

Scalability of a multi-physics system for forest fire spread prediction in multi-core platforms

Angel Farguell¹ · Ana Cortés¹  · Tomàs Margalef¹ · Josep R. Miró² · J. Mercader²

Published online: 20 March 2018
© The Author(s) 2018

Abstract Advances in high-performance computing have led to an improvement in modeling multi-physics systems because of the capacity to solve complex numerical systems in a reasonable time. WRF–SFIRE is a multi-physics system that couples the atmospheric model WRF and the forest fire spread model called SFIRE with the objective of considering the atmosphere–fire interactions. In systems like WRF–SFIRE, the trade-off between result accuracy and time required to deliver that result is crucial. So, in this work, we analyze the influence of WRF–SFIRE settings (grid resolutions) into the forecasts accuracy and into the execution times on multi-core platforms using OpenMP and MPI parallel programming paradigms.

Keywords Forest fire simulation · Multi-physics model · Multi-core · Scalability

1 Introduction

There are several factors that affect the evolution of a wildland fire. It is well known that wind is one of the key parameters to understand the forest fire propagation. Intuitively, the meteorological wind speed tends to drive the main direction of forest fire spread. However, in large forest fires that take place in complex terrain where the wind fluxes can vary due to the topography and the wind convections produced by the heat fluxes generated by the fire, the effective wind speed and wind direction can be unpredictable.

✉ Ana Cortés
ana.cortes@uab.es; ana.cortes@uab.cat

¹ Universitat Autònoma De Barcelona, Carrer de les Sitges s/n, Campus Bellaterra, 08193 Cerdanyola del Valles, Spain

² Servei Meteorològic de Catalunya, c/Berlin 38-48, 08029 Barcelona, Spain

For that reason, a multi-physics forest fire model that considers the feedback between an atmospheric model and a forest fire spread model could capture the micro-weather generated by a large forest fire and provide more accurate wildland fire predictions. However, this improvement in accuracy has a cost in terms of execution time. This limitation could be a serious drawback for using this system during a real event. Nonetheless, the predictability potential of a multi-physics forest fire spread system is supposed to be better than considering the atmosphere model and the forest fire as two isolated systems [3–5]. There are many works that analyze this problem showing that any strategy that takes into account the atmosphere effect into the forest fire evolution provides better forecasted results [2, 10]. In fact, there are three different approaches to tackle this relationship: the one-way, the two-way and the integration strategies. With a one-way approach, the atmosphere model provides an external forcing to the fire spread model, but the prognostic and diagnostic variables of the fire spread model (fire front position, heat fluxes,...) are not used to update the atmosphere model. Current operational wildfire simulators such as BehavePlus [1], FARSITE [7] and PROMETHEUS [15] mainly rely on this approach. The second approach, the two-way scheme, is a coupled system where the fire propagation considers meteorological information and also on the other way round. The state of the atmosphere model is modified by the outputs of the fire spread model, which leads to a modification of the micro-meteorology, thus of local winds near the wildland fire. This local wind variation has a feedback on the fire behavior itself that can have an increasing effect during the fire event. This coupling between a 3D representation of the atmosphere and 2D representation of the fire captures an important range of the fire behaviors ([3–5]). WRF–SFIRE [10] and FOREFIRE/Meso-NH [6] are relevant efforts on this coupled approach for mesoscale simulations. Finally, the integration scheme uses microscale integrated simulations that rely on Navier–Stokes equations and thereby describe mass, momentum and energy conservations equations. FIRETEC [9] and WFDS [11] are two examples of this approach. However, these schemes are computationally intensive and cannot be used for operational purpose, they are limited to research projects and reduced-size computational domains.

In this work, we focus on the multi-physics system WRF–SFIRE with the aim of analyzing its predictability compared to the time incurred in delivering the results. The main objective is to determine the viability of using WRF–SFIRE as an operational tool. The major constraint of these complex multi-physics system is its execution time. Fortunately, WRF–SFIRE has been parallelized using different parallel programming paradigms such as OpenMP and MPI. Therefore, an exhaustive analysis about the trade-off between results accuracy and time incurred in providing those results should be done.

