nor the *libgomp* [25] specifications present how the `auto` distribution maps the iterations of the parallel loops to the threads. However, it is possible to notice from *libgomp*'s source code [26] that `auto` distribution maps to the `static` distribution with a chunk size of roughly $N_{loop}/N_{threads}$.

In order to validate the 3D acoustic wave propagator used in our RTM program, we compared a seismic trace computed by our program with the 3D acoustic analytical solution, computed based on [27], in a homogeneous velocity model. The source was a Ricker wavelet with a peak frequency of 20 Hz. The distance between source and receiver was 200 m. The medium had a constant velocity of 2000 m/s. In this experiment, our propagator provided a very accurate approximation to the 3D waveform analytical solution with a mean squared error of 6×10^{-14} .

For all the following tests, $f_{peak} = 20$ Hz, the time sampling is 1 ms, the number of time steps is 3501, the spatial resolutions are $\Delta x_1 = \Delta x_2 = \Delta x_3 = 10$ m and the absorbing border thickness is 50 points in all directions of the 3D mesh. We built $c(\mathbf{x})$ by using a two layers model with a flat interface positioned at the center of the vertical dimension, where the top and bottom layers have velocities of 1400 m/s and 2000 m/s, respectively. The number of buffers (n_b) and checkpoints (n_c) depends on the size of the input, as shown in Table 1. The numbers of buffers were chosen in order to use up to 128 GB of RAM. We chose the number of checkpoints optimally according to [20].

The execution time measures were taken by using `MPI_Wtime()` function. The initial time is measured at the beginning of the algorithm just after launching the processes through `MPI` (Algorithm 1 Line 1). The final time is measured at the end of the algorithm just before finalizing the processes (Algorithm 1 Line 41).

TABLE 1. Number of buffers and checkpoints used in the experiments as function of the input size. The input size does not include the absorbing border. n_1, n_2 and n_3 are the number of samples for the spatial dimensions x_1, x_2 and x_3 , being the latter the vertical dimension. $n_2 = n_3 = 401$.

Input size ($n_1 \times n_2 \times n_3$)	n_b	n_c
$n_1 = 201$	170	3330
$n_1 = 401$	100	3400
$n_1 = 801$	56	1848

A. CSA AUTO-TUNING PARAMETERIZATION

In this subsection, we explain the methodology employed in CSA parameterization for the auto-tuning of RTM. We provide insights for configuring the CSA parameters in the RTM application. As previously discussed in sections IV and VI, for the proposed CSA-based auto-tuning of the RTM application, the input parameters that we chose to set were only the generation temperature, T_0^{gen} , and number of iterations, N . We fixed the acceptance temperature in $T_0^{ac} = 0.9$ and the number of optimizers in $m = 4$, as suggested in [10] and [16].

In the following experiment, we show that the configuration of the proposed auto-tuning method is robust to the chosen parameter values and computational architecture, by setting the values of T_0^{gen} and N . In contrast, section VII-B shows that the chunk size configuration is quite sensitive to its value and likely to change according to the architecture. In that sense, our experiments lead to the conclusion that it may be advantageous to set the auto-tuning parameters once, in a single architecture, instead of performing trial-and-error tests when moving across architectures.

To provide a methodology for choosing the auto-tuning parameters, we propose a parameterization experiment in a specific computational environment. The assumption here is that the parameters obtained in this particular architecture would provide consistent results for other architectures. This parameterization experiment yields the values of T_0^{gen} and N used in all the experiments of auto-tuning the RTM application. Note that a possible extension of the parameterization algorithm experiment is quite straightforward, for the use of the proposed auto-tuning in different applications. The parameterization experiment also shows that the auto-tuning algorithm is reasonably robust to the choice of T_0^{gen} and N values.

The computational system chosen for the parameterization experiment is composed of two additional compute nodes located in NPAD supercomputer, with the following characteristics:

- One CPU 7250 Intel Xeon Phi with 68 cores 1.4 GHz and 128 GB DDR4 (8 × 16 GB).

In this test, the velocity model characteristics were as described in the previous section and its dimensions were $201 \times 401 \times 401$. We tested combinations of $N = 40, 80, \text{ and } 160$ and $T_0^{gen} = 1, 10, 100, \text{ and } 1000$. For each

test, we executed 10 repetitions of the RTM of one seismic shot. We measured the execution time of the entire seismic shot, including the auto-tuning algorithm.

Fig. 1 shows the obtained results. Most of the shorter execution times occur for the smaller numbers of iterations ($N = 40$, and 80). This fact is most likely because the overhead of the CSA-based auto-tuning is smaller in these cases. It is also possible to observe that in the cases of $N = 40$, and 80 the variance of the executions was quite significant, with the minimum sometimes overlapping among the tests with different parameter combinations. We chose to use $N = 40$ as the initial value for the number of iterations since it has a shorter run time. Moreover, for this case, the average run time was quite close for both $T_0^{\text{gen}} = 1, 10$, and 100 . We chose to use $T_0^{\text{gen}} = 100$ since it presented the shorter run time among all the executions.




