

Faster Than Real Time Stochastic Fire Spread Simulations

A.R.Ervilha^{1,2}, F.A.Sousa¹, J.M.C.Pereira¹ and J.C.F.Pereira¹

Abstract: Faster than real time stochastic fire spread predictions are reported using a Non-Intrusive Spectral Projection (NISP) method based on Polynomial Chaos expansion and Graphic Processing Units (GPUs). The fireLib BEHAVE model together with a raster surface fire growth algorithm was implemented using the Compute Unified Device Architecture (CUDA) programming language. The uncertainty generated by the four random variables considered (wind speed, wind direction, fuel moisture, and fuel load) is quantified in the stochastic solution. Stochastic simulation of an idealized vegetation fire in a realistic complex terrain is obtained with speed-ups as high as 176 when compared to Central Processing Unit (CPU) and two orders of magnitude faster than real time fire propagation. The results include the fire front location and its error bar area, based on a 95% confidence interval, as well as temporal Probability Density Functions at selected points that quantify the uncertainty on the fire spread.

Keywords: Faster than real time; Graphic Processing Units; Non-intrusive spectral projection; Uncertainty quantification.

1 Introduction

Forest fires may be simulated using either deterministic or stochastic approaches depending on the purpose of the application. Fire spread predictions may be useful for firefighter training Kimmins, Blanco, Seely, Welham, and Scoullar (2010) and resource management such as detecting higher flammable zones, decision support, forest cleaning planning and to assess economic impacts towards structures fire risk. Additionally, if reliable much faster than real time predictions are obtained they may assist fire mitigation and fighting.

Deterministic models were thoroughly developed after the pioneering Rothermel (1972) work and extensive literature reviews exist on semi-empirical modeling,

¹ Instituto Superior Técnico, Technical University of Lisbon, Av. Rovisco Pais, 1, 1049-001, Lisbon, Portugal

² Email: eita.ervilha@ish.utl.pt

see e.g. Finney (1998), Perry (1998), Viegas (1998), Lopes, Cruz, and Viegas (2002), Pastor, Zárata, and Arnaldos (2003), Stratton (2006), and Papadopoulos and Pavlidou (2011).

The characterization of input variables is extremely important because they strongly affect forest fire simulations. Data bases of input parameters are scarce for real situations because their values are not constant neither in time nor in space. Consequently, model predictions have been mostly applicable to situations where the impact of possible errors is limited, such as training and fire management activities other than operational fire behavior prediction. Due to this uncertainty in the input parameters, namely in vegetation characteristics, wind fields and topographic conditions, it is highly recommended to simulate forest fire propagation with methods that consider the input variability in the model.

There is a lot of published work about stochastic and probabilistic fire modeling. Most encountered approaches are used in ecological and impact assessment studies which consider multiple fire events over long time periods. These studies address scenario simulation employing stochastic fire behavior and occurrence estimated from historical data Fried, Gilless, and Spero (2006), stochastic optimization methods to predict options that minimize the total cost of the wildfire [Ntaimo, Hu, and Sun (2008); Hu and Ntaimo (2009)] and ecological studies over several decades with fire occurrences introduced as probabilistic distributions based on past frequencies He and Mladenoff (1999). Furthermore, some models have been adjusted to account for random variables and mechanisms. Some techniques have been applied such as Cellular Automata (CA) with stochastic evolution rules Lichtenegger and Schappacher (2009), fire spread rate based on the probability functions of fire mechanism and input variables [Vorobov (1996); Boychuk, Braun, Kulperger, Krougly, and Stanford (2009)] and also conventional Monte Carlo techniques were used, for instance, in the elaboration of high resolution fire-risk maps Carmel, Paz, Jahashan, and Shoshany (2009).

All these works consider fire behavior to be intrinsically probabilistic and they are very little concerned about the way the uncertainty of the input data is propagated by forest fire spread models to the final forecast.

Knowing that it is not feasible to assume the input data to be 100% accurate, the uncertainty input in fire development should be quantified in the fire growth prediction. Stochastic Uncertainty Quantification (UQ) methods are able to effectively deliver this analysis resorting to a sample set of deterministic solutions obtained with the input's Probability Density Function (PDF). Several options are available for stochastic UQ as for instance Monte Carlo methods, but they are very computational demanding due to the large data sets required, namely for complex cases when the number of non-deterministic variables increases. Furthermore, Monte

Carlo methods do not allow for sensitivity analysis, this is, direct quantification of the uncertainty contribution of each input parameter in the solution. Non-Intrusive Spectral Projection (NISP) methods allow for this type of analysis while resorting to a much smaller sample set.

The NISP approach allows one to obtain a stochastic solution by evaluating the deterministic model for several samples of the uncertain input parameters (see e.g. Reagan, Najm, Ghanem, and Knio (2003); Xiu and Karniadakis (2003); Crestaux, Le Maître, and Martinez (2009); Mendes, Pereira, and Pereira (2010)). The solution is expanded in a series using Polynomial Chaos (PC) and the unknown coefficients of the terms are calculated with the solutions of the deterministic samples. Stochastic data solution, including error bars and PDFs, are then estimated by sampling the obtained solution function.

The implemented deterministic fire models are based on the semi-empirical Rothermel family of models, which are a validated and widely used approach to fire simulation. The fireLib Bevins (1996) version of the BEHAVE fire model Andrews (1986) was used as it is an easy to use and optimized Application Programming Interface (API), with freely available source code. Current models are able to produce a fire front solution from a large set of input data very quickly so that computational performance usually is not an issue. However, the great number of deterministic runs required by the UQ method raises constraints about computing time if the quasi-3D stochastic solution needs to be obtained much faster than real time. Our approach to overcome these constraints was by using Graphic Processing Units (GPUs).

GPUs have been used by the scientific community as soon as its capability to explore fine grained parallelism was perceived. Examples of those algorithms are Computational Fluid Dynamics (CFD) codes Tolke and Krafczyk (2008), Lattice Boltzman Habich (2008) and matrix solvers Bell and Garland (2008). Fine grained parallelism means that each task duration and data requirement is very small compared to the size of the problem. The semi-empirical fire models are very suitable for using with GPU and speed-ups up to two orders of magnitude can be achieved.

