Impact of Mixed-Precision: a Way to Accelerate Data-Driven Forest Fire Spread Systems

C. Carrillo[1 [0000−0003−3606−7517]], T. Margalef[1 [0000−0001−6384−7389]], A. Espinosa[1 [0000−0002−6460−3789]], and A. Cortés[1 [0000−0003−1697−1293]]

Computer Architecture and Operating Systems Department
Universitat Autònoma de Barcelona, Spain.
{carles.carrillo, tomas.margalef, antoniomiguel.espinosa, ana.cortes}@uab.cat

Abstract. Every year, forest fires burn thousands of hectares of forest around the world and cause significant damage to the economy and people from the affected zone. For that reason, computational fire spread models arise as useful tools to minimize the impact of wildfires. It is well known that part of the forest fire forecast error comes from the uncertainty in the input data required by the models. Different strategies have been developed to reduce the impact of this input-data uncertainty during the last few years. One of these strategies consists of introducing a data-driven calibration stage where the input parameters are adjusted according to the actual evolution of the fire using an evolutionary scheme. In particular, the approach described in this work consists of a Genetic Algorithm (GA). This calibration strategy is computationally intensive and time-consuming. In the case of natural hazards, it is necessary to maintain a balance between accuracy and time needed to calibrate the input parameters. Most scientific codes have over-engineered the numerical precision required to obtain reliable results. In this paper, we propose to use a mixed-precision methodology to accelerate the calibration of the input parameters involved in forest fire spread prediction without sacrificing the accuracy of the forecast. The proposed scheme has been tested on a real fire. The results have led us to conclude that using the mixed-precision approach does not compromise the quality of the result provided by the forest fire spread simulator and can also speed up the whole evolutionary prediction system.

Keywords: Mixed-precision, Forest Fire Spread Simulator, Input data Uncertainty

1 Introduction

The need to anticipate potential fire behavior and its resultant impacts has led to the development of the field of fire modeling. In recent decades, various physical and mathematical models have been developed to provide reliable forecasting of fire behavior and optimize firefighting resource management during outbreaks. Simulators implementing forest fire spread models need several input parameters

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

to describe the characteristics of the circumstances where the fire occurs to evaluate its future propagation. However, there are serious difficulties in acquiring precise values of certain parameters at the same places where the fire is taking place, often because the hazard itself distorts the measurements. To minimize the input data uncertainty when dealing with forest fire spread prediction, we focus on a Two-stage prediction scheme [1]. The fundamental aim of this approach is to introduce an intermediate adjustment stage to improve the estimation of certain input parameters according to the observed behavior of the forest fire. In particular, we focus on a well-known Two-stage Strategy based on Genetic Algorithms (GA) [7]. The GA generates an initial population where each individual consists of a particular configuration of the input parameters that will be fed into the underlying forest fire spread simulator. The population will evolve using the so-called genetic operators (elitism, mutation, selection, and crossover) to obtain an improved set of parameters that reproduce the observed past behavior of the fire best. To evaluate the simulation quality of each individual, the landscape where the fires are taking place is represented as a grid of cells, and each cell will have assigned a state according to its belonging to either a real burnt area or a simulated burnt area. There are four possible states for a given cell: cells that were burnt in both fires, the real and the simulated fire, *(Hits)*, cells burnt in the simulated fire but not in reality *(False Alarms)*, cells burnt in reality but not in the simulated fire *(Misses)* and cells that were not burnt in any case *(Correct negatives)* [8] [10]. These four possibilities are used to construct a 2×2 contingency table, [4]. A perfect simulation system would have data only on the main diagonal. In the particular case of forest fire, the difference between real and simulated fire is computed using Equation 1, [16]. The resulting error is used to rank the individuals in a list from lower to higher error.

$$\in = \frac{Misses + FA}{Hits + Misses} \tag{1}$$

Most scientific codes use double-precision by default for all their floating point variables. It is taken for granted that all real number variables must be implemented using double-precision types. But in many cases, the computational cost of this decision is not well weighted and other alternatives could be considered. In many cases, not all variables in a program need this double-precision. Different studies have demonstrated the potential benefits that mixed-precision approaches can provide to many different kinds of scientific codes since it is possible to achieve substantial speed-ups for both compute and memory-bound algorithms requiring little code effort and acceptable impact on functionality, [3], [15]. The main principle is to review programs from the observation that, in many cases, a single-precision solution to a problem can be refined to the degree where double-precision accuracy is achieved. However, an excessive precision reduction of a set of relevant variables can lead to accuracy losses that may finally produce unreliable results.

