ISEIS

Journal of Environmental Informatics 22(1) 1-12 (2013)

Journal of Environmental Informatics

www.iseis.org/jei

Assessing a Stochastic Fire Spread Simulator

Willard J. Braun[1] and Douglas G. Woolford[2,*]

[1]*Department of Statistical and Actuarial Sciences, The University of Western Ontario, London, Ontario N6A 5B7, Canada*
[2]*Department of Mathematics, Wilfrid Laurier University, Waterloo, Ontario N2L 3C5, Canada*

Received 16 December 2011; revised 28 June 2012; accepted 15 March 2013; published online 25 September 2013

ABSTRACT. Several wildfire simulators have been proposed in recent years. Often, such simulators have been judged to be reasonable if they are capable of predicting mean fire growth accurately. However, when the management objective is a burn probability map, this kind of assessment will be insufficient, since it does not address the uncertainties intrinsic to fire behaviour. The problem of comparing output from a stochastic fire growth simulator with real fire behaviour is difficult in general. Focussing on a single (important) aspect of the fire behaviour provides a way through this difficulty. In this paper, we propose a method to judge the appropriateness of a particular stochastic fire spread simulator by comparing the variability in the rates of spread in simulated output with the variability in the rates of spread seen in some experimental micro-scale fires. The methodology can be applied to other stochastic fire spread models and to large-scale fires.

Keywords: fire spread, lattice model, variability, rate of spread

1. Introduction

Wildfire risk assessment is an important problem. Stochastic models of fire spread are required in order to produce burn probability maps which can be used by fire managers to identify high risk regions.

Currently, the most frequently used fire growth simulators, such as PROMETHEUS (Tymstra et al., 2005) and FARSITE (Finney, 2004) are deterministic. In tests of these simulators against real fires, it has been found that they give reasonable behaviour, but the best that can be achieved with a deterministic simulator is accurate mean behaviour; it is not possible for the inherent unpredictability of fire to be captured with such a model.

The Burn-P3 simulator (Parisien et al., 2005) is based on PROMETHEUS, and it produces a form of burn probability map, using randomized weather streams. Although weather has a major effect on fire behaviour, it should be noted that there is much variation in fire behaviour that is left unexplained by weather. Thus, a certain degree of ad hoc tuning is required in order to obtain probability maps that might be thought of as reasonable (see Braun et al., 2010, for an example). Another approach was taken by Garcia et al. (2008) who attempted to introduce stochasticity to the PROMETHEUS model via a block bootstrap procedure. Much work remains to be done in order to test the accuracy of that method, or to implement it operationally.

In the meantime, several other stochastic models have been proposed. Boychuk et al. (2007) proposed a stochastic lattice-spread model and claimed that it captured many of the desirable features of actual fire spread. Burn probability maps were produced by repeatedly running the stochastic simulations and calculating the proportion of time that individual pixels would be burned. The virtues of the lattice-spread model were demonstrated primarily through the use of graphs comparing median probability contours with actual fires. Thus, the model was shown to be capable of achieving correct median behaviour, but there was no attempt to validate other probability levels. Boychuk et al. (2009) further showed that the model has the potential to be extended to accommodate spot-fires ignited by lofted firebrands, though an accurate mechanism for this was not worked out and remains an open problem.

In this paper, we will consider the simplest case of the stochastic lattice-spread model: fires burn in homogeneous fuel and only slope is varied. In addition, we introduce a second model with a different rule that governs how burning cells extinguish. Our purpose is to investigate whether such models have the potential to produce burn probability maps which reflect the uncertainty that is intrinsic to fire spread. The approach we take is to consider the variability in the rates of spread observed in a sample of experimental microfires. We wish to determine whether the lattice-spread model is capable of simultaneously matching the expected growth behaviour as well as the variability seen in these fires. Our statistical methodology is simple and is not restricted to the stu-

* Corresponding author. Tel.: +1 519 884 0710; fax: +1 519 884 0464.
E-mail address: dwoolford@wlu.ca (D. G. Woolford).

ISSN: 1726-2135 print/1684-8799 online
© 2013 ISEIS All rights reserved. doi:10.3808/jei.201300241

dy of given lattice-spread model; any stochastic simulator could be tested in this way.

As in Zhang et al. (1992), we are studying the characteristics of the lattice-spread model in the context of data collected on tiny fires burned under very controlled conditions. Zhang et al. burned ordinary paper. We have chosen to use wax paper in our experiment because it burns more cleanly. Other researchers have also conducted lab experiments to study fires: Martínez-de Dios et al. (2006) performed indoor burning experiments on an adjustable burning table and recorded experimental fires with both point and linear ignitions at different levels of slope using digital and infrared cameras. That work was extended to open field experiments by Martínez-de Dios et al. (2008) where they demonstrated how such a system could be used to measure and visualize characteristics of forest fires in real time. In Rossa and Viegas (2009) a burning table was used to study back rates of spread. The data and approach we are taking in this paper should be viewed as complementary to the lab experiments in such studies. Although our idea here is not new, it is important to point out that even though we are working with very homogeneous fuel in zero wind our experimental fires still exhibited a large amount of natural variability that needs to be modelled and this is the focus of our paper.

The experimental data that we analyze come from 31 small fires. The only measured factor was slope: several levels of which were considered, with replicates. Otherwise, the fires were burned under identical conditions. It will be seen that even under such carefully controlled conditions, differences in observed fire behaviour cannot be attributed only to differences in slope. Small imperfections in the wax paper and other unmeasured factors appear to give rise to relatively large amounts of variability between fires. In order to check whether the fire spread model gives rise to similar amounts of between fire variability, we simulate the 31 fires, holding the model parameter values fixed within each level of the slope factor. Simple regression models will be used to relate the model parameters to the slope. Analogous regression models can also be applied to the simulated data as well, leading to a way of comparing variability between the simulated data and the actual fire data. We will make the comparisons on the basis of the residuals coming from both sets of models as well as the corresponding estimated error standard deviations.

The rest of the paper proceeds as follows. The next section contains a description of the stochastic lattice-spread model as well as a modified version which may yield realizations which more closely resemble actual fires. Then Section 2 concludes with a discussion of the relation between the scale of the grid and the variability induced by the lattice-spread model. Section 3 describes the methods employed in this research: the micro-fire experiments are described as is the methodology for the comparisons with simulated replicates of these fires. Results follow in Section 4, and the paper closes with some concluding remarks and directions for further research.

2. The Fire Simulation Models

In this section, we summarize the basic features of the lattice-spread model proposed by Boychuk et al. (2007). We will also propose a variant of that model which yields somewhat different behaviour. We will refer to the original version of the model as Model A and the new variant as Model B.

Both models have the same basic structure. The landscape is assumed to be planar, the weather conditions are constant, and it assumed that there is no wind. The fuel type and density are also assumed to be homogeneous. On this landscape we overlay a regular square $n \times m$ lattice. Each of the grid cells can be in one of three possible states: unburned fuel (F), burning fuel (B) or burnt out (O).