The argument for using the GPU can be made regarding other aspects such as price and availability, making the GPU a very attractive computing platform. Based on those advantages, Nvidia released the Compute Unified Device Architecture (CUDA) framework, which consists of a new GPU architecture that enables code developers to program generalized scientific software, using only a few extensions of the C++ programming language, see e.g. Kirk and Hwu (2009).

The implementation of the fire model in the GPU required the porting (code development with CUDA) of fireLib, plus a suitable Fire Growth Model (FGM). Math-

ematical models of fire behavior, such as BEHAVE, were originally designed for 1D fire propagation. Thus, the FGM function arises with the purpose of extending these models to a quasi-3D fire front propagation in a real complex terrain.

FGMs can be classified as being raster or vector. More details about FGM can be found in Anderson, Catchpole, Mestre, and Parkes (1982) and Finney (1998) for vector models and Lopes, Cruz, and Viegas (2002) and Perera, Ouellette, Cui, Drescher, and Boychuk (2008) for raster models. The decision to port to GPU the fireLib API and raster type models was based on several reasons. Data parallelism can be more easily explored as fire propagates between neighbouring raster cells in a local and independent interaction. Also, although vector models are more accurate [Cui and Perera (2008); Papadopoulos and Pavlidou (2011)], their implementation raises several challenges as the fire front crossovers and branches have to be properly resolved at each iteration. The implementation effort is beyond the scope of the work presented here. The GPU implementation of fireLib models is reported in Sousa, Reis, and Pereira (2012).

The purpose of this paper is therefore to present faster than real time stochastic calculations using a NISP method and the existing deterministic fire modeling tools. Stochastic forest fire predictions on GPUs may allow to obtain thousands of simulations and their statistical treatment during the fire development, becoming useful for firefighting purposes.

In the present work one has only considered as random variables the wind speed, the wind direction, the fuel load, and the fuel moisture, but other random variables including the whole set of 17 Rothermel variables can be considered. However, the selected variables are those that immediately affect the fire spread model and everybody knows that are not known with 100% confidence. Furthermore, the reported results, either the simplified benchmarks or the realistic complex terrain fire scenario, are physically consistent.

In the following section the deterministic fire models ported to the GPU are described, together with verification and speed performance, and the stochastic uncertainty quantification method is presented. The models are compared and the best option is selected for the next section where the stochastic UQ and the NISP model coupling with the deterministic FGM are used to simulate three test cases of stochastic fire prediction with several random variables. The paper closes with summary conclusions about stochastic fire spread simulation feasibility.

2 Methods

2.1 Fire Growth Models

FireLib functions are used to compute 1D fire Rate Of Spread (ROS) based on local variables, such as wind speed and direction, slope and aspect, and fuel properties like moisture and fuel load [Rothermel (1972); Albini (1976); Andrews (1986); Bevins (1996)]. To simulate a fire over a raster computational terrain, the data is provided to the model organized as 2D layers, where each aforementioned variable's spatial distribution is represented by a layer.

The FGM uses fireLib's ROS to compute the ignition times for each raster cell according to an elliptical shape propagation Anderson, Catchpole, Mestre, and Parkes (1982). The result is a map of ignition times for the whole terrain. Two raster FGMs were implemented and compared, namely the Contagious Procedure (CP) and the Cellular Automata (CA).

The CP FGM can be found for instance in the fire modeling tools BFOLDS Perera, Ouellette, Cui, Drescher, and Boychuk (2008) and FireStation Lopes, Cruz, and Viegas (2002). It consists in a contagious algorithm which computes the cells ignition time t_{ig} based on the propagation time between neighboring cells yielded by L/ROS , where L is the centroid distance between the two cells. Fig. 1 shows the two neighborhood configurations that are used with this method, namely the 16 and 8 cells stencil. The time progression is not made with a fixed step, instead the iteration time t^{n+1} is the minimum time of all the newly ignited cells. Therefore, oscillations in the time step Δt are likely to occur along the iterative process. The set of rules leading the system can be summed up as follows.

- For the time step n with simulation time t^n , the cell is burning if $t_{ig} = t^n$. For these cells, all vicinity cells ignition times t_{igV} are computed as $t_{igV} = t_{ig} + L_V/ROS_V$, where L_V and ROS_V are the centroid distance that separates the neighboring cell and the propagating speed in the same neighbor direction, respectively. This value is compared with the ignition times due to other surrounding burning cells, replacing it if it is lower. This is done in order to respect the fire arrival times in multiple fire ignition situations or concave fire fronts.
- If $t_{ig} < t^n$ the cell is considered burned and is ignored.
- The t^{n+1} value is updated so that it is the smallest ignition time after t^n , $\{\min(t_{ig} > t^n)\}$.

A CA is a numerical procedure that advances a given system by deterministically computing a succession of states, based on a set of rules applied to all cells of




Figure 1: Neighborhoods and propagation direction.

the domain. CA algorithms are known about their fine grained parallelism, hence the interest in pursuing a GPU implementation. The literature reports several fire prediction models with CA and the current implementation is based on the work of Karafyllidis and Thanailakis (1997) and Encinas, Encinas, White, Rey, and Sánchez (2007).

The state variable S is defined as the ratio between the cell burned area and the total area ratio $S = A_b/A_t$, and a fixed time step is used. The rules are applied to each cell considering the influence of the 8 surrounding neighbors (see Fig. 2). The set of state rules are as follows.

- $S = 0$ represents an unburned cell.
- If $S = 1$ the cell is completely burned.
- $S \in]0, 1[$ represents a fraction of the burned area.
- The new cell state is represented by a transition rule, function of the previous state of the neighbors and the same cell:

$$S^{n+1} = S^n + \sum_V^{adj} \frac{ROS_V}{\frac{L_V}{\Delta t}} S_V^m + \sum_V^{diag} \frac{\pi ROS_V^2}{4 \left(\frac{L_V}{\Delta t}\right)^2} S_V^m \quad (1)$$




Figure 2: Neighborhood and propagation direction for the CA FGM.

where *adj* and *diag* refer to the adjacent and diagonal neighbors, respectively, and V are the vicinity cells. The variables ROS_V and L_V are the rate of spread and the cell center point distance, respectively. The neighbor state variable follows the rule $S_V^n = 0$ if $S_V^n < 1$, and $S_V^n = 1$ otherwise (see Sousa, Reis, and Pereira (2012) for more details).

A considerable amount of work has been put into the GPU implementations. In Sousa, Reis, and Pereira (2012) an efficient CUDA version is implemented and an extensive coverage of the GPU porting is presented.