In this paper, we present a straightforward methodology to review the need for double-precision variables to improve the global performance of the applications. We apply the methodology to the code of the forest fire spread simulator.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

We have modified it to select the set of variables that can use single-precision without disturbing the general evolution of the forest fire spread simulation and, even more, to determine the performance improvements. We will use well-known double-precision fire behavior results to compare and verify the quality of each mixed-precision approach generated. In particular, as a forest fire spread simulator, we used FARSITE (*Fire Area Simulator*) [9], which is a widely accepted and validated wildfire spread simulator [2] based on the Rothermel's model [17]. We have concentrated on the principal algorithms defining FARSITE floating number processing performance. Most complex modules include the fire perimeter expansion and the perimeter reconstruction that account for the 67% of the total execution time, 7% for the perimeter expansion, and 60% for the perimeter reconstruction modules. We apply the *Two-stage Strategy* to calibrate the input values required for the prediction during this work. Simulation is based on the execution of GA to evaluate hundreds of simulations, in our case, 10 generations and 128 individuals per generation. If the complexity of running these individuals can be reduced, even if the improvement is not very noticeable, it will be cumulative for all generations and the final prediction. To preserve the accuracy of the obtained prediction, we only modify the precision of those variables that do not significantly impact the fire evolution.

This paper is organized as follows. Section 2 exposes the proposed mixed-precision methodology applied to a forest fire spread prediction system. Subsequently, in section 3, the experimental study and the obtained results are reported. Finally, in section 4, the main conclusions and open lines of this work are summarized.

2 Mixed-precision methodology

Most computational science codes have overestimated the needed numerical precision of a model leading to a situation where simulators are using more precision than required without considering whether this precision is really needed. By using a more appropriate choice of 32-bit and 64-bit floating point arithmetic, the performance of some scientific applications was significantly enhanced without affecting the accuracy of the results, [11], [5], [13]. A careful way of managing floating-point precision that outstrips double-precision performance motivated us to study if using the mixed-precision variables allows reducing the execution time of the forest fire spread simulations without losing accuracy. To effectively compare the proposed mixed-precision model against a reference implementation (double-precision scheme) we need to ensure not to affect the quality of the results.

One Key point of our methodology is to evaluate the simulation quality when changing the variables of FARSITE that can use single precision. First, we need to define some metrics to quantify the quality of the mixed-precision simulation. The Forest area is divided into cells, and each cell will have a meaning. The cells around the map that have been burnt by neither the reference simulation nor the mixed-precision simulation are considered *Correct Negatives* (CN). Those cells

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

that have been burnt in both simulated fires are called *Hits*. The cells that are only burnt in the Reference simulation and are not burnt in the mixed-precision simulation are called *Misses*. The cells burnt in the mixed-precision simulation, but the Reference simulation does not reach them, are called *False Alarms* (FA), see Figure 1.

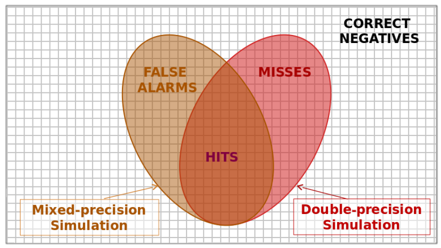

Fig. 1. Events involved in metrics related to model verification.

Three different quality metrics are used to verify the proposed mixed-precision implementation, [4]. The first metric is the *Bias score or frequency Bias* (BIAS), which represents the ratio of the number of correct forecasts over the number of correct observations, Equation 2. It represents the normalized symmetric difference between the real and simulated maps.

$$BIAS = \frac{Hits + FA}{Hits + Misses} \tag{2}$$

The second metric is the *False Alarm Rate* (FAR). The FAR measures the proportion of the wrong events forecast, see Equation 3. The FAR is sensitive to *False Alarms* but ignores the misses. A perfect comparison has a FAR value equal to 0.

$$FAR = \frac{FA}{(Hits + FA)} \tag{3}$$

The last metric is the *Probability of Detection of hits rate* (POD), see Equation 4. POD considers the observed and positively estimated events. Thus, it represents the probability of an event being detected. The POD is sensitive to hits but ignores the false alarms. Its ideal value for the POD is 1.

$$POD = \frac{Hits}{Hits + Misses} \tag{4}$$

In order to evaluate the impact of the variables of FARSITE that can use single-precision, we have to estimate how the use of mixed-precision affects the general prediction of forest fire behavior. We can summarize our methodology into three steps:

1. Define a reference execution without modifying the precision of any variable. The obtained simulation will be used as a reference against which we will validate any mixed-precision approach taken later on.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

2. Define an accuracy threshold. Table 1 displays the thresholds used for the different metrics when the mixed-precision simulations are compared against the reference.
3. Individual precision reduction of each variable and evaluating the impact of this reduction versus the reference implementation. Any variable analyzed can be defined in single-precision if the computed error complies with all three thresholds simultaneously. If the error does not satisfy one of the thresholds, the variable must keep its original double-precision.