The paper is organized as follows. Section 2 describes the WRF–SFIRE system. In Sect. 3, a real forest fire is used to analyze time requirements when running on a multi-core platform. The results accuracy depending on WRF–SFIRE initialization settings and a comparative analysis using OpenMP and MPI parallel programming paradigms are also included. Finally, the main conclusions of this work are reported in Sect. 4.

2 WRF–SFIRE

WRF–SFIRE is a forest fire simulator which couples the meteorological model WRF–ARW and the fire spread model SFIRE. Large-scale wildland fires are described as a propagation front separating the burning area to the vegetation. The quantity of interest is the speed at which the front propagates, which is called rate of spread. The model to evaluate this value is the Rothermel’s model [13]. The Rothermel’s model is based on the semiempirical Rothermel formula which computes the rate of spread corresponding to a point in the fire front depending on the effects of vegetation, slope and wind. A front-tracking solver that propagates the fire front at the Rothermel’s rate of spread is required. In WRF–SFIRE, this part is done by the level set method [12]. This method uses a three-dimensional construction, the level set function, to define the evolution of the fire front by level sets as time went.

Particularly, WRF–SFIRE solves this multi-physical problem related to forest fire propagation in a cyclic fashion. That is, the atmosphere state is obtained by running WRF for a certain period of time and, the obtained WRF information is used to expand the fire front numerically using the level set method. In each time step of the atmospheric model, the fire model is executed. The wind data generated by the WRF model is first interpolated for being used in the rate of spread formula included in the fire model. Then, the fire model makes its own time step computing the heat fluxes in each fire model cell. Afterward, these heat fluxes are inserted into the atmospheric model and, then, the atmospheric model completes the time step. In order to solve this constant feedback between both models properly, it is necessary to build a discretization of the domain where the hazard is taking place. The fire is a subgrid-scale process with respect to the atmosphere model, thus two meshes of different spatial resolution are required. The atmospheric model works on a logically quadrilateral 3D grid representing the atmosphere at different vertical levels, whereas the fire model uses a refined 2D fire mesh posed in the Earth surface level. The recommended mesh ratio is 1:10, so the fire grid resolution is 10 times the atmospheric grid resolution [10]. This grid scheme is depicted in Fig. 1.

In order to capture the small-scale meteorological processes, the atmospheric grid resolution should be as fine as possible. However, reducing the grid resolution implies to increase the total number of grid cells in the atmosphere model, and consequently the number of grid cells in the fire spread model given the ratio 1:10. Since this fact has a direct effect in the WRF–SFIRE execution time, to chose a proper grid resolution is a critical feature to analyze. In the subsequent section, a real case has been used to analyze the dependency between the grid resolution and the computational cost of the simulation to find the appropriate trade-off for operational application of the system. For this purpose, four different grid resolutions will be studied in this work.

3 Experimental study

The real case analyzed is a forest fire occurred in Catalonia (North-East of Spain) on July 8, 2005 [8]. The 2005 Cardona fire burned a total surface of 1438 ha., and it lasted 6 h. The fire started at 14:30, and it kept until 20:30 approximately. This particular case




Fig. 1 3D atmospheric grid and 2D fire grid with a mesh ratio of 1:10


was labeled by the firefighter as a forest fire driven by the winds generated by the fire itself; thus, it is a perfect example to show the benefits of using a multi-physics system that takes into account the feedback between the atmosphere and fire processes. To run WRF–SIFRE, it is required to determine the computational domain taking into account that the atmosphere model is represented using a 3D grid and that the fire is propagated at the surface, so on the 2D horizontal grid of the 3D atmospheric grid. Once the domain has been determined, the equations of the atmospheric model are resolved in all domain points. The size of the computational domain is fixed in this work and corresponds to 49 km^2 ($7 \text{ km} \times 7 \text{ km}$). This computational domain is large enough to include the evolution over the whole 6-h time period.