2.2 Verification and Speed Performance

The FGM implementation in the GPU can be verified in situations where the resulting fire front shape can be determined from the elliptical distribution computed by fireLib, which is true for constant condition scenarios. In these cases, the fire shape is exactly known by definition, and the resulting fire contours are explicitly yielded by fireLib.

The verification scenarios with constant parameters have terrain dimensions of 7×7 km² and are defined as follows.

- Scenario 1 corresponds to fuel model NFFL 1 Andrews (1986) with fuel moisture of 5%, flat terrain and no wind. A concentric progression of the fire front, starting from the ignition point, is to be expected with these conditions.

- Scenario 2 has the same fuel characteristics as the first scenario, but takes into account a uniform wind distribution of $2.3 \text{ m}\cdot\text{s}^{-1}$ with a 45° angle from North to South clockwise direction, slope value of 26.5 degrees with an aspect of 23° in the South West direction. The expected solution is now an ellipse tilted in the maximum spread rate direction.

Comparison between the burned area computed by the FGM and the exact elliptical shape Richards (1995) can be done with the Surface Error Index (*SEI*) as following:

$$SEI = 2 \frac{E \cup N}{E + N} \quad (2)$$

where E and N are the exact and numerical solutions of the ignition maps, respectively. For Scenario 1, the *SEI* obtained are 0.986, 0.948 and 0.976 for CP16, CP8 and CA models, respectively. Scenario 2 presents a model performance of 0.915, 0.690 and 0.772 for CP16, CP8 and CA models, respectively.

The *SEI* quantifies the error introduced by the FGM. A $SEI < 1$ means that the model is deviating from the elliptical shape predicted by the mathematical model for fire behavior built in fireLib. One can see that the model performance is related to the number of neighbors and the 16 cells CP model displays better results through both scenarios. This happens because the 16 cells stencil has more directions for the fire to propagate than the 8 cells one. With the increase of the number of neighbor directions in the stencil, the *SEI* would increase up to the limit of 1 with an infinite number of directions. The *SEI*'s results obtained for the FGM are expected according to the literature survey (see e.g. Cui and Perera (2008)).

Speed performance is assessed using real topographic data for the simulation of a fire that propagates during 8 days and it may be representative of the computing times of real situations. Tab. 1 lists the total computing time on the CPU and GPU as well as the speed-up, calculated as the ratio between CPU and GPU times, for three different grids.

The results show that all FGM satisfies the faster than real time condition, even with a computational domain comprising 4 million cells grid, which is a grid size far bigger than the usually employed in this kind of fire simulation. This particular calculation took 34 CPU hours to simulate an eight day real time fire.

Although the CPU simulation time is still faster than real time, the problem here lies in the feasibility to perform stochastic predictions also faster than real time. The excellent speed-up obtained with GPU implementations are very promising regarding the proposed objective. Two orders of magnitude in speed-up were obtained and the CP models show very promising results. The CA implementation

Table 1: Speed performance summary.

Model	Grid [M cells]	Total CPU time [s]	Total GPU time [s]	Speed-up
CP 16	0.26	139.24	0.79	174.1
CP 16	1.05	1171.24	6.33	184.7
CP 16	4.2	8541.81	49.91	171.0
CP 8	0.26	82.86	0.54	153.4
CP 8	1.05	644.09	3.96	162.2
CP 8	4.2	5090.03	30.84	164.9
CA	0.26	1598.67	49.08	32.6
CA	1.05	14346.7	341.85	42
CA	4.2	123776.0	2894.78	42.8

also delivered a good increase in performance but the overall computing times are bigger than the CP ones. Consequently, the 16 cells CP is the natural selection to use with the UQ method due to the verification results and computation times.

2.3 Stochastic Uncertainty Quantification Model

Uncertainty quantification in the fire front is done using Non-Intrusive Spectral Projection (NISP) methods, where the stochastic solution is built given the PDF of the input variables. The method assures that a solution variable can be correctly described from a set of input samples much smaller than with Monte Carlo family methods. Fig. 3 describes the process and the coupling between the FGM models and the NISP procedure.

As shown in Fig. 3, the present model comprises two sub-models and a parametric uncertainty methodology. The first sub-model is the fire spread model, that in the present case considers the fireLib and a fire growth model. The second sub-model comprises the stochastic model Non-Intrusive Spectral Projection method that is faster than Monte Carlo, see e.g. Ervilha, Pereira, and Pereira (2013), and allows for uncertainty propagation.

The NISP module is divided into two parts. The first one is responsible for building the deterministic input samples from the input variables PDF. Next, several FGM runs are issued at the same time to build a solution set from input samples. The resulting ignition maps are processed to compute the stochastic ignition time variable for each point in the map. A post-processing stage follows, in order to compute the PDF solution and the confidence intervals for the fire front position.

With the NISP approach a stochastic solution can be obtained by evaluating the deterministic model for several samples of the uncertain input parameters. The




Figure 3: Overall view of the coupling between the deterministic FGM runs and the NISP method.

solution is decomposed in a series using Polynomial Chaos (PC) expansion, see e.g. Xiu and Karniadakis (2003) and Crestaux, Le Maître, and Martinez (2009), and the unknown coefficients of the terms are calculated using the input samples. The solution PDF is then estimated by sampling the obtained solution polynomial. A random variable X can be described as a function of a random standard variable ξ in a PC expansion:

$$X(\xi) = \sum_{j=0}^p a_j^X \Phi_j(\xi) \tag{3}$$

where a_j^X are known coefficients and $\Phi_j, j = 0, \dots, p$, are orthogonal polynomials of order j . In the present study, ξ is a standard normal distribution $N(0, 1)$ and Φ_j are Hermite polynomials. There can also be N independent random parameters (X_1, \dots, X_N) , each one associated to a stochastic dimension $\xi_i, i = 0, \dots, N$, and the orthogonal polynomials would now be function of a vector of random variables $\vec{\xi} = (\xi_1, \dots, \xi_N)$. Because the polynomials are orthogonal, the inner product between

any two is $\langle \Phi_i, \Phi_j \rangle = 0$ if $i \neq j$ and

$$\langle \Phi_i, \Phi_j \rangle = \int \Phi_i \Phi_j W(\vec{\xi}) d\vec{\xi} = \langle \Phi_j^2 \rangle \delta_{ij} \tag{4}$$