Table 1. Thresholds for the validation metrics used to compare the reference and mixed-precision simulations.

$1.05 > \text{BIAS} > 0.95$
$\text{FAR} < 0.05$
$\text{POD} > 0.95$

In our case, an individual accuracy test was done for each variable to analyze the global impact of the mixed-precision implementation. For each test, the precision of a single variable was reduced, while the rest of the variables of the program kept their original form, and a new fire spread simulation was done. The obtained simulation was compared with the reference, and a new *BIAS*, *FAR*, and *POD* were calculated. If the simulation does not exceed the thresholds, the precision reduction for this variable was accepted. The variable maintained its original double-precision definition if the validation values were beyond the threshold limits. Then, the methodology analyzes the next variable and reduces its precision, and the validation process starts again. A new mixed-precision forest fire simulation model was obtained with the resulting final list of variables.

At the end of this analysis process, 266 different variables were tested, 221 for the point expansion algorithm and 45 for the perimeter reconstruction algorithm. We found that nearly 74% of the point expansion algorithm variables could safely use single-precision. For the perimeter reconstruction algorithm, only 15% of the variables could use single-precision without compromising the accuracy of the prediction.

This methodology is expensive in terms of execution time and resources because it is necessary to perform one simulation for each variable. If simulations consume 60 seconds, we need around 4.5 hours to apply the methodology completely. However, because this mixed-precision model is going to be used for calculating hundreds of simulations, we expect the benefits of overcoming this initial cost. Moreover, the mixed-precision fire spread model is not dependent on the forest fire scenario, which means we can use it in any future wildfire to predict its behavior.

3 Experimental Study and Results

To analyze the impact of the mixed-precision in the propagation of the fire front and into the forest fire spread simulator performance, we have selected an event

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

belonging to the database of EFFIS (*European Forest Fire Information System*) [6]. The study case took place in 2009 in the region of Nuñomoral, Spain. The forest fire began on July 25th, and the total burnt area was 3,314ha. In Figure 2, we show the fire perimeters at three different time instants: t_0 (July 26th at 11:27 am), t_1 (July 27th at 10:32am) and t_2 (July 28th at 11:15 am). The grid size of this fire is $100m$ X $100m$. All calculations reported here were performed using a cluster with 128 cores. The execution platforms are CPU Intel(R) Xeon(R) CPU E5-2620 v3 @ 2.40GHz, with 6 cores. This research used a population size of 128 individuals, and the number of iterations (generations) was set to 10. For this analysis, we assess the accuracy of the output when the mixed-precision is utilized compared to the reference simulation.

The Two-stage Strategy has two different steps: the calibration and prediction stages. In the calibration step, we have to perform hundreds of simulations to evaluate the input parameters that reproduce the real fire spread better. In order to improve the analysis in the prediction stage, the 5 individuals with the lowest error value after the calibration stage are used to perform fire predictions. The error of each simulation is calculated using Equation 1, see Figure 3. To effectively compare the double and mixed-precision implementations, we force both implementations to use the same individuals per generation during the calibration process. This ensures that the performance difference between both implementations is not a consequence of using a different combination of input parameters. Therefore, the same best individuals are selected for both implementations at the end of the calibration process.

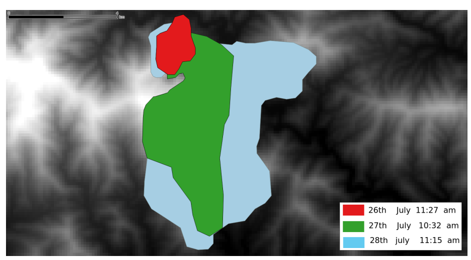

Fig. 2. Digital Elevation map of Nuñomoral fire area and the three different perimeters of the forest fire. The *Perimeter 1 - red* was used as initial perimeter (ignition Perimeter), *Perimeter 2 - green* was used in the calibration stage, and the perimeter to be predicted is *Perimeter 3 - blue*, [6].

In Figure 3(a), we can see the evolution of the computed error of the simulation through the generations of the calibration stage. The progression of the computed error reveals that the simulated perimeter of the fire becomes closer to the real fire perimeter. Moreover, we can see that the error is lower for the double-precision simulation than for the mixed-precision. Nevertheless, at the end of the calibration process, the error of the tenth generation is very close for both implementations.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

Figure 3(b) reveals the computed error of the predictions when the same five best individuals are used in both implementations. In all predictions, the error for both implementations is very close. In the example, we have used the third simulation, as it has the lowest error prediction.