FIGURE 1. One-shot RTM run time using the proposed auto-tuning varying the number of CSA iterations ($N = 40, 80$, and 160), and the initial generation temperature ($T_0^{\text{gen}} = 1, 10, 100$ and 1000). The marks on x-axis follows the *XiGY* format meaning $N = X$ and $T_0^{\text{gen}} = Y$. The size of the velocity model was $(n_1 \times n_2 \times n_3) = 201 \times 401 \times 401$. These measurements were taken at a CPU 7250 Intel Xeon Phi machine with 68 cores at 1.4 GHz.

Note that our parametrization study ran 120 instances of the RTM, i. e., we have tested 12 different parameter sets 10 times each. Also, the number of sets of parameters tested would be the same for different input sizes. Putting it into perspective, the number of RTM executions needed in an exhaustive search for the optimal chunk size depends on the size of the velocity model. For the model size and number of threads used in our experiments, $201 \times 401 \times 401$, and 32 , respectively, an exhaustive search would run approximately one million RTM instances.

The CSA configuration parameters used in the remaining experiments are summarized in Table 2. We show in Section VII-B that the proposed auto-tuning using these parameters was able to outperform the default OpenMP schedulers despite running in different computational architectures and for different input sizes.

TABLE 2. CSA parameters used in the numerical experiments.

T_0^{gen}	T_0^{ac}	N	m
100	0.9	40	4

B. PERFORMANCE ANALYSIS

In this subsection, we present three different types of performance experiments. The first type presents the speedup of the 3D RTM when using the proposed method in comparison with OpenMP schedulers *auto*, *static*, and *guided* for different computational architectures and sizes of the velocity model. The second type shows the absolute number of cache misses of the 3D RTM when using the proposed method and the OpenMP schedulers *auto*, *static*, and *guided*. Finally, the third type displays the proposed method's overhead measurements in the 3D RTM for different computational architectures and sizes of the velocity model.

1) SPEEDUP ANALYSIS

Following, we present two sets of experiments regarding the speedups of the proposed auto-tuning method.

The objective of the following set of experiments is to measure the performance of the proposed method in comparison with the OpenMP schedulers, *auto*, *static*, and *guided*, with their default chunk size, i. e., when the chunk size is not explicitly specified. All the following measured run times include the overhead of the proposed method.

The results presented in Table 3 display the performance of the proposed auto-tuning in comparison with the default OpenMP schedulers in the set of five computational architectures described in Section VII. For all the experiments, the dimension of the velocity model was $(n_1 \times n_2 \times n_3) = 401 \times 401 \times 401$. The proposed method presented a superior performance in all 75 tested scenarios but two, namely, when running the one-shot and the two-shots RTM at Yemoja using the OpenMP *guided* scheduler. Even in these cases, the performance loss was inferior to 2%. As Yemoja has the smallest number of cores per node among the machines tested, the maximum chunk size ($N_{\text{loop}}/N_{\text{threads}}$) is the largest of all scenarios, making the search domain wider for the proposed auto-tuning. On the other hand, the overhead of the OpenMP *guided* scheduler may be reduced as it manages a smaller number of processing units. The proposed method outperformed the default OpenMP *static* and *auto* schedulers in all tested scenarios.

Fig. 2 details the results shown in Table 3 for a 16-shot RTM. The proposed method outperforms the three OpenMP schedulers tested regardless of which machine was employed.

The following set of experiments, presented at Table 4, displays the performance of the proposed auto-tuning in comparison with the default OpenMP schedulers, *static*, *auto*, and *guided*, when varying the size of the problem both in means of the dimension of the velocity model and the number of shots. This set of tests was performed at Ogun.

TABLE 3. RTM speedup when using the proposed auto-tuning, compared with the OpenMP schedulers *auto*, *static*, and *guided* in five machines (Leuven, Yemoja, SDumont, NPAD, and Ogun), for 1, 2, 4, 8 and 16 seismic shots. For the OpenMP schedulers, the chunk size was not explicitly specified. For these experiments, the size of the velocity model was $(n_1 \times n_2 \times n_3) = 401 \times 401 \times 401$. Each point is a median of at least five executions.

		Computer environments														
		Leuven			Yemoja			SDumont			NPAD			Ogun		
		static	auto	guided	static	auto	guided	static	auto	guided	static	auto	guided	static	auto	guided
Number of shots	1	5.5%	7.3%	9.0%	8.6%	10.1%	-1.3%	16.5%	16.4%	2.1%	23.1%	22.7%	6.1%	15.3%	15.3%	6.5%
	2	4.5%	5.1%	7.1%	7.2%	6.1%	-1.8%	9.9%	10.2%	1.6%	18.8%	18.4%	6.3%	14.7%	14.7%	7.9%
	4	2.6%	4.0%	8.9%	4.4%	5.6%	1.2%	6.2%	6.9%	1.4%	16.0%	16.3%	8.4%	13.1%	12.9%	8.9%
	8	2.6%	2.9%	6.5%	10.4%	8.5%	3.1%	4.2%	4.0%	0.8%	16.6%	14.6%	9.2%	13.3%	13.5%	9.9%
	16	2.7%	2.9%	4.1%	4.5%	4.5%	4.8%	2.7%	2.9%	1.1%	14.2%	15.0%	10.0%	13.1%	13.3%	9.5%