where δ_{ij} is the Kronecker delta function, $W(\vec{\xi}) = w(\xi_1) \cdots w(\xi_N)$ is the weighting function of the polynomial $\{\Phi_j\}$, and $w(\xi_i), \forall i \in \{1, \dots, N\}$, is

$$w(\xi_i) = \frac{\exp\left(-\frac{\xi_i^2}{2}\right)}{\sqrt{2\pi}} \tag{5}$$

The solution variable can also be expanded in a multi-dimensional PC series as a function of the vector of random variables. Depending on the number of random input variables N and the maximum polynomial order p of the expansion, the number of terms in the polynomial is given by:

$$P + 1 = \frac{(N + p)!}{N!p!} \tag{6}$$

To better understand the PC expansion used in this work, one will describe this method for the first study case presented in the next section, where 2 stochastic input variables are considered, namely the wind speed U and the wind direction α . So, the stochastic ignition time t_{ig} is obtained through a PC series as a function of those random variables associated to the stochastic dimensions ξ_U and ξ_α , respectively. For a polynomial expansion of 2^{nd} order, the ignition time t_{ig} is:

$$t_{ig}(\vec{\xi}) = \sum_{j=0}^5 a_j^{t_{ig}} \Phi_j(\vec{\xi}) \tag{7}$$

where $a_j^{t_{ig}}, j = 0, \dots, 5$, are now unknown coefficients that must be computed as functions of the input parameters and $\Phi_j(\vec{\xi}), j = 0, \dots, 5$, are the Hermite orthogonal polynomials of order j . For this first study case, the expansion takes the concrete form:

$$t_{ig}(\xi_U, \xi_\alpha) = a_0^{t_{ig}} + a_1^{t_{ig}} \Phi_1(\xi_U) + a_2^{t_{ig}} \Phi_1(\xi_\alpha) + a_3^{t_{ig}} \Phi_2(\xi_U) + a_4^{t_{ig}} \Phi_1(\xi_U) \Phi_1(\xi_\alpha) + a_5^{t_{ig}} \Phi_2(\xi_\alpha)$$

(8)

where $\Phi_k(\vec{\xi}), k = 0, \dots, 2$, are the orthogonal polynomials belonging to the Hermite polynomial basis $\{\Phi_k\}$ and obtained as functions of the vector of random variables $\vec{\xi} = (\xi_U, \xi_\alpha)$.

To calculate the coefficients, a Galerkin projection of the PC in the polynomial basis $\{\Phi_j\}$ is made by multiplying each side of Eq. 7 by Φ_j and using the orthogonality relation of Eq. 4:

$$a_j^{t_{ig}} = \frac{\langle t_{ig}(\vec{\xi})\Phi_j \rangle}{\langle \Phi_j^2 \rangle} = \frac{\int t_{ig}(\vec{\xi})\Phi_j W(\vec{\xi}) d\vec{\xi}}{\langle \Phi_j^2 \rangle}, j = 0, \dots, P \tag{9}$$

The $a_j^{t_{ig}}$ coefficients are now written as a function of the solution variable and can be obtained in the NISP approach by evaluating the deterministic solutions $\{(t_{ig,d})^n\}_{n=1}^S$ as a function of the chosen input variables samples $\{(U, \alpha)^n\}_{n=1}^S$, where S are selected points. Finally, the integral in Eq. 9 is numerically solved by Gauss-Hermite quadrature. The number of selected samples are defined by the number of Gauss-Hermite quadrature points S_i for each variable $\xi_i, \forall i \in \{1, \dots, N\}$ and $(z_{r_i}, q_{r_i}), r = 1, \dots, S_i$, are the quadrature points and weights. For the multi-dimensional case there will be $S = \prod_{i=1}^N S_i$ number of samples. The Gauss quadrature rule yields an exact result with S_i collocation points for a polynomial's degree up to $2S_i - 1$. The integral on Eq. 9 has presently be considered to have a maximum degree of $2p$, because the PC expansion of $t_{ig}(\vec{\xi})$ has at most a degree of p so, the minimum Gauss collocation points can be related to the PC degree by $S_i \approx p + 1/2$. The coefficients in the above equation can now be presented in the form of a numerical expression:

$$a_j^{t_{ig}} \approx \frac{\sum_{r_1, \dots, r_N=1}^{S_1, \dots, S_N} t_{ig,d}(X_{r_1}, \dots, X_{r_N})\Phi_j(z_{r_1}, \dots, z_{r_N}) \prod_{i=1}^N q_{r_i}}{\langle \Phi_j^2 \rangle}, j = 0, \dots, P \tag{10}$$

The stochastic solution variable is now completely known with this method. A PDF of any dependent variable solution can be obtained and the characteristics mean value $\mu_{t_{ig}}$ and variance $\sigma_{t_{ig}}^2$ are given by Eq. 11 and Eq. 12, respectively.

$$\mu_{t_{ig}} = \langle t_{ig} \rangle = a_0^{t_{ig}} \tag{11}$$

$$\sigma_{t_{ig}}^2 = \langle t_{ig}^2 \rangle - \langle t_{ig} \rangle^2 = \sum_{j=1}^P (a_j^{t_{ig}})^2 \langle \Phi_j^2 \rangle \quad (12)$$

3 Results

3.1 Stochastic prediction results

This section presents results for several sets of random input variables (wind speed, wind direction, fuel load, and fuel moisture) to study the influence of the parametric uncertainty in the final solution. Other variables like terrain topography are considered to be exact. The random input variables have a Gaussian probability distribution, described by a mean (μ) and a standard deviation (σ).

The studied cases are divided into 3 main groups, namely cases A, B and C. Each group has different sets of random variables that are introduced in the model in order to carry out a sensitivity analysis study. For case group A, only wind speed (U) and wind direction (α) are random variables. In the second case group, the fuel data is considered uncertain and the wind characteristics are now assumed to be exact, which means that, for group B, the stochastic variables are fuel moisture (M) and fuel load (ω_0). Group C treats all the four inputs U , α , M and ω_0 as random variables. Additionally, the effect of terrain on the stochastic forecast is studied in C by comparing two cases with different orography (zero slope and complex terrain (see Fig. 4), respectively) under the same input conditions. All cases in A and B have zero slope terrain. The non random input parameters used in simulations are listed in Tab. 2.