Transitions between these states occur as follows: initially (i.e. at time $t = 0$), the grid cell at the point of ignition (i, j) is in state B, while all other cells are in state F. The fire burning in cell (i, j) will spread to each of its four cardinal nearest neighbors (i.e. north, south, east and west) in random amounts of time $T_{0,1}$, $T_{0,-1}$, $T_{1,0}$, and $T_{-1,0}$, provided it does not burn out first. Specifically, at time $T_{0,1}$, the cell at $(i, j+1)$ makes the transition from state F to B, if the cell is not already in state B. Similar transitions are made by cell $(i, j-1)$ at $T_{0,-1}$, cell $(i+1, 0)$ at $T_{1,0}$ and cell $(i-1, 0)$ at $T_{-1,0}$. These times are assumed to be independent and exponentially distributed at rate λ, when there are no slope or wind effects. When there are such effects, the rate will depend on direction, and there are four rate parameters λ_N, λ_S, λ_E and λ_W.

Once a cell has made a transition to state B, fire spreads from that cell to the sites of its nearest neighborhood at a new set of independent exponential random times. Under Model A, a cell in state B makes the transition to state O according to another independent time that is exponentially distributed at rate μ. Once in state O, a grid cell will make no further transitions.

Note that, because of the minimum property of independent and identically distributed exponential random variables, when there are k burning cells, the time until the first of these sites burns out is exponentially distributed with rate $k\mu$. This, together with the memoryless property of the exponential distribution, provides an equivalent way of specifying the burnout rule for Model A: a cell is randomly selected from the set of k sites that are burning at the time of the last event and that cell makes the transition to O after the expiration of a simulated exponential $k\mu$ time period. Under Model B, a site is chosen to burn out at this same time, but not a randomly selected site; instead, the site which has been burning longest is chosen to burn out. This may be a somewhat more realistic rule. This innovation has not been considered previously. In this paper we demonstrate that Model A has some deficiencies - in particular, its burn out rule can lead to simulations with large numbers of unburnt “islands” a phenomenon that does not occur in the actual micro fire experiments. Model B's burn out rule, which has not been considered previously, leads to more realistic behaviour.

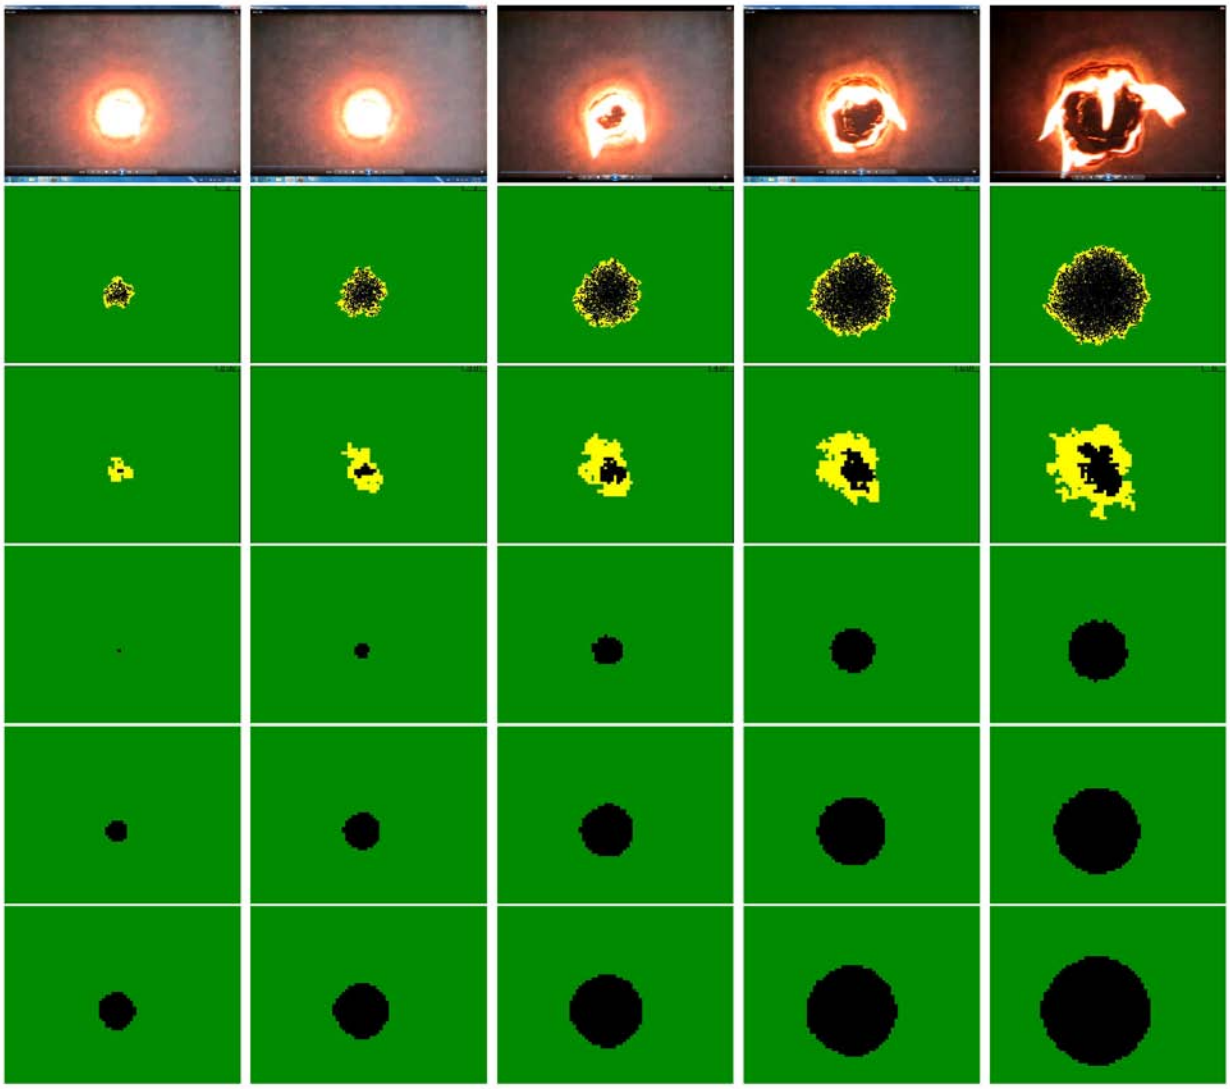

Figure 1. Top panel: A sequence of observed burn patterns for a fire on a sheet of wax paper (at 0 slope) observed at 1 second intervals. Time increases from left to right. Second row: A realization of a simulation of Model A under homogeneous conditions with $\mu = 0.666$, $\lambda_N = 4.448$, $\lambda_S = 4.369$, $\lambda_E = 2.984$, and $\lambda_W = 4.744$, observed at approximately 1 second intervals. The model was simulated on a 200 x 250 grid. Third row: A realization of a simulation of Model B under homogeneous conditions with $\mu = 0.2$, $\lambda_N = 1.336$, $\lambda_S = 1.312$, $\lambda_E = 0.896$, and $\lambda_W = 1.424$, observed at approximately 1 second intervals. The model was simulated on a 60 × 75 grid. Fourth row: Black area represents regions which burned in at least 90% of 1000 simulation runs. Fifth row: Black area represents regions which burned in at least 50% of 1000 simulation runs. Bottom row: Black area represents regions which burned in at least 10% of 1000 simulation runs.