TABLE 4. RTM speedup when using the proposed auto-tuning, compared with the OpenMP schedulers *auto*, *static*, and *guided* for three input sizes, $(n_1 \times n_2 \times n_3)$ equal to $201 \times 401 \times 401$, $401 \times 401 \times 401$, and $801 \times 401 \times 401$, and for 1, 2, 4, 8 and 16 seismic shots. For the OpenMP schedulers, the chunk size was not explicitly specified. These tests were performed at Ogun. Each point is a median of at least five executions.

		$n_1 = 201$			$n_1 = 401$			$n_1 = 801$		
		static	auto	guided	static	auto	guided	static	auto	guided
Number of shots	1	10.6%	9.9%	6.0%	15.3%	15.3%	6.5%	33.1%	33.0%	14.3%
	2	12.1%	13.0%	8.1%	14.7%	14.7%	7.9%	27.6%	27.4%	13.4%
	4	14.4%	14.1%	10.0%	13.1%	12.9%	8.9%	21.3%	21.2%	11.6%
	8	14.1%	14.2%	9.2%	13.3%	13.5%	9.9%	19.5%	20.2%	11.8%
	16	15.1%	14.9%	10.8%	13.1%	13.3%	9.5%	16.9%	17.2%	11.9%

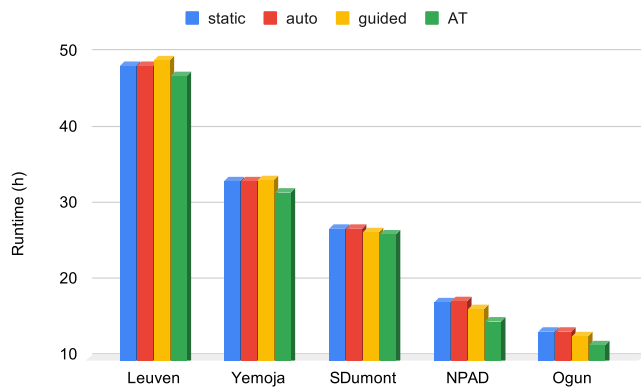


FIGURE 2. Execution time of a 16-shot RTM in a single node using the proposed auto-tuning (AT) and the OpenMP schedulers *auto*, *static*, and *guided*, in five machines (Leuven, Yemoja, SDumont, NPAD, and Ogun). For the OpenMP schedulers, the chunk size was not explicitly specified. For these experiments, the size of the velocity model was $(n_1 \times n_2 \times n_3) = 401 \times 401 \times 401$. Each measurement presented is a median of at least five executions.

The proposed method presents its best performance for larger sizes of the velocity model. Also, for the one-shot RTM, our method’s speedup increases with the size of the model.

The larger the input size, the bigger the chunks of the *static* scheduler and the initial chunk of the *guided* scheduler. For instance, when running over Ogun, the chunk sizes of the *static* distribution were approximately 1.9, 3.1, and 5.7 millions of loop iterations for the input sizes of $201 \times 401 \times 401$, $401 \times 401 \times 401$, and $801 \times 401 \times 401$ respectively. By working with larger chunks,




FIGURE 3. 16-shot RTM run time for the proposed auto-tuning, compared with the *auto* and the *static* scheduling types, for three input sizes, $(n_1 \times n_2 \times n_3)$ equal to $201 \times 401 \times 401$, $401 \times 401 \times 401$, and $801 \times 401 \times 401$. These tests were performed at Ogun. Each point is a median of at least five executions.

the data locality of the *static* distribution decreases, which explains its performance decrease for larger input sizes. On the other hand, the median of the chunk sizes chosen by the proposed auto-tuning method was 30.3, 122, and 232.5 thousands of loop iterations for the same input sizes, respectively. By processing suitably smaller chunks and distributing them dynamically, the likelihood that all threads work in a contiguous memory range increases, and so does the data locality in the proposed method. This temporal data locality increases the reuse of the data in cache memory.

Fig. 3 details the results shown in Table 4 for a 16-shot RTM. The proposed method outperformed the three OpenMP

schedulers tested regardless of the dimension of the velocity model. Moreover, the employment of the proposed method leads to higher performance increase as the dimension of the velocity model increases.

Regarding the number of shots, for most of the cases in Tables 3 and 4, the performance of the proposed method decreases as the number of shots increases. That indicates that the optimized chunk size for the first shot may not be optimal for the next ones and, thus, the proposed auto-tuning method could be applied to each shot in order to increase its efficiency. However, this topic is a matter of further investigation.

2) CACHE MISSES ANALYSIS

The following set of performance analysis experiments aimed to answer why optimizing the chunk size through the proposed method increases the parallel efficiency of a single shot RTM. For that, we employed the HPCToolkit performance tools [28] to measure the amount of cache misses for the three levels of cache available in NPAD, namely, L1, L2, and L3.