Regarding the atmospheric mesh initialization, the initial weather data used as initial conditions for the simulations is a weather data set provided by the SMC (Servei Meteorologic de Catalunya) [14], which are at 3 km horizontal resolution and interpolated at 19 vertical levels. The digital elevation map used has a resolution of 25 m, which is the minimum input data resolution used in the system. For that reason, we have selected as initial coarser fire grid resolution 24 m, which is the number closer to 25 that allows to have a fire and meteorological grid resolutions that fits the underlying domain size. Table 1 summarizes the four atmospheric and fire grid resolutions used in the experimental study as well as the number of points in the atmospheric grid in one vertical level.

The main objectives of the experimental study reported in this section consist of analyzing the relation between the prediction quality and the time incurred to obtain it. WRF–SFIRE has been parallelized using OpenMP and MPI, so in this work, we focused also on the parallelization. In particular, we have studied the relationship between the scalability improvements in terms of execution time when using shared

Table 1 Grid resolutions used in the four experiments performed

Atmospheric res. (m)	Fire res. (m)	Atmospheric grid points per vertical level
236	24	29 × 29
118	12	59 × 59
100	10	70 × 70
59	6	116 × 116

**Fig. 2** Cardona forest fire final burnt area compared to the forest fire prediction when using WRF–SFIRE with atmospheric grid resolutions equal to 236, 118, 100 and 59 m (color figure online)

and distributed memory versus the quality improvements due to the different grid resolutions used to solve the system.

3.1 Quality results

The main advantage of using a coupled multi-physics system to predict the evolution of a forest fire is the ability of capturing the effect of meteorological events that happen in a high resolution and, on the way back, the capacity to observe the influence of the heat fluxes generated by the fire into the atmosphere. Figure 2 shows the fire front position predicted by WRF–SFIRE when using different grid resolutions. The green shape corresponds to the observed final burnt area after the 6-h event, and the dotted lines are the time-evolving fire front positions. As we can observe, the atmospheric grid resolution plays a relevant role in the final result in terms of quality. Moving to higher spatial resolutions improves the match with the observation especially at the head of the fire. There is a significant change in behavior on the flanks of the fire only when reaching the higher resolutions (59 m). The resolution does not have a significant impact at the rear part of the fire as expected since the propagation is very low in this part of the fire. The main reason of these results is the capacity of WRF to better detect local winds convections due not only to the atmosphere effects but also

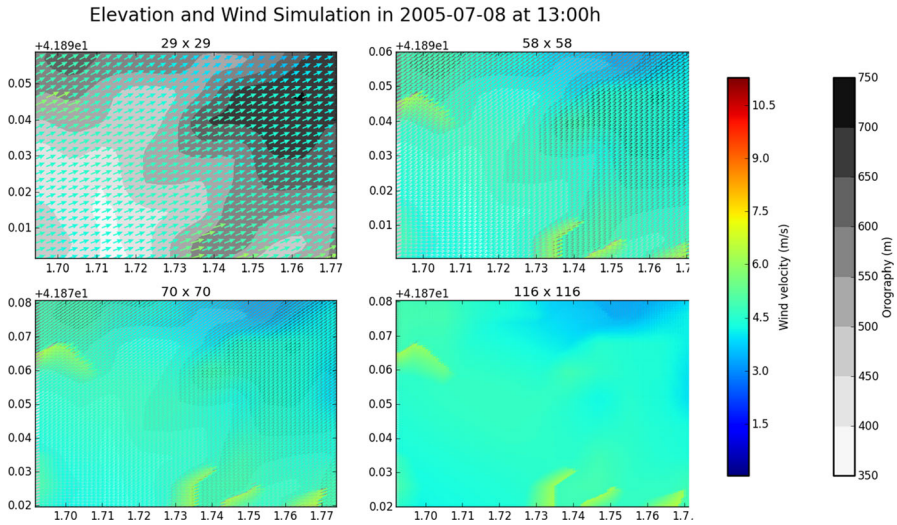


Fig. 3 Graphical representation of the wind speed and wind direction at fire initial time using WRF–SFIRE for the four experiments described in Table 1