Figure 4: Three dimensional rendering of the C2 case terrain, with elevation (in meters) displayed in gray scale. The upper side is pointed to north.

Table 2: Deterministic parameters used in the simulations.

Input variable	Deterministic value
Fuel bed depth	0.06 m
Moisture of extinction	12%
Fuel surface-area-to-volume ratio	11483 m ² .m ⁻³
Fuel density	513 kg.m ⁻³
Heat of combustion	18608 kJ.kg ⁻¹
Total silica content	0.01
Effective silica content	0.056

For all groups, the mean values (μ) of the random input variables considered are $\mu_U = 1.2 \text{ m.s}^{-1}$, $\mu_\alpha = 135^\circ$ clockwise from North, $\mu_{\omega_0} = 1.123 \text{ kg.m}^{-2}$, and $\mu_M = 7\%$. The parameters variability for each case is presented in Tab. 3 through the coefficient of variation $CV = \sigma/\mu$.

Table 3: Coefficient of variation considered for the input random variables used in group A (cases A1, A2 and A3), group B (cases B1, B2 and B3), and case group C.

Input random variable	A1	A2	A3	B1	B2	B3	C
U	20%	40%	20%	0%	0%	0%	20%
ω_0	10%	10%	20%	0%	0%	0%	10%
α	0%	0%	0%	15%	30%	15%	15%
M	0%	0%	0%	10%	10%	20%	10%

The studied cases considered have a computational domain corresponding to a terrain with $3 \times 3 \text{ km}^2$ and a flat terrain is used except for case C2. The location of the initial single ignition point is the same for all the studied cases.

The stochastic forecasts of the fire front evolution are listed in Tab. 4. The data correspond to the one hour burned area obtained for deterministic prediction, in which input variables have constant values, and stochastic simulation (ensemble mean and lower and upper limits depending on the Confidence Intervals (CI)). The last column refers to an adimensional factor (error factor) that represents the ratio between the burnt area's error bar (this is, the burned area between lower and upper limits) and the stochastic mean burned area.

The predictions display a stochastic mean burned area always smaller than the deterministic one. Although all random parameters are characterized by a symmetric PDF, the randomness influence originates a delay of the fire propagation.

Table 4: One hour forecast of burned area for 95% and 60% confidence intervals for all the studied cases.

Case study	CI [%]	Deterministic mean [km ²]	Stochastic mean [km ²]	Lower limit [km ²]	Upper limit [km ²]	Error area ^(a)	Error factor ^(b)
A1	95	0.21	0.17	0.05	0.66	0.61	3.59
A1	60	0.21	0.17	0.10	0.34	0.24	1.41
A2	95	0.21	0.13	0.03	1.74	1.71	13.15
A2	60	0.21	0.13	0.07	0.48	0.41	3.15
A3	95	0.21	0.12	0.03	0.89	0.86	7.17
A3	60	0.21	0.12	0.07	0.34	0.27	2.25
B1	95	0.17	0.16	0.13	0.19	0.06	0.38
B1	60	0.17	0.16	0.15	0.18	0.03	0.19
B2	95	0.17	0.16	0.13	0.19	0.06	0.38
B2	60	0.17	0.16	0.15	0.18	0.03	0.19
B3	95	0.17	0.15	0.07	0.22	0.15	1.00
B3	60	0.17	0.15	0.12	0.21	0.09	0.60
C1	95	0.17	0.13	0.04	0.52	0.48	3.69
C1	60	0.17	0.13	0.08	0.27	0.19	1.46
C2	95	0.13	0.11	0.04	0.40	0.36	3.27
C2	60	0.13	0.11	0.07	0.20	0.13	1.18

^(a)error area = upper limit - lower limit

^(b)error factor = $\frac{\text{error area}}{\text{stochastic mean}}$

Focusing the analysis on the flat terrain cases (A, B and C1), the error factor allows to assess the impact of each random variable on the ignition times of fire propagation. Cases A and C1 are characterized by higher values of the error factor, denoting the great impact of the wind properties variability in comparison with fuel variability (B cases). The area's error bar for the cases A and C1 are up to 10 times the stochastic mean burnt area, while the area's error bar in B cases is approximately equal to the stochastic mean value.

The effect of the variability of each random parameter show that for case A the effect of the wind velocity variability is higher than the wind direction variability and Tab. 4 lists the increase of the error factor when σ_U rises from 20% to 40% (case A2) over the increase of σ_α from 10% to 20% (case A3). This value is almost 4 times bigger for A2 and 2 times bigger for A3, when comparing to A1. The results for the B cases show that moisture variability has a far greater impact than the fuel load's. When the later increases from $\sigma_{\omega_0} = 15\%$ in B1 to 30% in B2, the difference

in error factor is smaller than 1%. In comparison, increasing σ_M from 10% to 20% in case B3 results in an error factor twice the size of B1's. C1 simulation yields similar results than A1 for error factor, indicating the smaller influence of the fuel properties variability over the wind properties. This is consistent with A and B findings.

The C2 simulation introduces a realistic terrain (under C1 initial conditions) to quantify the influence of the orography in fire propagation. Fig. 4 shows a 3D rendering of the 3 by 3 km² terrain, where the arrow points to the initial ignition spot.

Fig. 5(b) shows the stochastic mean for the fire front position plotted in one hour intervals, presenting a slower backfire region (temporal isolines very close) and increasing speed when going upslope (larger distance between temporal isolines).




Figure 5: Case C2. (a) Aspect map as a reference for the forecast simulations. (b) Stochastic mean for the fire front plotted with 1 hour intervals.

Fig. 6(a) and Fig. 6(b) show the ensemble mean together with the error bar with upper and lower fire locations. The 60 minutes predictions correspond to 95% and 60% CI, respectively. The stochastic mean average burned area and the deterministic counterpart can be found in Tab. 4. Fig. 6(c) shows the 90 minutes forecast for a 95% CI. The error area rises relatively to the one hour plot, with a 12% increase in the error factor. For the cases with flat terrain, this difference does not occur, i.e. the error factor is practically the same along the simulation time.

A relevant output of faster than real time stochastic prediction is the error bar of the fire front line according to a confidence interval, usually 95%. The error bar gives upper and lower locations of the fire line. The model mimics the physical and ecophysiological processes that happen in a wildfire and their influence in the fire




Figure 6: Forecast of the fire front. Stochastic mean is in dashed, deterministic mean is in dot dashed, and CI are the solid lines.

spread because the uncertainty in almost all the variables is the only thing that is certain.