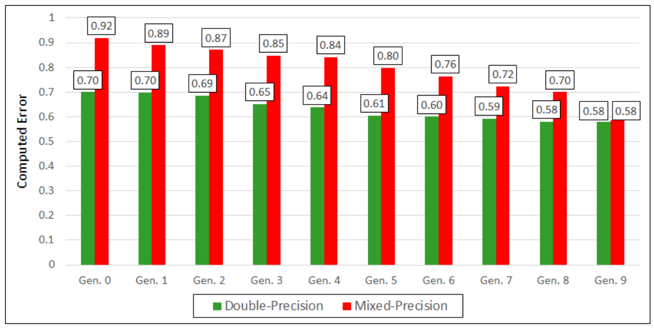

(a) Calibration Stage Error

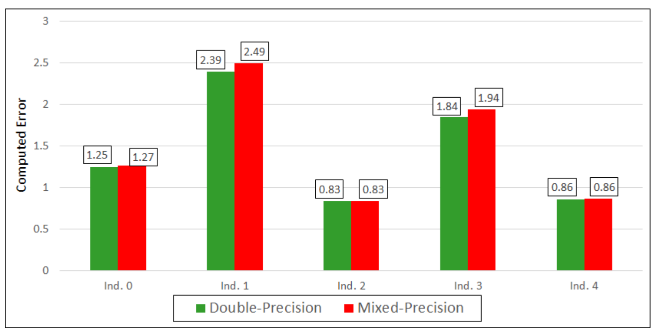

(b) Prediction Stage Error

Fig. 3. Computed error per individual using Equation 1 of the Nuñomoral's fire when double-precision (red) and mixed-precision (blue) are used.

Figure 4 details the difference in the fire evolution between double (green line) and mixed-precision (red line) when the second individual simulation is used. We can see that the difference between the perimeter evolution of both implementations increases as the fire extension widens. As we said, the precision reduction produces a small difference when Rothermel's model is applied. This difference is accumulative and increases with the velocity of the propagation of the fire front. We notice that in some parts of the maps, the perimeter of the fire spreads very fast.

Table 2 compares the best five fire predictions when the two implementations are used considering the difference between sets. It shows that the number of cells burnt by both simulations is very close. In this case, the computed validation values for the best individual prediction are $BIAS = 1.012$, $FAR = 0.018$, and

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

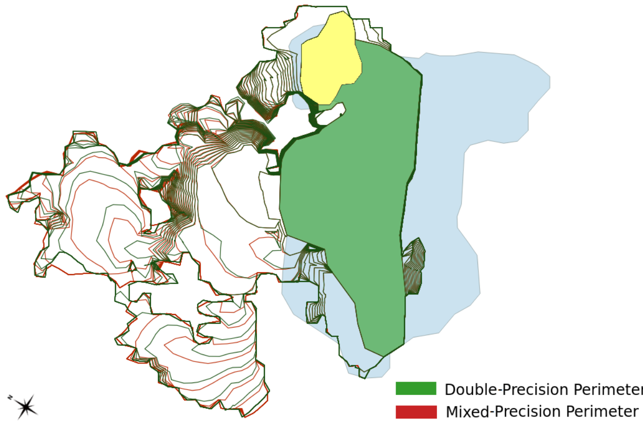

Fig. 4. Obtained simulations of the Nuñomoral's fire after the prediction stage when the third simulation is used. The initial perimeter is drawn in *green* area. The *dark red* line indicates the fire propagation when mixed-precision is used, and the *dark green* line indicates the fire propagation when we use double-precision. The *light blue* area represents the final perimeter of the fire.

$POD=$ 0.993. A BIAS value higher than one means that the mixed-precision implementation burns more area than the reference simulation. The mixed-precision implementation produces the same prediction of the fire behavior when the third individual prediction is used. Its validation values are $BIAS=$ 1.000, $FAR=$ 0.000 and $POD=$ 1.000. The Computed validation confirms that both simulations are similar and well below the defined thresholds.

In any case, a consequence of the mixed-precision implementation is that, due to its lower accuracy, the fire perimeter shape is simplified in some circumstances. This outcome has an indirect effect in reducing the execution time. When the shape is smoother, see Figure 4, we need less number of points to represent the fire front; therefore, the time consumed by the point expansion and the perimeter reconstruction algorithms decreases, which generates an execution performance improvement. The penalty for paying is an error increment versus the reference but kept within the acceptable threshold values.

Table 2. Obtained results of comparing the prediction individuals between the double-precisionand when mixed-precision using the *Two-stage Strategy*.

	1st Pred.	2nd Pred.	3th Pred.	4th Pred.	5th Pred.
BIAS	1.012	1.002	1.000	0.998	1.002
FAR	0.018	0.032	0.000	0.017	0.005
POD	0.993	0.970	1.000	0.980	0.997

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

As mentioned above, when a wildfire is simulated, the most basic information is related to the evolution tendency of fire behavior, which is the most relevant information the firefighter can use to tackle these hazards efficiently.