It should be pointed out that these models are much more general than described above. They can handle non-homogeneous conditions with regard to weather, fuel type, moisture content, and topography. These issues, as they apply to Model A, are described in detail in the Boychuk et al. (2007) paper. In that paper, it was argued that the model produced realistic fire shapes. The model is not unreasonable from a scientific point of view, since it is based on a simple fire propagation mechanism: areas of unburned fuel immediately adjacent to burning regions will ignite because of the availability of fuel and oxygen, combined with sufficient radiant heat. However, there has been no attempt at model validation from a statistical perspective. In what follows, we will present a methodology that could be used to validate this model for wildfires, and illustrate its application using some experimental data which we describe in the next section.

2.1. Scale: Grid Resolution

The size of the grid cells is a critical parameter, since it represents the spatial scale of the simulator. The interplay between scale and variability in the lattice-spread model can

Figure 2. Top panel: A sequence of observed burn patterns for a fire on a sheet of wax paper (at 0 slope) observed at 1 second intervals. Time increases from left to right. Second and third rows: Realizations of simulations of Models A and B under homogeneous conditions with $\mu = 0.2$, $\lambda_N = 1.336$, $\lambda_S = 1.312$, $\lambda_E = 0.896$, and $\lambda_W = 1.424$, observed at approximately 1 second intervals. Both models were simulated on a 60 x 75 grid. Third row: A realization of a simulation of Model B under homogeneous conditions with $\mu = 0.2$, $\lambda_N = 1.336$, $\lambda_S = 1.312$, $\lambda_E = 0.896$, and $\lambda_W = 1.424$, observed at approximately 1 second intervals. Both models were simulated on a 60 x 75 grid. Fourth row: Black area represents regions which burned in at least 90% of 1000 simulation runs. Fifth row: Black area represents regions which burned in at least 50% of 1000 simulation runs. Bottom row: Black area represents regions which burned in at least 10% of 1000 simulation runs.

be illustrated very simply using a one-dimensional version of the model: we imagine the fire travelling uni-directionally along a fuse or wick, letting T denote the time until the fire has travelled a distance d from its point of ignition.

A lattice model applied to this problem would have us postulate that T is the sum of k independent exponential random variables (where k is the number of lattice cells the fire must traverse before arriving at its destination). If we fix the mean of T to be some constant value m, we have that $m = k\lambda_k$, where λ_k is the rate for each independent and identically distributed exponential component of T. The variance of T is easily seen to be k/λ_k^2, or m^2/k. Hence, the variance of the transit time for a sojourn of interest depends on size of the cells in the lattice.

It is not possible to write down a closed-form expression for the mean-variance relation in the more complicated model discussed in this paper. Nevertheless, such a relationship undoubtedly exists - the modelled variance of the time

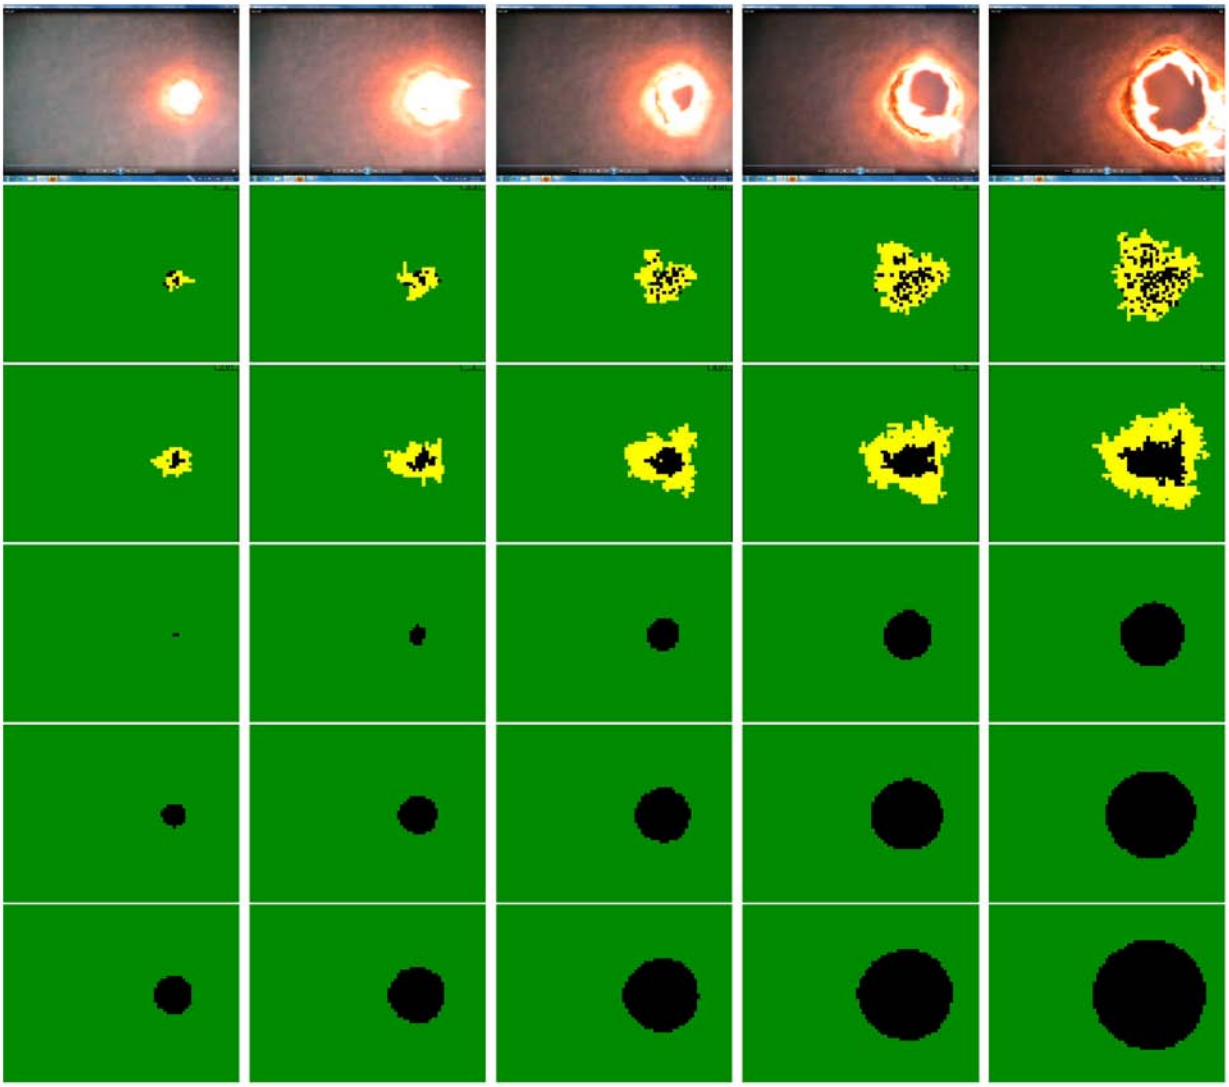

Figure 3. Top panel: A sequence of observed burn patterns for a fire on a sheet of wax paper (at slope 0.03) observed at 1 second intervals. Time increases from left to right. Second and third rows: Realizations of simulations of Models A and B under homogeneous conditions with $\mu = 0.2$, $\lambda_N = 1.336$, $\lambda_S = 1.312$, $\lambda_E = 0.808$, and $\lambda_W = 1.665$, observed at approximately 1 second intervals. Both models were simulated on a 60 × 75 grid. Fourth row: Black area represents regions which burned in at least 90% of 1000 simulation runs. Fifth row: Black area represents regions which burned in at least 50% of 1000 simulation runs. Bottom row: Black area represents regions which burned in at least 10% of 1000 simulation runs.

until the fire spreads a certain amount will depend on the size of the cells in the lattice. Thus, ideally, information about the variance should be used to determine the grid cell size.