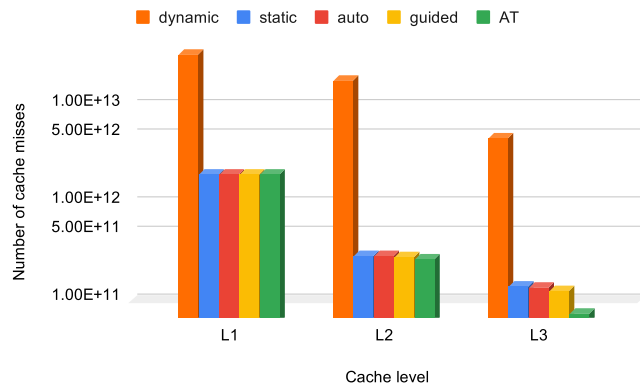


FIGURE 4. Absolute number of cache misses measured by HPCToolkit for the execution of a single shot RTM using the proposed auto-tuning (AT) and the OpenMP schedulers, namely, *auto*, *static*, *dynamic*, and *guided*. For the OpenMP schedulers, the chunk size was not explicitly specified. The size of the velocity model was $(n_1 \times n_2 \times n_3) = 401 \times 401 \times 401$. These measurements were taken at NPAD.

Fig. 4 shows that using the OpenMP *dynamic* scheduler without specifying the chunk size leads to a significantly higher number of cache misses when compared to the other schedulers. This result, as well as the experiments presented by [23], [29], [30], shows that using the OpenMP *dynamic* scheduler with the default unitary chunk size and other very small chunk sizes for the RTM leads to loss of performance. The reasons for that are the high number of cache misses due to false sharing, and the overhead to manage the distribution of tasks. Given that, we decided to restrict the chunk size search domain of our method to start at 50, as mentioned before. Also, for the same reason, we omitted the results of the *dynamic* scheduler with the default chunk size from the other tests.

The number of cache misses for the *auto*, *static*, and *guided* OpenMP schedulers with the default chunk size were nearly the same for the experiment shown in Fig. 4

considering cache levels L1 and L2. For cache level L3, the *guided* scheduler has 7.46% and 8.15% less cache misses compared to *auto* and *static* schedulers respectively. The proposed method has 1%, 7%, and 43% less cache misses compared to the *guided* scheduler for cache levels L1, L2, and L3 respectively.

3) OVERHEAD ANALYSIS

The following set of performance analysis experiments aimed to measure the overhead of the proposed auto-tuning in an RTM. For that, we employed the MPI function `MPI_Wtime` before and after Line 11 of Algorithm 1. In these experiments, we verified how the proposed method's overhead behaves when varying the size of the problem. For an RTM, the size of the problem can vary by either changing the size of the velocity model or the number of shots.

Table 5 presents the overhead of the proposed method when varying the size of the velocity model in two computer environments, namely, NPAD and Ogun. For all experiments of this set, the overhead was inferior to 2%. Table 5 shows that, for this experiment, the overhead does not change significantly when the size of the velocity model changes. As the proposed method performs its iterations using the same size of the model used by the RTM iterations, the size of the model should not affect the technique's overhead.

TABLE 5. Proposed auto-tuning overhead running over a single shot of the RTM in two different computer environments (NPAD, and Ogun), and three input sizes, $(n_1 \times n_2 \times n_3)$ equal to $201 \times 401 \times 401$, $401 \times 401 \times 401$, and $801 \times 401 \times 401$. Each value is the median of at least five executions.

	$n_1 = 201$	$n_1 = 401$	$n_1 = 801$
NPAD	1.33%	1.24%	1.28%
Ogun	0.78%	0.84%	0.86%

Table 6 shows the overhead of the proposed method when varying the number of shots in five computer environments, namely, Leuven, Yemoja, SDumont, NPAD, and Ogun. Again, for all experiments of this set, the overhead was inferior to 2%. The proposed method's overhead decreases as the number of shots increases because the auto-tuning is executed only for each node's first shot. The optimized chunk size is employed for all shots processed in the same node.

VIII. RELATED WORK

Several works have introduced auto-tuning techniques on multicore systems. Katagiri *et al.* [31], [32] presented ppOpen-AT, a framework for code optimization guided by directives. Sena *et al.* [33] proposed a method to determine a near-optimal workload chunk size by testing a set of possible values in a few time steps of an RTM and choosing the one with the shortest execution time. Andreolli *et al.* [34], [35] introduced an approach to tune seismic applications automatically by compiling and running each set of parameters chosen by a genetic algorithm, including chunk size and compilation flags. Kamil *et al.* [36] presented a framework to

TABLE 6. Proposed auto-tuning overhead running over an RTM for five amounts of shots (1, 2, 4, 8, and 16) and five computer environments (Leuven, Yemoja, SDumont, NPAD, and Ogun). The size of the velocity model employed was $(n_1 \times n_2 \times n_3) = (401 \times 401 \times 401)$. Each cell is the median of at least five executions.