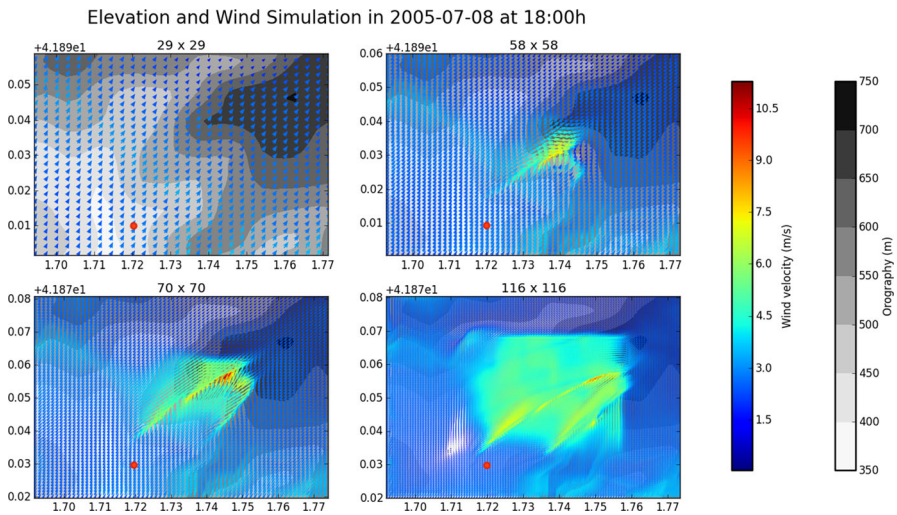


Fig. 4 Graphical representation of the wind speed and wind direction 5 h later the fire starts using WRF–SFIRE for the four experiments described in Table 1. The red dots are the ignition points

because of the forest fire. Figures 3 and 4 show the wind speed and wind direction at two different moments of the fire propagation. Figure 3 depicts wind conditions before fire ignition, and Fig. 4 shows wind parameters 5 h later. In both cases, the four above-mentioned grid configurations were used, to determine, not only the influence of the coupled system into the atmosphere behavior, but also the precision in detecting such interaction. As it can be observed, in Fig. 3, independently on the grid resolution, the only wind observed is the one generated by the WRF model. However, when

considering the effect of the atmospheric grid resolution once the fire has started (see Fig. 4), one can clearly observe that the results are quite different. Increasing the atmospheric resolution allows the model to capture the winds generated by the fire itself.

However, a question that arises at this point is can we increase the mesh resolution as much as we want in order to obtain better prediction results? The answer is not straightforward. Figure 2 depicts the predicted fire evolutions when using 236, 118, 100 and 59 m as atmospheric grid resolution. The first comparison was performed using only the values 236, 118 and 59 because each value implies twice resolution respect to the next one. This variation clearly affects the quality of the results. However, as we will see in the next section, high resolution simulations imply higher execution times, for that reason, we also test an intermediate case with the atmospheric grid resolution equal to 100 m. As we can see, this case provides a forest fire spread similar to the one obtained with a grid resolutions equal to 59 m but, as it is following explained, the execution time is significantly lower. Therefore, it is important to be able to determine a relation between grids resolutions and prediction quality if we want to use this multi-physics system in an operational way.