3. Methods

3.1. Collecting Data from Micro-Fire Experiments

Data for testing the usefulness of the grid-based fire spread model was obtained from a sequence of small-scale experimental fires. These experiments were conducted indoors, at a temperature of approximately 10 ^{o}C, and the wind was considered negligible. The material used for fuel in each case was a rectangular sheet of dry wax paper. For each experiment, this sheet was suspended horizontally or at various moderate slopes. For a non-zero slope, the sheet was always inclined from East to West (the slope was always 0 in the North-South direction). The slopes used for the various experimental runs are listed in Table 1.

The paper was ignited from below at a variety of locations, and an Olympus Stylus® 600 Camera, which was suspended on a tripod above the paper, was used to record the experimental fire until most of the paper was consumed.

3.2. Data Extraction and Measurement of Rate of Spread

Table 1. Measured Rates of Spread (Inches/s) in Each Direction: East, West, North and South for each of 31 Experimental Fires at Given Slopes

Run	Slope	*ROSE*	*ROSW*	*ROSS*	*ROSN*	Time
1	0.00	0.74	0.90	0.49	1.06	7
2	0.03	(0.51)	0.93	0.78	0.95	6
3	0.03	0.68	1.02	1.12	0.88	4
4	0.03	(0.51)	0.90	1.18	0.95	4
5	0.03	(0.51)	1.06	(0.81)	(0.84)	5
6	0.03	0.45	1.18	0.60	0.95	4
7	0.03	(0.51)	1.02	1.32	0.70	4
8	0.03	(0.51)	0.98	0.94	0.80	5
9	0.03	0.50	(1.04)	1.05	0.98	4
10	0.03	(0.51)	1.65	0.80	0.77	4
11	0.07	(0.51)	1.37	0.72	0.55	6
12	0.07	0.60	(1.22)	0.66	0.54	5
13	0.07	0.33	1.07	0.68	0.68	6
14	0.02	0.68	0.88	0.60	0.78	5
15	0.02	0.82	1.18	0.93	1.32	4
16	0.03	0.46	1.06	0.62	0.82	5
17	0.03	0.33	1.33	1.17	0.70	3
18	0.05	0.43	1.00	1.35	0.97	4
19	0.05	0.42	0.82	0.95	1.05	4
20	0.08	0.38	(1.26)	0.78	0.85	6
21	0.13	0.50	(1.59)	0.82	(0.71)	5
22	0.13	0.48	(1.59)	0.78	0.86	5
23	0.13	0.30	(1.59)	(0.89)	(0.71)	7
24	0.17	0.50	1.96	(0.93)	(0.66)	5
25	0.17	0.33	(1.92)	(0.93)	(0.66)	6
26	0.17	0.32	(1.92)	(0.93)	(0.66)	5
27	0.14	0.26	(1.68)	(0.90)	(0.69)	5
28	0.09	0.42	(1.35)	(0.86)	(0.76)	5
29	0.09	0.35	(1.35)	(0.86)	(0.76)	6
30	0.00	0.45	1.02	0.57	0.60	6
31	0.00	0.58	0.60	0.70	1.04	5

* Measured over the given time period of each (measured in seconds). Values in brackets were imputed from regressions of log(*ROS*) on slope for each direction, using the available completely observed data.

Windows® Movie Maker was used to freeze-frame the movie at approximately 1 second intervals to obtain clear image-captures with time-stamps. These captured images were then saved as .png files. For illustrative purposes, images for some of the fires are shown in the top panels of Figures 1, 2 and 3. The first two figures show fires burned under identical conditions with no slope; this gives a visual impression of the sampling variability between fires. Figure 3 shows a fire burned at a slope of 0.03 (inclined from East to West).

Approximate values of the rate of spread (*ROS*, in inches/s) were obtained for each fire by measuring the distance travelled by the fire front in the four cardinal directions. Although crude, this method is not unlike the method commonly used in experimental forest fires where rates of spread are estimated by timing the passage of the flaming front between two landmarks (Forest Encyclopedia Network, 2011).

Table 1 contains the *ROS* measurements for the given fires. In some cases, the fire was ignited near the edge of the sheet; in such cases, it was not possible to measure *ROS* in all directions. This resulted in partially missing observations.

Simple imputation was used to replace the missing values; simple regressions were computed for log(*ROS*) for each of the four directions, using the complete data as described at the beginning of the next subsection. The missing rates of spread were then imputed by exponentiating the predicted values of log(*ROS*) coming from the fitted models. The imputed rates of spread appear in brackets in Table 1.

3.3. Relating *ROS* to Slope

Using the non-missing data, regressions of log(*ROS*) on slope were calculated for each of the East and West directions:

$$\log(ROSW) = -0.116 + 4.56 \times \text{slope} \quad (s = 0.1972) \qquad (1)$$

$$\log(ROSE) = -0.580 - 3.01 \times \text{slope} \quad (s = 0.2552) \qquad (2)$$

Here, *s* (the term in brackets) corresponds to the estimated residual standard deviation for each model. Because the slope in the northerly and southerly directions was 0, simple averages (on the log scale) were calculated for these directions:

$$\log(ROSN) = -0.181 \quad (s = 0.2246) \qquad (3)$$

$$\log(ROSS) = -0.198 \quad (s = 0.2831) \qquad (4)$$

Note that, because of missing values, the above regressions were carried out on 20, 24, 23 and 24 observations, respectively. The log scale was used, because it was the recommended Box-Cox transformation (Box and Cox, 1964) for all directions but North. A square root transformation would be more appropriate in the North direction, but the log scale is only moderately sub-optimal, and using the same form of transformation in all directions leads to a more parsimonious over-all model. Furthermore, standard diagnostic procedures (residual-versus-fitted plots, normal QQ-plot of residuals, influence plots) indicate a substantial improvement in the behaviour of the residuals when a log transformation is used as compared with no transformation at all.

Note also, that even though the slopes in both the North and South directions are 0, we have not imposed an equality constraint on the *ROS* in these directions, since we allow for the possibility that the wax paper has a “grain” or some other property which could lead to a systematic bias. In fact, the *ROS* estimates given above in (3) and (4) do not support such a conclusion, since their difference lies within one standard error.