		Number of shots				
		1	2	4	8	16
Computer environments	Leuven	1.96%	0.99%	0.50%	0.27%	0.16%
	Yemoja	0.76%	0.38%	0.20%	0.10%	0.05%
	SDumont	0.72%	0.40%	0.21%	0.11%	0.06%
	NPAD	1.24%	0.70%	0.37%	0.19%	0.10%
	Ogun	0.84%	0.46%	0.24%	0.12%	0.06%

generate auto-tuned C, Fortran, and CUDA codes for stencil applications specified in sequential Fortran 95. In common, these works [31]–[36] perform the auto-tuning before the execution time. By doing so, these methods may find a system status different from the tuning time as aspects such as memory availability may change over time. Conversely, the proposed method performs auto-tuning at run time, enabling it to capture more realistic parameters.

Many authors have proposed run-time auto-tuning approaches. Padoin *et al.* [37], [38] employed load balancing techniques along with processors frequency control tools to improve the energy efficiency of imbalanced parallel applications in multicore systems. Both works use processors frequency scaling techniques to slow down less loaded cores to save energy. Tchiboukdjian *et al.* [39] introduced a scheduler for applications with linear access to shared memory. They aim to improve locality by guaranteeing that all data in the cache are used before being replaced. Olivier *et al.* [40], [41] proposed an OpenMP hierarchical task scheduler for multicore systems using a work-stealing strategy [42]. Different from [37]–[41], the proposed method focuses on tuning the scheduling of the OpenMP directive for parallel loops automatically. According to Diaz *et al.* [43], OpenMP is a predominant approach in shared memory architectures both in industry and academia, which enforces the relevance of our work.

Some works also introduced auto-tuning techniques which can schedule parallel loops using OpenMP. Kale and Gropp [44] employed a strategy combining static and dynamic scheduling to tune implementations of regular meshes applications in clusters of symmetric multiprocessing (SMP) machines and by Donfack *et al.* [45] to tune highly-optimized dense matrix factorization methods. Kale *et al.* [46] proposed improving a combined static and dynamic scheduling on SMP machines by enforcing spatial locality. Bak *et al.* [47] introduce a load-balancing method based on the integration of Charm++ and OpenMP to dynamically distribute user-created tasks. Differently, we employ CSA to automatically determine the size of the chunks of parallel loop iterations. The relation between chunk size and total execution time in geophysical methods is unknown but has

been shown to have many local optima [23], [30]. By using a global optimization strategy, our method aims to increase the probability of finding the global optimum.

Tiwari and Hollingsworth [48] proposed an online auto-tuning method to tune programmer-defined parameters. For that, a set of nodes of a distributed system generate, compile, and run code versions using different sets of parameters. The optimization of the sets of parameters is defined via simplex. Differently from our work, Tiwari and Hollingsworth’s method relies on a distributed system. Although it can be performed in a single node, the results presented by Tiwari and Hollingsworth show that the simplex did not converge and increased the time of execution when employing a single node. Also, the overhead of Tiwari and Hollingsworth’s method increases as the size of the problem decreases while our results show that our method’s overhead is not correlated to the size of the problem.

We have shown before that CSA [10] is a promising method to obtain a near-optimal load balance for a 3D FDM. Experiments in [23], [30] showed that the optimal load balancing for shared memory environments depends on the hardware and software employed. In this paper, we use the auto-tuning method proposed in our previous works [23], [30] to the RTM and provide overhead, parallel performance, and CSA parameters analysis.

IX. CONCLUSION

We have proposed a CSA-based auto-tuning strategy for seeking the optimal chunk size of OpenMP’s dynamic loop scheduler to reduce the execution time of a 3D reverse time migration algorithm. The proposed approach is designed to work wherever the code is executed, being robust for changes in computational environment parameters, such as the number of threads, processors, memory hierarchy, and compiler.

We showed that the proposed auto-tuning reduces the number of cache misses compared to the default OpenMP schedulers. Also, for all experiments performed in this paper, the overhead of the proposed method was inferior to 2%.

Experiments running an RTM showed that, in most cases, the proposed auto-tuning method outperforms the OpenMP *static*, *auto*, and *guided* schedulers in different computational resources for different amounts of seismic shots. For this set of experiments, the proposed auto-tuning method reached speedups up to 23.1%.

Another set of tests with RTM presented the proposed method’s performance in comparison with the OpenMP *static*, *auto*, and *guided* schedulers for a varying input size and a different number of seismic shots. For this set of experiments, the proposed auto-tuning outperformed the OpenMP schedulers, reaching up to 33% speedup. The proposed method performed better for larger inputs, which improves the parallel scalability of 3D RTM.

In summary, the proposed method has shown to be robust and scalable for the 3D RTM. As future work, we are interested in testing the proposed method’s ability to improve the performance of similar wave-based algorithms,

such as full-waveform inversion (FWI) and other iterative applications.