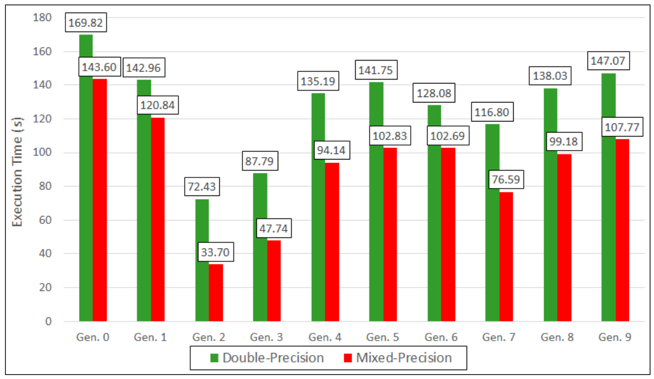

(a) Calibration time per generation.

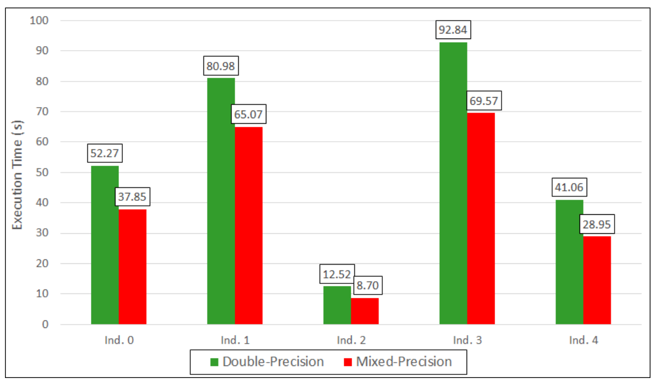

(b) Prediction time for best simulations.

Fig. 5. Execution time for two different implementations, the reference (red) and the mixed-precision (blue) in Nuñomoral's fire.

In Figure 5 we illustrate the execution time for predicting the forest fire behavior. Figure 5(a) displays the execution time per generation in the calibration stage. In our particular case, all the individuals of the same generation are executed in parallel; therefore, the execution time of a generation is determined by the slowest individual.

This figure shows a considerable reduction in the execution time when the mixed-precision is used. The highest performance improvement is obtained in the third generation, where the execution time is reduced by 53.5% when the mixed-precision is applied, representing a speed up of 2.15. The worst performance improvement is obtained in the first and second generations, where the computed

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

speed up is around 1.18 in both cases. Figure 5(b) illustrates the execution time of each individual from the prediction stage. As in the calibration stage, the usage of mixed-precision shows a significant reduction in the execution time. The maximum speed up, 1.44, is obtained by the third best individual. If we focus only on the best individual, the speed up calculated is 1.38, which represents a substantial performance improvement. Therefore, we can confirm that using mixed-precision methodology reduces the execution time.

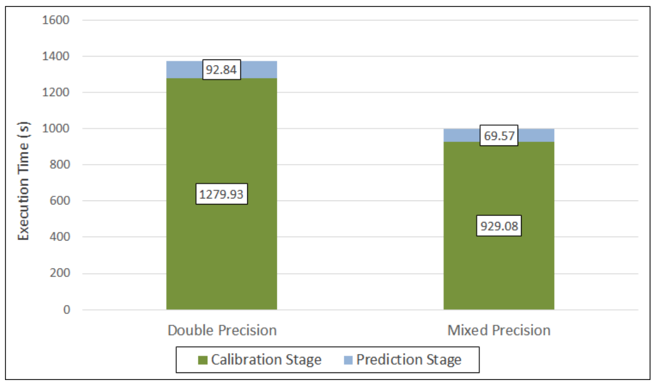

Fig. 6. Execution time of the *Two-stage Strategy* for the different scenarios: double and mixed-precision in Nuñomoral's fire.

Figure 6 details the execution time invested in completing a forest fire spread prediction. The green column represents the execution time of the *calibration Stage* and the blue column the *prediction stage* for the two different implementations: double and mixed-precision. We can see that, in this fire scenario, the reduction of the execution time is notable, 374.12 seconds, 350.85 seconds for the Calibration Stage, and 23.27 for the Prediction Stage, which represents an improvement around the 27.3% of the execution time, 27, 4% for the Calibration Stage and 25.1% for the Prediction Stage. The computed speeds up are 1.40 for Calibration Stage and 1.33 for the Prediction Stage, respectively. The speed up of the whole execution is 1.37. We note that for the case of Nuñomoral's fire, the execution time reduction is significant.