3.2 WRF–SFIRE scalability

As it has been introduced, WRF–SFIRE is a multi-physics system that couples the WRF atmospheric model and the SFIRE forest fire spread model. WRF requires a 3D grid to solve the atmosphere processes, meanwhile SFIRE works in a 2D mesh. The execution platform used for the experiments using shared memory is a multi-core system composed of 2 sockets integrating Intel Xeon processors with 8 cores and multithreading. On the other hand, the execution platform when running the model in its MPI version consists of a cluster of 24 compute nodes each equipped with two Intel Xeon E5-2670 Sandy Bridge CPUs with 16 cores each. In order to analyze the parallelization improvements and the scalability of the model, we have previously executed the study case using WRF–SFIRE in its baseline serial configuration. Figure 5 shows the execution time spent when simulating Cardona fire using a single thread approach for all tested grid resolutions. Moreover, the execution time spent for each individual model (WRF and SFIRE) is also depicted. As we can see, lower resolutions imply higher executions times varying from, approximately, 17 h to 7 days depending on the resolution used. Whatever grid resolution we use, the execution times obtained are prohibitively for operational purposes. It is also remarkable what happened when using the atmospheric grid resolution of 236 m. As we can see, the total execution time for this case and for the experiment with atmospheric grid resolution equal to 118 m are quite similar. This could be seen as an anomaly; however, there is a reason to explain this behavior. As it was previously mentioned, WRF uses a 3D grid to solve the atmospheric equations and, for all the experiments, the number of vertical level was initially set to 29. However, when executing the case with a grid resolution equal to 236 m the system was not able to converge to a solution due to the reduced number of grid points used. For that reason, in this particular case, the number of vertical levels was increased to 200 to include the minimum required number of points that allows the

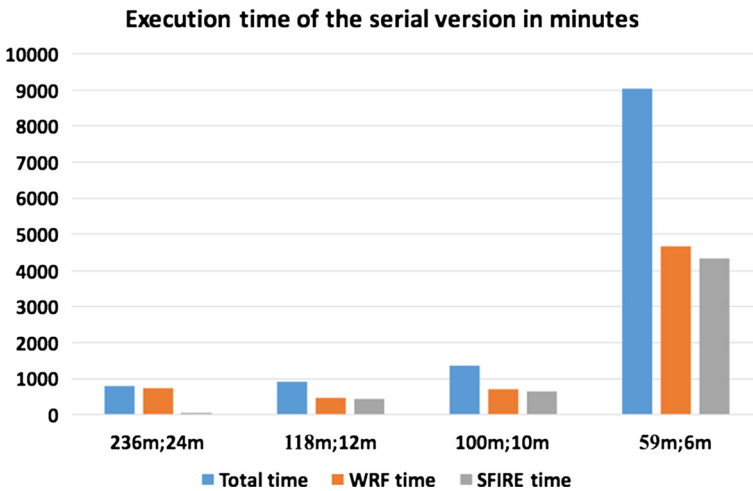


Fig. 5 Total serial execution time simulating Cardona fire using WRF–SFIRE with atmospheric grid horizontal resolutions equal to 236, 118, 100 and 59 m (color figure online)

system to converge. Since the final number of points in the atmospheric grid has been similar to the case where the grid resolution is equal to 118 with 29 vertical levels, the total execution time is quite similar. However, if we analyze the execution times of WRF and SFIRE separately, we can observe that for the case of 236 m grid resolution the time required to execute WRF is almost 90% of the total execution time, whereas in the case of 118 m grid resolution the total execution time is balanced among both models.

Fortunately, WRF–SFIRE has been parallelized using OpenMP and MPI so, in order to be able to use the system during an ongoing event, we have studied the behavior of the multi-physics system on multi-core platforms. The MPI parallelization is done through patches, that is, the domain is divided into a fixed number of parts according to the amount of cores available for a given simulation. These patches used to be square division of the domain using the horizontal plane as a reference. Since each patch needs information from its neighbors' patches to run the model, each patch includes extra points (called halo) to incorporate those points from the four borders that are required to execute on one iteration of the model. After finishing each iteration, the patches must exchange the results from the points in the halo, therefore, a synchronization barrier is required. This scheme implies that all MPI processes proceed in a synchronized fashion what can imply a non-depreciable communication time if the patch size is not well evaluated. In the case of OpenMP parallelization, the domain is divided into tiles. The number of tiles will be set to the amount of threads that will be run the model. Figure 6 shows the patch division and tile division of a domain based on a horizontal plane with 16×16 points in one vertical level.