ACKNOWLEDGMENT

The authors acknowledge the National Laboratory for Scientific Computing (LNCC/MCTI, Brazil), the High-Performance Computing Center at UFRN (NPAD/UFRN), the Center of Technology of the UFRN (CT/UFRN), and the Manufacturing and Technology Integrated Campus of the National Service of Industrial Training (SENAI CIMATEC) for providing HPC resources of the SDumont, NPAD, Leuven, Yemoja, and Ogun supercomputers, which have contributed to the research results reported within this article. Finally, the authors thank Jorge Lopez from Shell and the anonymous reviewers for providing essential comments on this article.

REFERENCES

- [1] P. Kearey, M. Brooks, and I. Hill, *An Introduction to Geophysical Exploration*, 3rd ed. Malden, MA, USA: Blackwell Science, 2002.
- [2] E. Baysal, D. D. Kosloff, and J. W. Sherwood, "Reverse time migration," *Geophysics*, vol. 48, no. 11, pp. 1514–1524, 1983.
- [3] D. D. Kosloff and E. Baysal, "Migration with the full acoustic wave equation," *Geophysics*, vol. 48, no. 6, pp. 677–687, Jun. 1983.
- [4] J.-H. Zhang, S.-Q. Wang, and Z.-X. Yao, "Accelerating 3D Fourier migration with graphics processing units," *Geophysics*, vol. 74, no. 6, pp. WCA129–WCA139, Nov. 2009.
- [5] M. Araya-Polo, F. Rubio, R. de la Cruz, M. Hanzich, J. M. Cela, and D. P. Scarpazza, "3D seismic imaging through reverse-time migration on homogeneous and heterogeneous multi-core processors," *Sci. Program.*, vol. 17, nos. 1–2, pp. 185–198, 2009.
- [6] D. A. Nunes-Do-Rosário, S. Xavier-de-Souza, R. C. Maciel, and J. C. Costa, "Parallel scalability of a fine-grain prestack reverse time migration algorithm," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 12, pp. 2433–2437, Dec. 2015.
- [7] K. Naono, K. Teranishi, J. Cavazos, and R. Suda, *Software Automatic Tuning: From Concepts to State-of-the-Art Results*. New York, NY, USA: Springer, 2010.
- [8] A. M. C. Demetrios, D. De Sensi, A. F. Lorenzon, K. Georgiou, J. Nunez-Yanez, K. Eder, and S. Xavier-de-Souza, "Performance and energy trade-offs for parallel applications on heterogeneous multi-processing systems," *Energies*, vol. 13, no. 9, p. 2409, May 2020, doi: 10.3390/en13092409.
- [9] L. Dagum and R. Menon, "OpenMP: An industry standard API for shared-memory programming," *IEEE Comput. Sci. Eng.*, vol. 5, no. 1, pp. 46–55, Jan. 1998.
- [10] S. Xavier-de-Souza, J. A. Suykens, J. Vandewalle, and D. Bollé, "Coupled simulated annealing," *IEEE Trans. Syst. Man, Cybern. B, Cybern.*, vol. 40, no. 2, pp. 320–335, Apr. 2009.
- [11] J. M. Carcione, G. C. Herman, and A. P. E. ten Kroode, "Seismic modeling," *Geophysics*, vol. 67, no. 4, pp. 1304–1325, 2002.
- [12] C. Cerjan, D. Kosloff, R. Kosloff, and M. Reshef, "A nonreflecting boundary condition for discrete acoustic and elastic wave equations," *Geophysics*, vol. 50, no. 4, p. 705, 1985.
- [13] Ö. Yilmaz, *Seismic Data Analysis: Processing, Inversion, and Interpretation of Seismic Data* (Investigations in Geophysics), no. 2. Tulsa, OK, USA: Society of Exploration Geophysicists, 2001.
- [14] L. Clarke, I. Glendinning, and R. Hempel, "The MPI message passing interface standard," in *Programming Environments for Massively Parallel Distributed Systems*, K. M. Decker and R. M. Rehmann, Eds. Basel, Switzerland: Birkhäuser, 1994, pp. 213–218.
- [15] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [16] K. Gonçalves-e-Silva, D. Aloise, and S. Xavier-de-Souza, "Parallel synchronous and asynchronous coupled simulated annealing," *J. Supercomput.*, vol. 74, no. 6, pp. 2841–2869, Jun. 2018.
- [17] I. A. S. Assis, A. D. S. Oliveira, T. Barros, I. M. Sardina, C. P. Bianchini, and S. X. De-Souza, "Distributed-memory load balancing with cyclic token-based work-stealing applied to reverse time migration," *IEEE Access*, vol. 7, pp. 128419–128430, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8822671/>
- [18] Y. Wang, "Frequencies of the ricker wavelet," *Geophysics*, vol. 80, no. 2, pp. A31–A37, Mar. 2015.
- [19] W. W. Symes, "Reverse time migration with optimal checkpointing," *Geophysics*, vol. 72, no. 5, pp. SM213–SM221, Sep. 2007.