In order to better understand the computational reasons for the performance improvements of the mixed-precision in the *Two-stage Strategy*, we used a profiling tool, *Linux prof*, to get the basic performance metrics of the experiments and compare the results of the execution. We will use CPU Cycles Per machine code Instruction (CPI) as a metric for comparing the computational cost of both implementations. The CPI is the ratio between the number of CPU cycles over the number of instructions executed. It reflects the average number of CPU cycles needed to complete an instruction. Thus, CPI indicates how much latency is in the system and can be a useful measure of how an application is perform-

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

ing. A low CPI indicates that a few cycles are needed to execute an instruction. However, if the CPI is high, the execution of a single instruction requires a large number of CPU cycles, which indicates that the application has poor performance. Because CPI is affected by changes in the number of CPU cycles an application takes or changes in the number of instructions executed, it is best used for comparison when only one part of the ratio changes. The goal is to decrease the CPI in certain parts of the code as well as the whole application. In our particular case, we want to understand the impact of mixed precision in the complexity of calculating specific functionality. CPI helps to determine the reduction of the processing complexity in those parts of the code where the mixed precision has been implemented, as in the point expansion and perimeter reconstruction algorithms. Figure 7(a) displays the average CPI for each generation in the calibration stage. It shows that the usage of mixed-precision implementation improves the average CPI in all generations. As we can see, the highest CPI reduction is found in the fourth generation, with a reduction of 28%. This substantial computational cost reduction is due to a high number of individuals with an extended perimeter. In these individuals, the weight of the point expansion algorithm is large, so we have many potential instructions that can be simplified with mixed precision.

This notable difference is a consequence that in the reference execution, the fire perimeter tends to spread slightly further than the mixed precision simulations. In some cases, these small differences in fire propagation have the consequence of a sudden acceleration of the fire propagation in areas of the terrain with the worst possible combination of the factors that drive the fire evolution. These factors are the slope, aspect, fuel, and wind. When these four factors are combined in a specific adversarial way, the acceleration of the propagation of the fire front increases exponentially when the fire arrives at this area. In this case, double-precision simulation implementation introduces artificially complicated cases that significantly increase the computational cost of the simulations. When the mixed-precision implementation is used, some complex calculations are not executed as we avoid some particularly complex simulation cases: high error simulations with extra perimeter points and more work to simulate an extensive area propagation. On top of this, it is important to consider how these special cases affect the final predicted burned area. Ultimately, most of these expensive cases are discarded as they produce high error rates well above the quality threshold limits. In summary, mixed-precision avoids the execution of long, complex high-error cases whose results are not useful for the final prediction.

Another consideration that justifies this computational cost reduction is due to the *Distance Resolution*, see [9]. When the fire front propagation is accelerated, the perimeter points will propagate forward a long distance. As explained above, the Distance Resolution limits the maximum expansion of a single perimeter point in a single time iteration. For that reason, the propagation of these perimeter points is divided into a long set of shorter propagation sub-steps, which implies a greater consumption of resources by the perimeter expansion

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

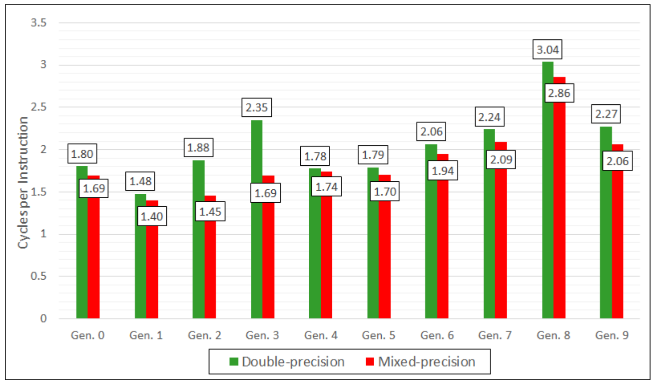

(a) Calibration CPI per generation (avg.)

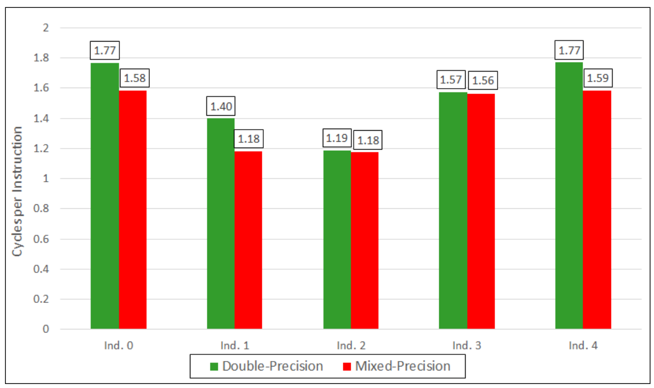

(b) Prediction CPI for best simulations

Fig. 7. CPI of the *Two-stage Strategy* for two different implementations, reference and mixed-precision in Nuñomoral's fire.

algorithm. When mixed-precision is applied, the propagation is much simpler as it needs a few steps. If we evaluate the whole calibration stage, the obtained CPI is 1.91 for the reference simulation and 1.67 for the mixed-precision implementation, representing a reduction of 12.60% of the computational cost. So, we can conclude that the utilization of the mixed-precision produces a notable reduction of the computational cost maintaining the quality of the simulation. In Figure 7(b), we can observe the CPI of the prediction stage. It shows that, as in the calibration stage, the utilization of the mixed-precision provides a fair reduction of the computational cost.