All of the above calculations were undertaken in R (R

Core Development Team, 2011) with the use of the boxcox() function (Venables and Ripley, 2002).

3.4. Relating the Exponential Rates to the Fitted *ROS* Values

Our next step was to use the fire spread simulator to generate replicates of the 31 recorded fires. The exponential rates for the simulator were taken to be proportional to the rates of spread predicted from the models obtained in the previous section. The proportionality constant was chosen by trial and error from comparisons between the simulated realizations and the observed fires. The objective here was to achieve a rough matching between the fire growth behaviour in the observations and in the simulations. Specifically, we ran individual simulations at the 200 × 250 grid resolution, using different parameter settings. For each simulation run, the rate of spread in each cardinal direction was measured, and the proportion of simulated burned area was noted. Parameter values for subsequent simulation runs were increased if the simulated rate of spread values were lower than the observed rate of spread values, and decreased if the simulated values were higher. The burnout rate was chosen similarly, but using only the output from the final observation time: if the proportion of black area in the simulated fire appeared to be larger than in the observed fire, the burnout rate for the subsequent simulation run was decreased. If the simulated burned area appeared to be too small, then the updated burnout rate was increased. The procedure was stopped after only a few iterations, when a reasonable qualitative match was observed between the simulated and observed fires. The 5 parameter estimates (4 spread rates and 1 burn-out rate) were then scaled down to the lower resolutions by simple division, e.g., for the 60 by 75 resolution, each parameter estimate was divided by 3.33. Statistical estimation of the model parameters is ongoing work.

The exponential rate parameters λ_N, λ_S, λ_E and λ_W are given in Table 2 at the grid resolution of 200 × 250. The burn-out rate was set at $\mu = 0.667$. The ignition locations for each simulated fire were taken at the same locations as for the original fires.

Exponential rate parameters used for other grid dimensions were obtained by dividing the rates in Table 2 by constants in order to adjust for different sizes of grid cells: for a 60 × 75 resolution, the rates were divided by 3.33, for a 80 × 100 resolution, the rates were divided by 2.5, for a 100 × 125 resolution, the rates were divided by 2, and for a 120 × 150 resolution, the rates were divided by 1.67. The burn-out rates were similarly scaled and are given in Table 3.

Simulations were conducted in order to replicate the set of 31 fires under Models A and B using each of the grid dimensions listed in Table 3. Since there is insufficient space to display pictures of all of these simulated fires, we have arbitrarily selected three for illustrative purposes.

One of the simulated fires (under Model A, corresponding to a slope of 0, and using a 200 × 250 grid) is pictured in the second row of Figure 1. This simulation realization was chosen to demonstrate the appearance of output when a fine resolution is used. A simulation realization for Model B using the coarser 60 × 75 resolution is pictured in the bottom row of Figure 1. The second and third panels of Figures 2 and 3 exhibit simulation realizations using the 60 × 75 grid dimensions under Models A and B. The exponential growth rate parameters used for these simulations were appropriate for slopes of 0 and 0.03, respectively. Notice the appearance of many unburnt “islands” in Model A's simulated output.

Table 2. Estimated Exponential Rates to be Used in the Simulations of the Replicates of the 31 Fires

Run	Slope	λ_N	λ_S	λ_E	λ_W
1	0.00	4.45	4.37	2.98	4.74
2	0.03	4.45	4.37	2.69	5.54
3	0.03	4.45	4.37	2.69	5.54
4	0.03	4.45	4.37	2.69	5.54
5	0.03	4.45	4.37	2.69	5.54
6	0.03	4.45	4.37	2.69	5.54
7	0.03	4.45	4.37	2.69	5.54
8	0.03	4.45	4.37	2.69	5.54
9	0.03	4.45	4.37	2.69	5.54
10	0.03	4.45	4.37	2.69	5.54
11	0.07	4.45	4.37	2.42	6.52
12	0.07	4.45	4.37	2.42	6.52
13	0.07	4.45	4.37	2.42	6.52
14	0.02	4.45	4.37	2.84	5.12
15	0.02	4.45	4.37	2.78	5.29
16	0.03	4.45	4.37	2.76	5.33
17	0.03	4.45	4.37	2.76	5.33
18	0.05	4.45	4.37	2.56	5.99
19	0.05	4.45	4.37	2.56	5.99
20	0.08	4.45	4.37	2.37	6.72
21	0.13	4.45	4.37	2.04	8.45
22	0.13	4.45	4.37	2.04	8.45
23	0.13	4.45	4.37	2.04	8.45
24	0.17	4.45	4.37	1.80	10.24
25	0.17	4.45	4.37	1.80	10.24
26	0.17	4.45	4.37	1.80	10.24
27	0.14	4.45	4.37	1.96	8.95
28	0.09	4.45	4.37	2.27	7.19
29	0.09	4.45	4.37	2.27	7.19
30	0.00	4.45	4.37	2.98	4.74
31	0.00	4.45	4.37	2.98	4.74

*Using grid dimensions 200 × 250. These rates are based on scaled fitted values from the regressions in (1) to (4).

Table 3. Burn out Rates at each Grid Cell Resolution

Grid Dimensions	μ
60 × 75	0.200
80 × 100	0.267
100 × 125	0.333
120 × 150	0.400
200 × 250	0.667

Rates of spread in each direction were measured for each of the simulated fires. For the different grid dimensions, the rates of spread have been scaled so as to be directly comparable with the rates measured on the actual fires as listed in Table 1. As a proof of concept, simple burn probability maps were computed using model B and the 60 by 75 grid resolution, using the same parameters as for the model B simulation. The 4th row of plots shows (in black) the locations that would burn with 10% probability by the given time, the 5th row displays the corresponding 50% plots and the bottom row contains the 90% plots. The 10% plots give an indication of the minimal expected fire area, while the 90% indicate how large such a fire could usually be expected to be. Note that in 2 of the 3 cases, the actual burned area (in time sequence) lies between the median (50%) area and the 90% area. In the third case, the actual fire outside the area predicted by the 90% probability map. The accuracy of this map has likely been compromised somewhat by inaccuracy in the model parameter estimates. Since a proper parameter estimation methodology is still under development, we will not discuss these probability maps further in this paper.