- [20] A. Griewank and A. Walther, "Algorithm 799: Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation," *ACM Trans. Math. Softw.*, vol. 26, no. 1, pp. 19–45, Mar. 2000.
- [21] T. Katagiri, K. Kise, H. Honda, and T. Yuba, "Fiber: A generalized framework for auto-tuning software," in *High Performance Computing*, A. Veidenbaum, K. Joe, H. Amano, and H. Aiso, Eds. Berlin, Germany: Springer, 2003, pp. 146–159.
- [22] A. F. A. Furtunato, K. Georgiou, K. Eder, and S. Xavier-De-Souza, "When parallel speedups hit the memory wall," *IEEE Access*, vol. 8, pp. 79225–79238, 2020.
- [23] T. Barros, J. B. Fernandes, I. A. Souza-de-Assis, and S. Xavier-de-Souza, "Auto-tuning of 3D acoustic wave propagation in shared memory environments," in *Proc. 1st EAGE Workshop High Perform. Comput. Upstream Latin Amer.* Boadilla del Monte, Spain: Santander, 2018, pp. 1–5. [Online]. Available: <http://www.earthdoc.org/publication/publicationdetails/?publication=94579>
- [24] (2015). *OpenMP Architecture Review Board, OpenMP Application Program Interface Version 4.5*. [Online]. Available: <https://www.openmp.org/wp-content/uploads/openmp-4.5.pdf>
- [25] GCC Team. (Jun. 2019). *GNU Libgomp*. [Online]. Available: <https://gcc.gnu.org/onlinedocs/libgomp/>
- [26] GCC Team. (Jun. 2020). *GNU Compiler Collection (GCC)*. [Online]. Available: <https://github.com/gcc-mirror/gcc/blob/master/libgomp/loop.c>
- [27] A. T. de Hoop, "A modification of Cagniard's method for solving seismic pulse problems," *Appl. Sci. Res. B*, vol. 8, no. 1, pp. 349–356, Dec. 1960.
- [28] L. Adhianto, S. Banerjee, M. Fagan, M. Krentel, G. Marin, J. Mellor-Crummey, and N. R. Tallent, "HPCTOOLKIT: Tools for performance analysis of optimized parallel programs," *Concurrency Comput., Pract. Exper.*, vol. 22, no. 6, pp. 685–701, 2010.
- [29] I. A. Souza-de-Assis, D. A. Nunes-do-Rosário, and R. C. S. M. Xavier-de-Souza, "Um Algoritmo Paralelo Eficiente de Propagação de Onda Acústica 3D," in *Proc. Rio Oil Gas Expo Conf.*, Rio de Janeiro, Brazil, 2014, pp. 1–8.
- [30] J. B. Fernandes, Í. A. Souza-de-Assis, T. Barros, and S. Xavier-de-Souza, "Automatic scheduler for 3D seismic modeling by finite differences," in *Proc. Rio Oil Gas Expo Conf.*, Rio de Janeiro, Brazil, 2018, pp. 1–10. [Online]. Available: https://stt.ibp.org.br/eventos/2018/rioil2018/pdfs/Riooil2018_1901_201806151345rioog_end_paper.pdf
- [31] T. Katagiri, S. Ohshima, and M. Matsumoto, "Auto-tuning of computation kernels from an FDM code with ppOpen-AT," in *Proc. IEEE 8th Int. Symp. Embedded Multicore/Manycore SoCs*, Sep. 2014, pp. 91–98.
- [32] T. Katagiri, S. Ohshima, and M. Matsumoto, "Directive-based auto-tuning for the finite difference method on the Xeon Phi," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. Workshop*, May 2015, pp. 1221–1230.
- [33] A. C. Sena, A. P. Nascimento, C. Boeres, V. Rebello, and A. Bulcao, "An approach to optimise the execution of RTM algorithm in multicore machines," in *Proc. IEEE 7th Int. Conf. eScience*, Dec. 2011, pp. 403–410.
- [34] C. Andreolli, P. Thierry, L. Borges, C. Yount, and G. Skinner, "Genetic algorithm based auto-tuning of seismic applications on multi and manycore computers," in *Proc. EAGE Workshop High Perform. Comput. Upstream*, Sep. 2014, p. 426.
- [35] C. Andreolli, P. Thierry, L. Borges, G. Skinner, and C. Yount, "Characterization and optimization methodology applied to stencil computations," in *High Performance Parallelism Pearls*, J. Jeffers and J. Reinders, Eds. Boston, MA, USA: Elsevier, 2015, ch. 23, pp. 377–396. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780128021187000236>
- [36] S. Kamil, C. Chan, L. Oliker, J. Shalf, and S. Williams, "An auto-tuning framework for parallel multicore stencil computations," in *Proc. IEEE Int. Symp. Parallel Distrib. Process. (IPDPS)*, 2010, pp. 1–12.
- [37] E. L. Padoin, M. Castro, L. L. Pilla, P. O. A. Navaux, and J.-F. Mehaut, "Saving energy by exploiting residual imbalances on iterative applications," in *Proc. 21st Int. Conf. High Perform. Comput. (HiPC)*, Dec. 2014, pp. 1–10.