The statistical output of the predicted fields allows to know the probability of fire occurrence at a selected point and to decide accordingly about actions to have in real wildfires.

Stochastic fire forecast allows the quantification of the fire front occurrence probability at any location, as a function of time. This becomes particularly interesting when specific locations are of utmost importance and require priority protection, such as personnel and equipment staging areas or building structures. The fire behavior is studied for two different points, namely P1 and P2, located at (1.13 km, 1.87 km) and (0.75 km, 1.5 km), respectively (see Fig. 5(a)).

Tab. 5 and Fig. 7(a), Fig. 7(b) and Fig. 7(c) present the stochastic information

and the normalized PDF obtained for both points in all studied cases, respectively. Apart from the stochastic mean, one can obtain other stochastic information. The standard deviation is obtained through the square root of variance (second order moments) and represents the range values considered, being low if data points tend to be close to the mean and high for disperse points. The third order moment (skewness) is a measure of the PDF symmetry, acquiring positive values for negative bias and negative values otherwise. Finally, kurtosis is a fourth standardized moment and it is representative of the PDF shape. The PDF translates the probability of fire ignition at the point under consideration for a certain time (time coordinate axis). The current work presents normalized PDFs and, consequently, the integral of each PDF over time is equal to one. This means that one can easily know what is the probability of a certain ignition time at a specific location directly through the percentage of the PDF area until that considered time.

Table 5: Mean, standard deviation (*sd*), skewness (*skew*) and kurtosis (*kurt*) of the ignition times PDFs at points P1 and P2.

Case study	P1 mean	P1 sd	P1 skew	P1 kurt	P2 mean	P2 sd	P2 skew	P2 kurt
A1	59.23	21.35	1.12	1.71	169.57	60.39	0.78	0.83
A2	70.84	46.94	1.44	2.85	184.80	88.13	0.88	1.07
A3	74.89	35.49	1.69	5.09	185.12	99.25	1.23	2.06
B1	56.35	3.12	0.96	1.25	178.64	9.88	0.96	1.25
B2	56.44	3.13	0.95	1.23	178.93	9.93	0.96	1.23
B3	58.61	11.46	2.07	5.99	185.80	36.33	2.07	5.99
C1	67.05	24.47	1.12	1.72	191.83	69.29	0.80	0.88
C2	73.64	30.81	1.38	2.67	236.31	77.49	0.71	0.70

As expected, for all cases the stochastic mean ignition time for P1 is always smaller than for P2 because P1 is located on the direction of the maximum spread propagation. Consequently, the fire front is expected to arrive at P1 before P2. This fact is responsible for the location of the PDF for P1, which is shifted into the left in Fig. 7(a), Fig. 7(b) and Fig. 7(c).

The plots for the A group (Fig. 7(a)) show that the standard deviation is much smaller in P1 than in P2. This is due to the ignition point proximity of P1. For instance, the results for case A1 show that at point P1 after 50 minutes there is around 50% probability that the fire front reach this particularly point and after 150 minutes the probability is 100%. The increase in the number of random input variables translates in an increase of the standard deviation of the PDF, although some variables have higher impact than others, which is consistent with previous




Figure 7: PDF of ignition times in P1 and P2.

findings. Notice that the standard deviation for case A2 at location P1 is more than twice that for case A1 (see Tab. 5). Additionally, at location P2 the wind speed variability has an impact similar to that of wind direction. This shows that the influence of wind speed variability is higher at location P1, which corresponds to P1 being in the direction of maximum fire propagation.

The PDFs of B cases are shown in Fig. 7(b). In agreement with Tab. 4, B2 yields the same result for stochastic mean as B1. Both points display much narrow PDF than in the A cases, again in accordance to the burned area results of Tab. 4. Relatively to the B3, the PDF curve is shaped after a time step, which means that the highest probability of ignition time is exactly in this slice and then the probability decreases over time. The wide variation in the PDFs shape in B3 case translates the highest influence of moisture fuel variance over the fuel load distribution. Notice that for all B group cases, both ensemble mean and standard deviation in P2 are about three

times bigger than in P1, while symmetry and shape factors are similar. So, beyond their different location that implies obviously different mean time ignition values, both points reflect the same fire front behavior but with sparse results.

Case C2 (Fig. 7(c)) quantifies the influence of real terrain on the simulations, which delays or speeds fire propagation as the topographic features are encountered. However, given the terrain influence, the comparison between different location point's PDF is not feasible to be done and important conclusions come mainly from fire front forecast plots (as for instance Fig. 6), which provide stochastic mean contours and error bar burned area according to the considered CI. At each location it is possible to determine the probability of fire occurrence as a function of time. For example, in Fig. 7(c), the PDF shows that the fire front takes always more than 50 min to reach point P2. Then, the probability of ignition rises until 210 min, after which it decreases to nearly zero at 500 min.

3.2 Parametric Uncertainty Quantification

The present results advance the knowledge of wildfire spreading simulation because they allow the uncertainty quantification on the fire spread. The influence of a random variable's uncertainty can be accurately quantified by comparing the stochastic coefficients (see Eq. 8), allowing a sensitivity analysis of the influence of each variable into the stochastic results.

For instance, taking the ignition time solution under consideration with wind speed and wind direction as the random input parameters, Eq. 8 shows the expansion solution as a function of first and second order stochastic coefficients. The coefficients $a_1^{t_{ig}}$ and $a_2^{t_{ig}}$ denote first order wind speed and first order wind direction, respectively, while $a_3^{t_{ig}}$ translates the second order wind speed influence, $a_4^{t_{ig}}$ contains the cross influence of wind speed and direction and, finally, $a_5^{t_{ig}}$ stands for the second order wind direction. One should note that other solution random parameters has different but similar expansions.

Fig. 8(a) and Fig. 8(b) show the stochastic coefficients obtained for the group A scenario at P1 and P2, respectively. The small influence of the first order wind direction coefficient (α) at location P1 is translated into a very small influence of this parameter variability into the stochastic ignition time results. This is justified because P1 is in the maximum propagation direction and the variability induced is symmetric. However, the second order coefficient of α has a considerable value at this location for case A3. Fig. 8(b) shows that the coefficients related to α for P2 have the expected large values, contributing to a decrease in the ignition time result.