Table 4. Regression Results for Simulated Data - West Direction. Compare with Model in Equation (1)

Grid Dimensions	β_{0W}	β_{1W}	σ_W
Model A			
60 × 75	-0.0552	2.48	0.179
80 × 100	-0.141	3.09	0.286
100 × 125	-0.126	3.07	0.129
120 × 150	-0.133	4.18	0.120
200 × 250	-0.104	3.70	0.133
Model B			
60 × 75	-0.126	2.94	0.180
80 × 100	-0.107	2.66	0.166
100 × 125	-0.036	3.15	0.145
120 × 150	-0.083	3.98	0.129
200 × 250	-0.061	3.72	0.148

Table 5. Regression Results for Simulated Data - East Direction. Compare with Model in Equation (2)

Grid Dimensions	β_{0E}	β_{1E}	σ_E
Model A			
60 × 75	-0.522	-4.16	0.429
80 × 100	-0.535	-2.35	0.215
100 × 125	-0.548	-1.69	0.209
120 × 150	-0.523	-2.31	0.166
200 × 250	-0.480	-1.87	0.138
Model B			
60 × 75	-0.512	-1.87	0.281
80 × 100	-0.520	-1.16	0.285
100 × 125	-0.435	-1.34	0.161
120 × 150	-0.366	-2.34	0.197
200 × 250	-0.344	-2.12	0.095

3.5. Comparisons with Regressions on the Simulated Data

For each of the 5 sets of experimental runs listed in Table 3, the measured rates of spread were regressed against slope, analogously to the technique used for the actual fires. That is, the following models were fit to each of the 5 data sets:

$$\log(ROSW) = \beta_{0W} + \beta_{1W} \times \text{slope} + \varepsilon \quad (5)$$

$$\log(ROSE) = \beta_{0E} + \beta_{1E} \times \text{slope} + \varepsilon \quad (6)$$

$$\log(ROSN) = \beta_{0N} + \varepsilon \quad (7)$$

$$\log(ROSS) = \beta_{0S} + \varepsilon \quad (8)$$

The error terms were assumed to have constant variances σ^2_W, σ^2_E, σ^2_N and σ^2_S, respectively.

Table 6. Regression Results for Simulated Data - North Direction. Compare with Model in Equation (3)

Grid Dimensions	β_{0N}	σ_N
Model A		
60 × 75	-0.234	0.252
80 × 100	-0.223	0.182
100 × 125	-0.229	0.144
120 × 150	-0.190	0.137
200 × 250	-0.163	0.120
Model B		
60 × 75	-0.160	0.245
80 × 100	-0.128	0.161
100 × 125	-0.125	0.158
120 × 150	-0.155	0.142
200 × 250	-0.088	0.081

Table 7. Regression Results for Simulated Data - South Direction. Compare with Model in Equation (4)

Grid Dimensions	β_{0S}	σ_S
Model A		
60 × 75	-0.307	0.352
80 × 100	-0.297	0.235
100 × 125	-0.194	0.158
120 × 150	-0.171	0.139
200 × 250	-0.163	0.120
Model B		
60 × 75	-0.160	0.245
80 × 100	-0.226	0.190
100 × 125	-0.152	0.093
120 × 150	-0.143	0.128
200 × 250	-0.132	0.085

Each of the above models was fit only to simulated observations which corresponded to non-missing observations in the actual fire data set. This ensured that exactly the same number and type of observations were used in the simulation-based regression models as in the actual-fire-based regression models. Thus, model (5) was fit only to the obser-

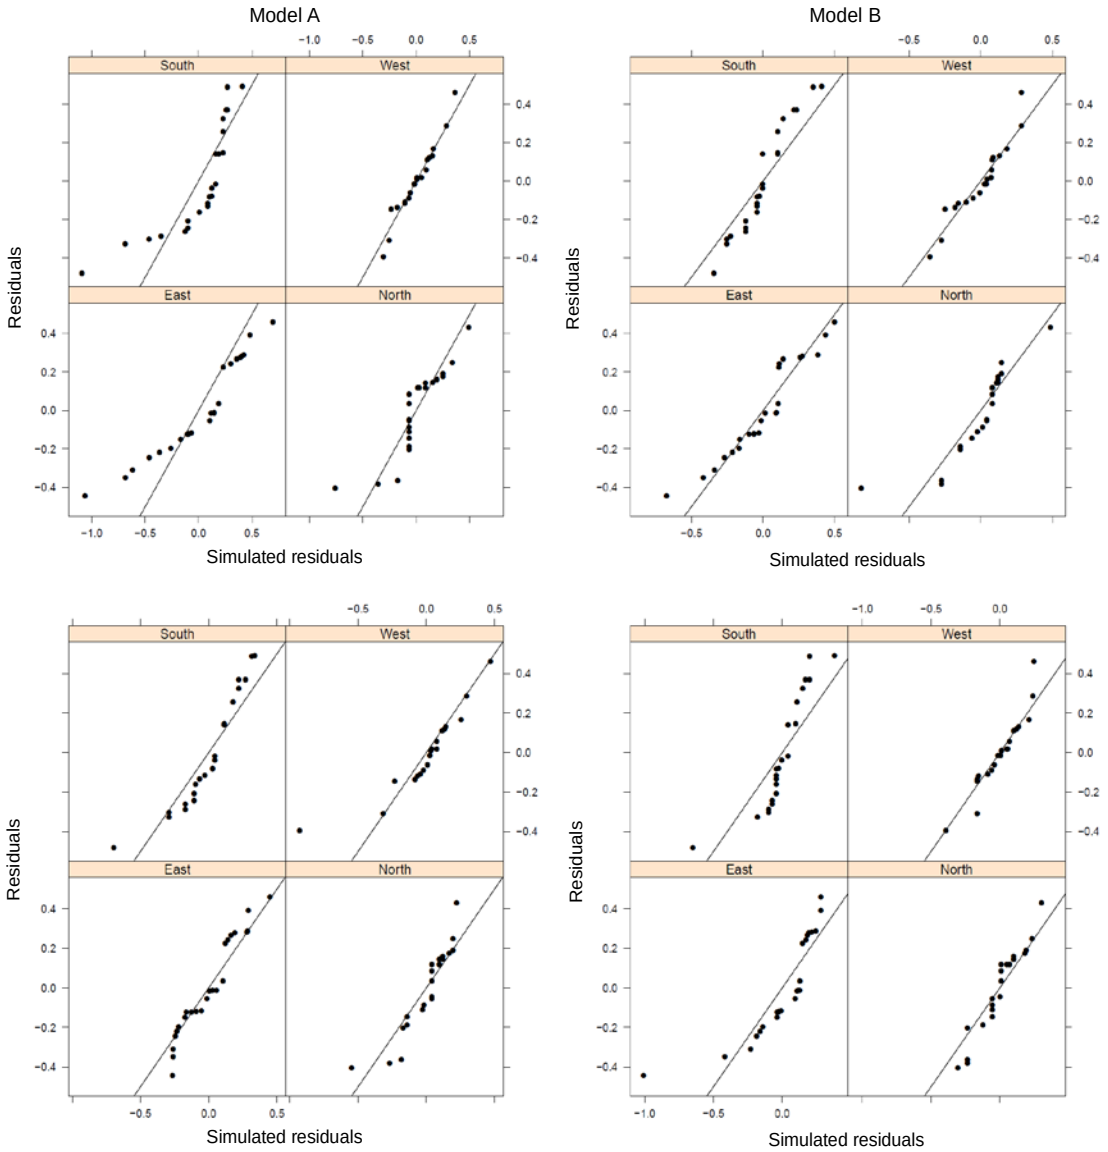

Figure 4. QQ-plots of residuals from regression models (5) to (8) fitted to observed data and to data simulated at a coarse resolution using from Model A (left side) and Model B (right side). Top panel: simulation grid resolution is 60 × 75. Bottom panel: simulation grid resolution is 80 × 100.

vations which corresponded to the 20 non-missing observations used to fit model (1). Model (6) was fit to 24 observations, model (7) was fit to 23 observations and model (8) was fit to 24 observations.