In this scenario, the highest CPI decrease is found in the second-best case, where the computational cost decreases from 1.40 CPI in the reference simulation to 1.18 when the mixed-precision implementation is used. This represents a reduction of the 15.5% of the computational cost. In a real forest fire emergency, the velocity to predict future fire behavior is crucial; for that reason, only the best

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

individual is used to predict the fire spread. If we focus on the best individual, we see a reduction of the 10.5% of the CPI. This computational cost reduction improves performance when the mixed-precision approach is employed.

4 Conclusions

Forest fire is a significant natural hazard that every year causes important damages. Models and their implementation in simulators can estimate their behavior, but they are not exempt from a certain degree of error. The quality results of these models depend not only on the propagation equations describing the behavior of the fire but also on the input data required to initialize the model. Typically, this data is subjected to a high degree of uncertainty and variability during the evolution of the event due to the dynamic nature of some of them, such as meteorological conditions or moisture contents in the vegetation. Consequently, there is a need for strategies to minimize this uncertainty to provide better predictions. Previous studies have demonstrated that using the *Two-stage Strategy* undoubtedly increases the accuracy of the predictions. However, due to its characteristics, *Two-stage Strategy* implies that the execution time invested in predicting the fire propagation increases significantly. Different works expose that the generalized use of double-precision in most scientific codes could be notably reduced. For this reason, using mixed-precision would pay back in terms of performance improvement, requiring little coding effort.

In this work, we present a mixed-precision methodology applied to improve the performance of a forest fire propagation simulation. In particular, we focus on the forest fire spread simulator FARSITE, which is a widely accepted tool in the scientific community related to this field. However, the proposed methodology is simulator independent and could be applied to any other. A relevant outcome from the experimental study is that around 74% of the variables used to calculate the point expansion algorithm can be defined in single-precision. On the other hand, only 15% of the variables used in the perimeter reconstruction module could be used with single-precision because that part represents, in some cases, around 60% of the total simulation time. For that reason, the performance improvement using the mixed-precision is limited in this particular case.

The experimental study performed analyzes the impact of the mixed-precision on the *Two-stage Strategy* where a genetic algorithm is used in a data-driven way to determine the most suitable values of the input simulator parameters based on the current evolution of the forest fire. The proposed mixed-precision methodology has been tested using a real fire that took place in 2009 in the region of Nuñomoral, Spain. The results suggest that using mixed-precision can provide remarkable performance improvements without compromising the simulation results in terms of quality. The implementation of the mixed prediction reduces the execution time by around the 30% of the whole process. We have to take into account that the improvement is accumulative. One way to improve the prediction quality of the fire spread is to increase the number of generations and the number of individuals per generation. In this case, the potential

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

improvement provided by the mixed-precision implementation will be higher. Moreover, all future simulations will benefit from this implementation without additional effort. In terms of quality, for short fire front propagation, the difference between both obtained perimeters is negligible. However, this difference is accumulative; therefore, when the fire propagation increases, this difference is more discernible. As we saw, in some cases, the double-precision implementation introduces an artificial complexity in the simulations that significantly increases the computational cost of those executions and, therefore, the time invested into the *Two-stage Strategy*. Finally, although the evolutionary mixed-precision data-driven forest fire spread system described in this work has been focused on accelerating the forecast of wildfire's evolution, it should be considered that the complete methodology can be generalized for any other natural hazard forecast systems.

This study is a starting point for utilizing the mixed-precision methodology in the forest fire spread simulation field. Future works include the utilization of different tools, see [14], [18] and [12], to evaluate a different set of variables that can be used in single-precision, to check out the best combination to optimize the performance benefit with the lowest accuracy reduction. In addition, using accelerators like GPUs could increment the impact of the methodology on the performance of the simulator, as the savings in mixed-precision could be higher.

Acknowledgements This work has been granted by the Spanish Ministry of Science and Innovation MCIN AEI/10.13039/501100011033 under contract PID2020-113614RB-C21 and by the Catalan government under grant 2021 SGR-00574.