Taking as random input variables the fuel load (ω_0) and the fuel moisture (M), the





Figure 8: Stochastic coefficients for case group A.


expansion coefficients show that P1 and P2 have similar coefficient behavior, see Fig. 9(a) and Fig. 9(b), indicating that uncertainty propagation pattern is not dependent of the location. Fig. 9(a) and Fig. 9(b) show that the fuel load’s variability has a negligible impact on the solution, even with a twofold increase in σ_{ω_0} from B1 to B2. The moisture fuel variability presents a small impact in the stochastic mean of time ignition simulations, but its contribution is mostly to the solution’s variability and therefore notable in the PDF development (see Fig. 7(b)).

For the group C with four random variables, the PC expansion (see Eq. 7) generates 15 stochastic coefficients, including the stochastic mean represented by the first term $a_0^{t_{ig}}$. The findings for the case C1 (flat terrain) are consistent with A and B results because the fuel load has a negligible contribution as well as the wind direction for P1 (see Fig. 10(a)). When the complex terrain is considered (case C2), the wind direction parameter is no longer negligible in the stochastic solution for this point. At P2 location (see Fig. 10(b)) both cases C1 and C2 show that the main influence in forest fire propagation is from wind features (speed and direction) over the fuel ones (load and moisture).

The present uncertainty propagation analysis can be applied for any other scenarios and point locations. This is of special relevance under firefighting mode because it will allow to locate the virtual sensor at, for example, the urban-wildland interface and to quantify the uncertainty (variability of the parametric input random variables) on the ignition time at that location as well as the parametric variability hierarchy.



(a) P1 (b) P2
Figure 9: Stochastic coefficients for case group B.



(a) P1 (b) P2
Figure 10: Stochastic coefficients for case group C.

The limitations of the present stochastic model are inherent to the sub-models related with fire spread and also the assumption of the input random parameter variability. However, the present NISP model may be used with other fuel spread models and for a better guess of the meteorological conditions wind forecast models can be used.

The model could be improved by allowing a discontinuous fire front line as an input parameter uncertainty. For large fires, the variability of the fire front line can be captured, or will be captured soon, using optical diagnostics sensors, unmanned vehicles, or satellite pictures, almost as in Google street view (see <https://maps.google.com/>).

3.3 Speed-up

The stochastic calculations considered 5 Gauss-Hermite quadrature points for each stochastic variable. So, for stochastic calculations with 4 random inputs a total number of 625 (5^4) different deterministic runs is required. Each deterministic simulation calculates the fire spread during 8 hours of real time evolution.

The CP16 model was computed on the GPU and the computing time is listed in Tab. 1 and Tab. 6. The CPU time to provide the 625 required runs is not faster than real time and, consequently, not feasible to stochastic fire predictions. Moreover, a total computing time would be far greater with Monte Carlo calculation instead of the presented NISP model. The GPU results are much faster than real time, 60 times faster, providing the calculations of 625 runs being obtained with a speed-up of 176 times the CPU results and, consequently, the GPU and NISP method are suitable for the purpose of a decision support tool in firefighting context.

Table 6: Comparison between CPU and GPU speed performances.

GPU [s]	CPU [s]	$\frac{\text{CPU}}{\text{GPU}}$	$\frac{\text{CPU}}{\text{real time}}$	$\frac{\text{GPU}}{\text{real time}}$
493.75	87025	176.25	3.02	0.02

4 Conclusions

The paper presented details of a stochastic simulation of vegetation fires. The calculations were obtained using a Non-Intrusive Spectral Projection method in which the solution is expanded in a series using Polynomial Chaos. The unknown coefficients of the expansion terms are calculated from deterministic solutions using a conventional fire growth model. In the present case the fireLib functions were used for the calculation of the rate of spread together with raster contagious or cellular automata algorithms.

To decrease the computing time required for hundreds or thousands of deterministic simulations over, typically, one hour real fire propagation, the low cost Graphics Processing Unit (GPUs) were used.

The stochastic simulations of idealized vegetation fires with input parametric uncertainty involving wind speed, wind direction, fuel load, and fuel moisture as random

variables were simulated under flat and complex realistic terrain. The following conclusions can be withdrawn.

1. The GPU have allowed to obtained much faster stochastic fire spread predictions than the real time fire propagation.
2. Under flat or complex terrain the wind speed and wind direction uncertainties have stronger influence on the fire propagation than the fuel contents. The relative influence of the uncertainties were quantified.
3. The proposed methodology can be used for any scenario. The output provides the time evolution of the ensemble mean fire front location and burned areas error bar for a certain confidence interval and the probability density functions at each point as a function of time. In addition, the hierarchy of input parametric uncertainties in the fire spread simulation was quantified. These estimators may add relevant information because wildland fire growth is an intrinsic stochastic process.

Acknowledgement: The authors would like to thank the support from Project EXTREME - PTDC/EME-MFE/114343/2009. The first and third authors would like to acknowledge the fellowships received from Fundação para a Ciência e Tecnologia (FCT).

References

- Albini, F. A.** (1976): *Estimating wildfire behaviour and effects*. USDA Forest Service, International Mountain and Range Experiment Station, Gen. Tech. Rep. INT-GTR-30, Ogden, UT, 1976.
- Anderson, D. H.; Catchpole, E. A.; Mestre, N. J. D.; Parkes, T.** (1982): Modelling the spread of grass fires. *ANZIAM J.*, vol. 23, no. 4, pp. 451–466.
- Andrews, P. L.** (1986): *BEHAVE: Fire behaviour prediction and fuel modelling system - burn subsystem, part1*. USDA Forest Service, International Mountain Forest and Range Experiment Station, Gen. Tech. Rep. INT-194, Ogden, UT, 1986.
- Bell, N.; Garland, M.** (2008): Efficient sparse matrix-vector multiplication on cuda (nvidia). Technical Report NVR-2008-004, NVIDIA Corporation, 2008.
- Bevins, C. D.** (1996): *fireLib User Manual and Technical Reference. Systems for Environmental Management*, 1996.
- Boychuk, D.; Braun, W.; Kulperger, R.; Krougly, Z.; Stanford, D.** (2009): A stochastic forest fire growth model. *Environ. and Ecol. Stat.*, vol. 16, pp. 133–151–.