[38] E. L. Padoin, L. L. Pilla, M. Castro, P. O. Navaux, and J. F. Méhaut, "Exploration of load balancing thresholds to save energy on iterative applications," in *Proc. Latin Amer. High Perform. Comput. Conf.*, in Communications in Computer and Information Science, 2017, pp 76–88.

[39] M. Tchiboukdjian, V. Danjean, T. Gautier, F. Le Mentec, and B. Raffin, "A work stealing scheduler for parallel loops on shared cache multicores," in *Proc. Eur. Conf. Parallel Process.*, in Lecture Notes in Computer Science: Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics, 2011, pp. 99–107.

[40] S. L. Olivier, A. K. Porterfield, K. B. Wheeler, and J. F. Prins, "Scheduling task parallelism on multi-socket multicore systems," in *Proc. 1st Int. Workshop Runtime Oper. Syst. Supercomput. (ROSS)*, 2011, pp. 49–56.

[41] S. L. Olivier, A. K. Porterfield, K. B. Wheeler, M. Spiegel, and J. F. Prins, "OpenMP task scheduling strategies for multicore NUMA systems," *Int. J. High Perform. Comput. Appl.*, vol. 26, no. 2, pp. 110–124, 2012, doi: 10.1177/1094342011434065.

[42] R. D. Blumofe and C. E. Leiserson, "Scheduling multithreaded computations by work stealing," *J. ACM*, vol. 46, no. 5, pp. 720–748, Sep. 1999, doi: 10.1145/324133.324234.

[43] J. Diaz, C. Muñoz-Caro, and A. Niño, "A survey of parallel programming models and tools in the multi and many-core era," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 8, pp. 1369–1386, Aug. 2012.

[44] V. Kale and W. Gropp, "Load balancing for regular meshes on SMPs with MPI," in *Recent Advances in the Message Passing Interface*, R. Keller, E. Gabriel, M. Resch, and J. Dongarra, Eds. Berlin, Germany: Springer, 2010, pp. 229–238.

[45] S. Donfack, L. Grigori, W. D. Gropp, and V. Kale, "Hybrid static/dynamic scheduling for already optimized dense matrix factorization," in *Proc. IEEE 26th Int. Parallel Distrib. Process. Symp.*, May 2012, pp. 496–507.

[46] V. Kale, A. Randles, and W. D. Gropp, "Locality-optimized mixed static/dynamic scheduling for improving load balancing on SMPs," in *Proc. 21st Eur. MPI Users' Group Meeting (EuroMPI/ASIA)*, 2014, pp. 115–116.

[47] S. Bak, H. Menon, S. White, M. Diener, and L. Kale, "Multi-level load balancing with an integrated runtime approach," in *Proc. 18th IEEE/ACM Int. Symp. Cluster, Cloud Grid Comput. (CCGRID)*, May 2018, pp. 31–40.

[48] A. Tiwari and J. K. Hollingsworth, "Online adaptive code generation and tuning," in *Proc. IEEE Int. Parallel Distrib. Process. Symp.*, May 2011, pp. 879–892.



JOÃO B. FERNANDES received the B.S. degree in science and technology and the B.S. and M.Sc. degrees in computer engineering from the Universidade Federal do Rio Grande do Norte (UFRN), Brazil, in 2016, 2018, and 2020, respectively, where he is currently pursuing the Ph.D. degree in computer engineering. His research interest includes high-performance computing.



TIAGO BARROS received the B.S., M.Sc., and Ph.D. degrees in electrical engineering from the University of Campinas, Brazil, in 2008, 2012, and 2018, respectively. He is currently a Postdoctoral Researcher in computer engineering with the Universidade Federal do Rio Grande do Norte, Brazil. His research interests include digital signal processing applied to seismic data analysis and computational geophysics, focusing on seismic processing and inversion, especially in tomographic and waveform inversion methods.



ÍTALO A. S. ASSIS received the B.S., M.Sc., and Ph.D. degrees in computer engineering from the Universidade Federal do Rio Grande do Norte (UFRN), Brazil, in 2013, 2015, and 2019, respectively. He is currently a Professor with the Department of Informatics and Applied Mathematics (DIMAp), UFRN, where he is also a Postdoctoral Researcher in high-performance computing applied to exploration geophysics with the Laboratory of Parallel Architectures for Signal Processing-LAPPS. His research interests include high-performance computing and optimization.



SAMUEL XAVIER-DE-SOUZA (Senior Member, IEEE) was born in Natal, Brazil. He received the degree in computer engineering from the Universidade Federal do Rio Grande do Norte (UFRN), Brazil, in 2000, and the Ph.D. degree in electrical engineering from Katholieke Universiteit Leuven, Belgium, in 2007. He worked with IMEC, Belgium, as a Software/Hardware Engineer. He was also with Flemish Supercomputing Center, Belgium, as a High-Performance Computing Consultant. In 2009, he joined the Department of Computer Engineering and Automation, UFRN, where he is currently an Associate Professor. He is also the Founder and the Director of the High-Performance Computing Center-NPAD, UFRN. He received the Royal Society-Newton Advanced Fellowship, in 2016.

• • •