The parameter estimates are displayed in Tables 4-7. Also displayed are the estimates of the error standard deviations for each case.

3.6. Residual QQ-Plots

Recall, our focus is to compare the variability in the simulated rates of spread to that which was observed. The residuals from the *ROS* regressions for the actual fires in equations (1) through (4) contain information about the variability in the fires which can be compared with the corresponding residuals for *ROS* regressions for the simulated fires. The error standard deviation estimates give a numerical summary measure, but additional distributional information can be obtained from quantile-quantile (QQ) comparisons.

Figures 4 and 5 contain QQ-plots of the residuals from the regressions of the actual fire rates of spread against the residuals from the regressions of the simulated fire rates of spread. These plots were constructed for the five different grid dimensions studied and for each of the two models.

4. Results and Discussion

The full set of fires and corresponding simulation realiza-

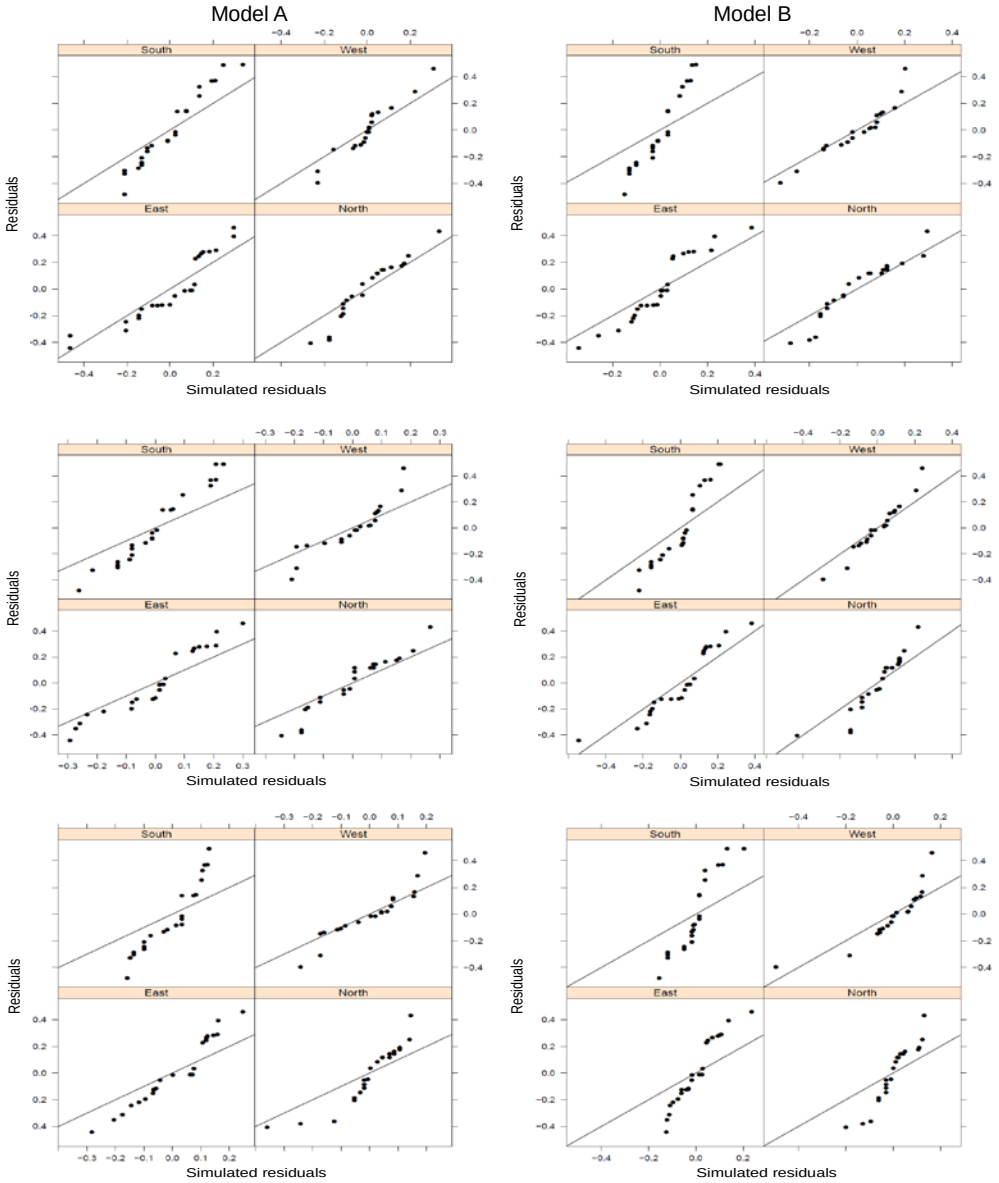

Figure 5. QQ-plots of residuals from regression models (5) to (8) fitted to observed data and to data simulated at a fine resolution using Model A (left side) and Model B (right side). Top Panel: simulation grid resolution is 100 × 125. Middle Panel: simulation grid resolution is 120 × 150. Bottom Panel: simulation grid resolution is 200 × 250.

tions are documented in Braun and Woolford (2011). What can be observed in these pictures is that when the ignition point is sufficiently far from the edge, the simulated fire behaviour apparently matches the actual fire behaviour. Thus, on the basis of visual comparisons between the simulated and actual fire realizations, it is possible to reinforce the conclusion of Boychuk et al. (2007) that the lattice-spread model is capable of reproducing observable fire growth. When the actual fire reaches the edge, the simulator appears to be incapable of reproducing the correct fire behaviour, tending to spread more slowly than the actual fire. This is likely due to the fact that more oxygen is available for consumption at the edge, and the simulation model does not account for this process change.

Note also the tendency for the Model A output to have many isolated burning grid cells in the interior, while the burning cells in Model B are concentrated at the fire perimeter. Model B's behaviour appears to more closely match the experimental fires. The occurrence of these “islands”, which are not generally observed in the actual fire experiments fires, is greatly reduced with the use of Model B's new burn out rule.

Pictures of actual fires and simulated fires cannot be used to make variability comparisons. For this purpose, the regressions of log(*ROS*) on slope are more useful. First, it should be pointed out that the slope and intercept estimates listed in Tables 4 through 7 do not always match the slope and intercept estimates for the regressions based on the actual fires. In

the cases of the North and South directions, there is a good correspondence between the simulations and the actual fires. However, while the intercepts for the West direction match reasonably well, the slopes for the simulations tend to be considerably smaller in magnitude than the slope estimate based on the actual fires. With one exception, the same can be said for the East direction. These discrepancies are due to the fact that the parameters used in the simulation have been chosen in a fairly ad hoc manner; we have not yet pursued a careful estimation strategy. In addition, as observed earlier, the simulated fires tend to spread slower than the actual fires when the edge has been reached by the fire front. This is also a likely contributor to a systematic under-estimation.