- Carmel, Y.; Paz, S.; Jahashan, F.; Shoshany, M.** (2009): Assessing fire risk using monte carlo simulations of fire spread. *For. Ecol. and Manag.*, vol. 257, pp. 370–377.
- Crestaux, T.; Le Maître, O.; Martinez, J.** (2009): Polynomial chaos expansion for sensitivity analysis. *Reliab. Eng. & Syst. Saf.*, vol. 94, no. 7, pp. 1161–1172.
- Cui, W.; Perera, A. H.** (2008): *A study of simulation errors caused by algorithms of forest fire growth models*. Ontario Forest Research Institute, Sault Ste. Marie, ON. Forest Research Report No. 167, 2008.
- Encinas, A. H.; Encinas, L. H.; White, S. H.; Rey, A. M.; Sánchez, G. R.** (2007): Simulation of forest fire fronts using cellular automata. *Adv. in Eng. Softw.*, vol. 38, no. 6, pp. 372–378.
- Ervilha, A. R.; Pereira, J. M. C.; Pereira, J. C. F.** (2013): On the parametric uncertainty quantification of the Rothermel's rate of spread model. *Submitted to Environmental Modelling & Software*. (unpublished data).
- Finney, M.** (1998): *FARSITE: Fire Area Simulator - Model Development and Evaluation*. USDA Forest Service, Rocky Mountain Research Station, Res. Pap. RMRS-RP-4 Revised, Ogden, UT, 1998.
- Fried, J. S.; Gilless, J. K.; Spero, J.** (2006): Analysing initial attack on wildland fires using stochastic simulation. *Int. J. of Wildland Fire*, vol. 15, pp. 137–146.
- Habich, J.** (2008): Performance evaluation of numeric compute kernels on nvidia gpus. Master's thesis, Friedrich Alexander Universität, 2008.
- He, H. S.; Mladenoff, D. J.** (1999): Spatially explicit and stochastic simulation of forest landscape fire disturbance and succession. *Ecol.*, vol. 80, no. 1, pp. 81–99.
- Hu, X.; Ntaimo, L.** (2009): Integrated simulation and optimization for wildfire containment. *ACM Trans. Model. Comput. Simul.*, vol. 19, no. 4.
- Karafyllidis, I.; Thanailakis, A.** (1997): A model for predicting forest fire spreading using cellular automata. *Ecol. Model.*, vol. 99, no. 1, pp. 87–97.
- Kimmins, J. P.; Blanco, J. A.; Seely, B.; Welham, C.; Scoullar, K.** (2010): *Forecasting forest futures: a hybrid modelling approach to the assessment of sustainability of forest ecosystems and their values*. Earthscan Ltd., London, UK.
- Kirk, D.; Hwu, W.** (2009): *Programming massively parallel processors: A hands-on approach*. Elsevier Morgan Kaufmann.
- Lichtenegger, K.; Schappacher, W.** (2009): Phase transition in a stochastic forest fire model and effects of the definition of neighbourhood. *Int. J. of Mod. Phys. C*, vol. 20, no. 8, pp. 1247–1269.

Lopes, A.; Cruz, M.; Viegas, D. (2002): Firestation - an integrated software system for the numerical simulation of fire spread on complex topography. *Environ. Model. and Soft.*, vol. 17, no. 3, pp. 269–285.

Mendes, M.; Pereira, J. M. C.; Pereira, J. C. F. (2010): Calculation of pre-mixed combustion within inert porous media with model parametric uncertainty quantification. *Combust. and Flame*, vol. 158, no. 3, pp. 466–476.

Ntaimo, L.; Hu, X.; Sun, Y. (2008): Devs-fire: Towards an integrated simulation environment for surface wildfire spread and containment. *Simul.: Trans. of the Soc. for Model. and Simul. Int.*, vol. 84, pp. 137–155.

Papadopoulos, G. D.; Pavlidou, F.-N. (2011): A comparative review on wildfire simulators. *IEEE Syst. J.*, vol. 5, no. 2, pp. 233–243.

Pastor, E.; Zárate, L.; Arnaldos, E. P. J. (2003): Mathematical models and calculation systems for the study of wildland fire behaviour. *Prog. in Energy and Combust. Sci.*, vol. 29, no. 2, pp. 139–153.

Perera, A.; Ouellette, M.; Cui, W.; Drescher, M.; Boychuk, D. (2008): *BFOLDS 1.0: A spatial simulation model for exploring large scale fire regimes and succession in boreal forest landscapes*. Ontario Forest Research Institute, 2008.

Perry, G. (1998): Current approaches to modelling the spread of wild land fire: a review. *Prog. in Phys. Geogr.*, vol. 22, pp. 222–224.

Reagan, M.; Najm, H.; Ghanem, R.; Knio, O. (2003): Uncertainty quantification in reacting-flow simulations through non-intrusive spectral projection. *Combust. and Flame*, vol. 132, pp. 545–555.

Richards, G. D. (1995): A general mathematical framework for modelling two-dimensional wildland fire spread. *Int. J. of Wildland Fire*, vol. 5, no. 2, pp. 63–72.

Rothermel, R. C. (1972): *A mathematical model for predicting fire spread in wildland fuels*. USDA Forest Service, International Forest and Range Experiment Station, Res. Pap. INT-115, Ogden, UT, 1972.

Sousa, F.; Reis, R.; Pereira, J. C. F. (2012): Simulation of surface fire fronts using firelib and cplus. *Environ. Model. Soft.*, vol. 38, pp. 167–177.

Stratton, R. D. (2006): *Guidance on spatial wildland fire analysis: models, tools, and techniques*. USDA Forest Service, Rocky Mountain Research Station, Gen. Tech. Rep. RMRS-GTR-183, Ogden, UT, 2006.

Tolke, J.; Krafczyk, M. (2008): Teraflop computing on a desktop pc with gpus for 3d cfd. *Int. J. Comput. Fluid Dyn.*, vol. 22, no. 7, pp. 443–456.

Viegas, D. X. (1998): Forest fire propagation. *Philos. Trans. of the R. Soc. of London, Series A: Math., Phys. and Eng. Sci.*, vol. 356, no. 1748, pp. 2907–2928.

Vorobov, O. (1996): Random set models of fire spread. *Fire Technol.*, vol. 32, no. 2, pp. 137–173.

Xiu, D.; Karniadakis, G. (2003): Modelling uncertainty in flow simulations via generalized polynomial chaos. *J. of Comput. Phys.*, vol. 187, pp. 137–167.