Since the simulations that provide better quality results are the ones with atmospheric grid resolutions equal to 100 and 59 m, we have focused the study on these two cases. Figure 7 shows for an atmospheric grid resolution of 100 m, the ideal execution time, the real total execution time and the execution time of the two models




Fig. 6 Patches division and tiles division of a domain depicted in the horizontal plane




Fig. 7 Execution time simulating Cardona fire using WRF–SFIRE with atmospheric grid resolutions 100 m and OpenMP with 2, 4, 8 and 16 threads

(WRF and SFIRE) when using 2, 4, 8 and 16 threads locating each thread in a different core. Figure 8 depicts the same information than Fig. 7 but for an atmospheric grid resolution of 59 m. As we can observe, the scalability of the system is quite good because the total execution time has almost the same tendency that the ideal case. However, in terms of absolute execution time incurred in both cases, it is clear that the one that could reach the execution time requirements for operational purposes, is the experiment with atmospheric grid resolution equal to 100 and running using either 8 or 16 threads. Consequently, a deeper study to be able to assess in advance the appropriate grid settings and the hardware requirements to cope with operational constraints must be done using shared memory parallel strategy.

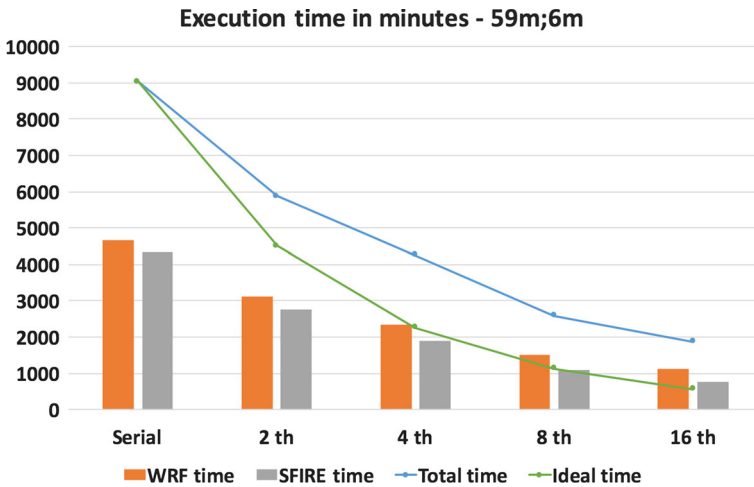


Fig. 8 Execution time simulating Cardona fire using WRF–SFIRE with atmospheric grid resolutions 59 m and OpenMP with 2, 4, 8 and 16 threads




Fig. 9 Execution time simulating Cardona fire using WRF–SFIRE with atmospheric grid resolutions 100 m and MPI with 48, 96 and 144 cores

In the case of using MPI, Figs. 9 and 10 show the ideal execution time, the real total execution time and the execution time of the two models when using 48, 96 and 144 cores in a MPI distributed memory parallelization. It is possible to observe that execution time is reduced a lot, but it is very different from the ideal execution time. Furthermore, the CPU hours spent are really huge because this execution time is distributed among all the cores. Therefore, it is important also to take into account which cases provide an accurate result while minimizing the computational cost in terms of CPU hours. So, it seems that using an adequate setup and parallelization strategy, it could be possible to reach operational times.




Fig. 10 Execution time simulating Cardona fire using WRF–SFIRE with atmospheric grid resolutions 59 m and MPI with 48, 96 and 144 cores

4 Conclusions

Atmospheric wind is one of the parameters that most affect the forest fire behavior. However, this atmospheric phenomenon can be hardly affected by the forest fire micro-weather generated by the fire itself. WRF–SFIRE is a coupled atmospheric–fire system that takes into account the continuous relation between both models. However, the penalty of continuously evaluating the evolution of the atmosphere processes and the fire processes is the execution time. WRF–SFIRE has an OpenMP and MPI parallel implementation that allows to exploit multi-core systems. These parallel approaches have been tested using as study case: the Cardona fire. This forest fire was a fire driven by the winds generated by the fire itself, so to run the simulations with high resolutions is the best approach to cope the local effect of the fire in the atmosphere. However, to reach operational times, it is necessary to exploit to the maximum the OpenMP and MPI parallelization of WRF–SFIRE taking into account convergence features of the system that can go from modifying the model grid resolutions to improving the vertical levels of the meteorological grid in order to have enough points to allow the system to converge. The results obtained using both OpenMP and MPI parallel programming paradigms isolate one of each other, show that using the required grids settings and distributed memory parallelization strategy, it can be possible to reach operational times. Therefore, next step will consist of hybridizing the system using both MPI and OpenMP approaches simultaneously trying to put together the benefits of both schemes.