References

1. Abdalhaq, B., Cortés, A., Margalef, T., Luque, E.: Enhancing wildland fire prediction on cluster systems applying evolutionary optimization techniques. Future Generation Comp. Syst. **21**(1), 61–67 (2005), `http://dx.doi.org/10.1016/j.future.2004.09.013`
2. Arca, B., Laconib, M., Maccionib, A., Pellizzaroa, G., Salisb, M.: P1. 1 validation of farsite model in mediterranean area (2014)
3. Baboulin, M., Buttari, A., Dongarra, J.J., Kurzak, J., Langou, J., Langou, J., Luszczek, P., Tomov, S.: Accelerating scientific computations with mixed precision algorithms. Computer Physics Communications **180**(12), 2526–2533 (2009), `https://doi.org/10.1016/j.cpc.2008.11.005`
4. Bennett, N.D., et alter., B.: Characterising performance of environmental models. Environmental Modelling & Software **40**, 1–20 (2013). `https://doi.org/http://dx.doi.org/10.1016/j.envsoft.2012.09.011`
5. Buttari, A., Dongarra, J., Langou, J., Langou, J., Luszczek, P., Kurzak, J.: Mixed precision iterative refinement techniques for the solution of dense linear systems. The International Journal of High Performance Computing Applications **21**(4), 457–466 (2007), `https://doi.org/10.1177/1094342007084026`
6. Centre, J.R.: European forest fire information system. http://forest.jrc.ec.europa.eu/effis/ (August 2011)

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5

7. Denham, M., Wendt, K., Bianchini, G., Cortés, A., Margalef, T.: Dynamic data-driven genetic algorithm for forest fire spread prediction. J. Comput. Science **3**(5), 398–404 (2012), http://dx.doi.org/10.1016/j.jocs.2012.06.002
8. Ebert, B.: Forecast verification:issues, methods and faq. http://www.cawcr.gov.au/projects/verification (June 2012)
9. Finney, M.A.: Farsite: Fire area simulator—model development and evaluation. Research Paper RMRS-RP-4 Revised **236** (1998)
10. Gjertsen U, O.V.: The water phase of precipitation: a comparison between observed, estimated and predicted values. http://dx.doi.org/10.1016/j.atmosres.2004.10.030 (October 2005)
11. Jia, X., Song, S., He, W., Wang, Y., Rong, H., Zhou, F., Xie, L., Guo, Z., Yang, Y., Yu, L., Chen, T., Hu, G., Shi, S., Chu, X.: Highly scalable deep learning training system with mixed-precision: Training imagenet in four minutes. CoRR **abs/1807.11205** (2018), http://arxiv.org/abs/1807.11205
12. Lam, M.O., de Supinksi, B.R., LeGendre, M.P., Hollingsworth, J.K.: Poster: Automatically adapting programs for mixed-precision floating-point computation. In: 2012 SC Companion: High Performance Computing, Networking Storage and Analysis. pp. 1424–1424 (2012). https://doi.org/10.1109/SC.Companion.2012.232
13. Maynard, C., Walters, D.: Mixed-precision arithmetic in the endgame dynamical core of the unified model, a numerical weather prediction and climate model code. Computer Physics Communications (2019), http://www.sciencedirect.com/science/article/pii/S0010465519302127
14. Menon, H., Lam, M.O., Osei-Kuffuor, D., Schordan, M., Lloyd, S., Mohror, K., Hittinger, J.: Adapt: Algorithmic differentiation applied to floating-point precision tuning. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis. SC '18, IEEE Press (2019). https://doi.org/10.1109/SC.2018.00051, https://doi.org/10.1109/SC.2018.00051
15. Prikopa, K.E., Gansterer, W.N.: On mixed precision iterative refinement for eigenvalue problems. In: Proceedings of the International Conference on Computational Science, ICCS 2013, Barcelona, Spain, 5-7 June, 2013. Procedia Computer Science, vol. 18, pp. 2647–2650. Elsevier (2013), https://doi.org/10.1016/j.procs.2013.06.002
16. Rodríguez, R., Cortés, A., Margalef, T.: Injecting dynamic real-time data into a DDDAS for forest fire behavior prediction. In: Allen, G., Nabrzyski, J., Seidel, E., van Albada, G.D., Dongarra, J.J., Sloot, P.M.A. (eds.) Computational Science - ICCS 2009, 9th International Conference, Baton Rouge, LA, USA, May 25-27, 2009, Proceedings, Part II. Lecture Notes in Computer Science, vol. 5545, pp. 489–499. Springer (2009), https://doi.org/10.1007/978-3-642-01973-9_55
17. Rothermel, R.: A mathematical model for predicting fire spread in wildland fuels (US Department of Agriculture, Forest Service, Inter- mountain Forest and Range Experiment Station Ogden, UT, USA 1972)
18. Rubio-González, C., Nguyen, C., Nguyen, H.D., Demmel, J., Kahan, W., Sen, K., Bailey, D.H., Iancu, C., Hough, D.: Precimonious: Tuning assistant for floating-point precision. In: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis. SC '13, Association for Computing Machinery, New York, NY, USA (2013). https://doi.org/10.1145/2503210.2503296, https://doi.org/10.1145/2503210.2503296

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_5