The more important comparison, from the standpoint of the current paper, is the residual standard deviation. As conjectured in the introduction to this paper, there is a tendency for the residual standard deviation to decrease as the grid dimensions increase.

The residual/simulated residual QQ-plots in Figures 4 and 5 provide an elaboration on the preceding statement. It appears that the error distributions do not match well at all for high grid dimensions (Figure 5) but that there is a reasonable correspondence at the coarser resolutions (Figure 4).

On the basis of the results given in the tables and figures, one might suggest that the fires could be simulated reasonably well with Model B at a 60 by 75 grid resolution, or Model A at the 80 by 100 resolution. It is clear that both of these models are capturing the rate of spread variability adequately. In fact, it might be argued that they are capturing other aspects of the rate of spread distribution as well. It is worth emphasizing that a simple visual assessment may have led an "expert" to conclude that a higher resolution was adequate; this would perhaps lead to correct mean behaviour, but the variability between fires would have been under-estimated substantially.

The choice between Model A and B would have to be made on a different basis, such as the visual comparison discussed above.

5. Conclusions

This work is part of an ongoing investigation into the suitability of a simple lattice-based model for stochastically modelling forest fire spread. Ultimately, we wish to fit such a model to sequences of satellite-based photographs of wildfires or infrared photographs obtained from helicopters such as those analyzed by Dold et al. (2010). Then, simulations of the model could be used to produce the maps of fire spread risk that are in demand by forest fire managers. Before doing so, a systematic methodology must be developed for estimating the exponential rate parameters. This is the subject of an upcoming paper by the authors.

The model needs to be further validated on data collected under experimentally controlled conditions. The current paper represents the analysis of one such experiment. Further follow-up studies on additional experimental fires are necessary prior to drawing firm conclusions on the validity of this process as a model for fire spread. In particular, experimental fires incorporating wind need to be studied.

What can be firmly concluded on the basis of the current work is that the lattice-spread model and its variant have the flexibility to capture the nature of the variability exhibited in the kinds of micro-fires studied here. A critical element here is scale. We have shown that the "natural" grid-cell size is related to the variability in the observed rates of spread. Given the appropriate scale, the stochastic spread simulator is then capable of generating images of fires which have the same range of fire spread rates which could be observed in the micro-fires burned under the given conditions.

There is no reason for us to think that this relation between variability and scale would be any different than what would be observed in large scale experiments. Generally many variables would affect large scale fire spread and the rate parameters would need to account for locally observed covariates (e.g., wind, topography, fueltype and density). There will be uncertainty induced by the imprecision in how these factors would be measured and that would also translate into uncertainty that would need to be modelled in the way that we modelled it. Even if we considered all this, there's still natural variability that we need to account for. Looking at variability allows us to choose scale. This is an important step toward demonstrating that such a model will indeed produce valid burn probability maps.

Acknowledgments. The authors gratefully acknowledge support from GEOIDE ("Stochastic Modelling of Forest Dynamics"). The authors also acknowledge technical assistance from Qiang Fu, Yu Han and Ethan Shrubsole.

References

Box, G.E.P., and Cox, D.R. (1964). An analysis of transformations (with discussion). *J. Roy. Stat. Soc. B*, 26, 211-252.

Boychuk, D., Braun, W.J., Kulperger, R.J., Krougly, Z.L., and Stanford, D.A. (2007). A stochastic model for forest fire growth. *INFOR*, 45(1), 339-352.

Boychuk, D., Braun, W.J., Kulperger, R.J., Krougly, Z.L., and Stanford, D.A. (2009). A stochastic forest fire growth model. *Environ. Ecol. Stat.*, 16, 133-151. http://dx.doi.org/10.1007/s10651-007-0079-z

Braun, W.J., Jones, B.L., Lee, J.S.W., Woolford, D.G., and Wotton, B.M. (2010). Forest fire risk assessment: an illustrative example from Muskoka, Ontario. *J. Prob. Stat.*, 1-16. http://dx.doi.org/10.11 55/2010/823018

Braun, W.J., and Woolford, D.G. (2011). Full set of comparisons between simulations and actual fires used for assessing a fire spread simulator. http://www.stats.uwo.ca/faculty/braun/research/microfires/viscomp.pdf

Dold, J., Tsitsopoulos, V., Khan, I., Scott, K., McMorrow, J., Lowe, E., Danson, F.M., Ramirez, A., Doerr, S., Bryant, R., Harris, M., Tollitt, T., Allen, K., Paugam, R., Freeborn, P., Smith, T., Davies, H., Wooster, M., Legg, C., Gibson, S., Elliott, A., and Yearsley, S. (2010). Report on field experiments in Northumberland, March 2010-a multidisciplinary approach to assess fire behaviour and effects in a temperate climate. *VI International Conference on Forest Fire Research*, Coimbra, Portugal. D.X. Viegas (Ed.)

Finney, M.A. (2004). FARSITE: fire area simulator-model development and evaluation. Res. Pap. RMRS-RP-4, Ogden, UT: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station.

Forest Encyclopedia Network. 2011. http://www.forestencyclopedia.net

Garcia, T., Braun, W.J., Bryce, R., and Tymstra, C. (2008). Smoothing and bootstrapping the PROMETHEUS Fire Spread Model. *Environmetrics*, 19(8), 836-848. http://dx.doi.org/10.1002/env.907

Martínez-de Dios, J.R., Andre, J.C., Goncalves, J.C., Arrue, B.C., Ollero, A., and Viegas, D.X. (2006). Laboratory fire spread analysis using visual and infrared images. *Int. J. Wildland Fire*, 15(2), 179-186. http://dx.doi.org/10.1071/WF05004

Martínez-de Dios, J.R., Arrue, B.C., Ollero, A., Merino, L., and Gómez-Rodríguez. (2008). Computer vision techniques for forest fire perception. *Image Vision Comput.*, 26(4), 550-562. http://dx.doi.org/10.1016/j.imavis.2007.07.002

Parisien, M.A., Kafka, V.G., Hirsch, K.G., Todd, J.B., Lavoie, S.G., and Maczek, P.D. (2005). Mapping wildfire susceptibility with the Burn-P3 simulation model. Information Report NOR-X-405.

R Development Core Team. (2011). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0. http://www.R-project.org

Rossa, C.G., and Viegas, D.X. (2009). Propagation of wind and slope backfires. *18th World IMACS / MODSIM Congress*, Cairns, Australia 13-17 July 2009. http://mssanz.org.au/modsim09

Tymstra, C. (2005). Prometheus: the Canadian wildland fire growth model. Forest Protection Division, Alberta Sustainable Resource Development, Edmonton, AB. http://www.firegrowthmodel.com

Venables, W.N., and Ripley, B.D. (2002). Modern applied statistics with S. Fourth edition. Springer. http://dx.doi.org/10.1007/978-0-387-21706-2

Zhang, J., Zhang, Y.-C., Alstrom, P., and Levinsen, M.T. (1992). Modeling forest fire by a paper-burning experiment, a realization of the interface growth mechanism. *Physica A*, 189(3-4), 383-389. http://dx.doi.org/10.1016/0378-4371(92)90050-Z