Acknowledgements This work has been supported by Ministerio de Economía y Competitividad (Spain) under contract numbers TIN2014-53234-C2-1-R and TIN2016-81840-REDT

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution,

and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Andrews PL (2007) Behaveplus fire modeling system: past, present, and future. In: 7th Symposium on Fire and Forest Meteorology, Paper J2.1
2. Brun C, Artés T, Margalef T, Cortés A (2012) Coupling wind dynamics into a DDDAS forest fire propagation prediction system. *Procedia Comput Sci* 9:1110–1118. <https://doi.org/10.1016/j.procs.2012.04.120>
3. Clark TL, Coen J, Latham D (2004) Description of a coupled atmosphere-fire model. *Int J Wildland Fire* 13(1):49–63
4. Clark TL, Jenkins MA, Coen J, Packham D (1996) A coupled atmosphere fire model: convective feedback on fire-line dynamics. *J Appl Meteorol* 35(6):875–901. [https://doi.org/10.1175/1520-0450\(1996\)035<0875:ACAMCF>2.0.CO;2](https://doi.org/10.1175/1520-0450(1996)035<0875:ACAMCF>2.0.CO;2)
5. Clark TL, Jenkins MA, Coen J, Packham D (1996) A coupled atmosphere-fire model: role of the convective froude number and dynamic fingering at the fireline. *Int J Wildland Fire* 6(4):177–190
6. Filippi JB, Bosseur F, Pialat X, Santoni PA, Strada S, Mari C (2011) Simulation of coupled fire/atmosphere interaction with the MesoNH-ForeFire models. *J Combust* 2011:540390. <https://doi.org/10.1155/2011/540390>
7. Finney MA (1998) Farsite: fire area simulator-model development and evaluation. U.S. Department of Agriculture: Forest Service Research, Paper RMRS-RP-4
8. Generalitat de Catalunya, Departament d'Interior, GRAF: Informe de l'incendi de cardona (08-07-05) (2005) http://interior.gencat.cat/web/.content/home/030_arees_dactuacio/bombers/foc_forestal/consulta_incendis_forestals/informes_incendis_forestals/2000-2009/2005/20050708_REC_Cardona.pdf. Accessed 20 Feb 2018
9. Linn R, Reisner J, Colman JJ, Winterkamp J (2002) Studying wildfire behaviour using firetec. *Int J Wildland Fire* 11:233–246
10. Mandel J, Beezley JD, Kochanski aK (2011) Coupled atmosphere-wildland fire modeling with WRF 3.3 and SFIRE 2011. *Geosci Model Dev* 4(3):591–610 10.5194/gmd-4-591-2011
11. Mell W, Jenkins MA, Gould J, Cheney P (2007) A physics-based approach to modelling grassland fires. *Int J Wildland Fire* 16:1–22
12. Osher S, Piecher K, Fedkiw R (2004) The level set methods and dynamic implicit surfaces. *Appl Mech Rev* 57:B15
13. Rothermel RC (1972) A mathematical model for predicting fire spread in wild land fuels. USDA Forest Service Research Paper INT-115
14. Servei Meteorològic de Catalunya (METEOCAT): Official website. <http://www.meteo.cat>. Accessed 20 Feb 2018
15. Tymstra C, Bryce RW, Wotton BM, Taylor SW, Armitage, OB (2010) Development and structure of prometheus: the Canadian wildland fire growth simulation model. Information report NOR-X-417. ISBN 978-1-100-14674-4