

Spatial Modeling of Post-Frontal Fire Behavior

Final Report RMRS-99557-RJVA

Mark A. Finney
Systems for Environmental Management
PO BOX 8868
Missoula MT 59808

Abstract

A simulation model was developed to calculate smoke and heat production from fires growing in two dimensions. The work required modifications of the *FARSITE* model (Finney 1998) to incorporate spatially the combustion history for each vertex on the fire perimeter. This combustion history was simulated locally at each vertex using the Burnup model (Albini and Reinhardt 1997). It was then applied along the vertex trajectory and integrated repeatedly over the area represented by that line for successive timesteps until all combustion activity ceased. Simulations produce the heat and weight-loss for the area represented by each vertex (and for each fire) as the simulation progresses. The simulation of post-frontal products depends on having accurate fuel moistures.

Introduction

Post-frontal combustion is defined as fire activity occurring after the passage of the advancing edge of the fire. It is neglected from most models of fire behavior, largely because it is a time-dependent process that lacks data to describe it and models to predict it. Current simulation models of fire behavior and growth (*e.g.* BEHAVE, *FARSITE*) only compute characteristics of the advancing edge of the fire. Post-frontal combustion, however, is essential to the production of smoke and to understanding patterns of heat-evolution from large fires that are believed critical to lofting of the smoke and interactions with the atmosphere (Rothermel 1991, 1994, Linn and Harlow 1998).

The *FARSITE* model (Fire Area Simulator, Finney 1998) incorporates five sub-models of fire behavior: surface fire (Rothermel 1972, Albini 1976), crown fire (Rothermel 1991, Van Wagner 1977, 1993), fire acceleration (CFBPS 1993), fuel moisture (Rothermel et al. 1986), and spotting from torching trees (Albini 1979). All fire behavior calculations apply to the edge of a fire. If the edge of the fire is specified as a line or polygon, the model calculations are performed at the vertices and serve to locally propagate that edge as elliptical wavelets according to Huygen's principle (Anderson et al. 1982, Richards 1989, Finney 1998). The elliptical wave shape (Anderson 1983) is determined locally at each vertex by the additive combination of windspeed and slope (Finney 1998). Fire growth is simulated as the aggregate movement of all vertices along the fire edge during a finite timestep. Spatial data on fuels and topography required for these calculations are obtained from gridded maps from a geographic information system (GIS). Weather is obtained from data streams that discern time-dependent changes in wind speed, wind direction, temperature, and humidity.

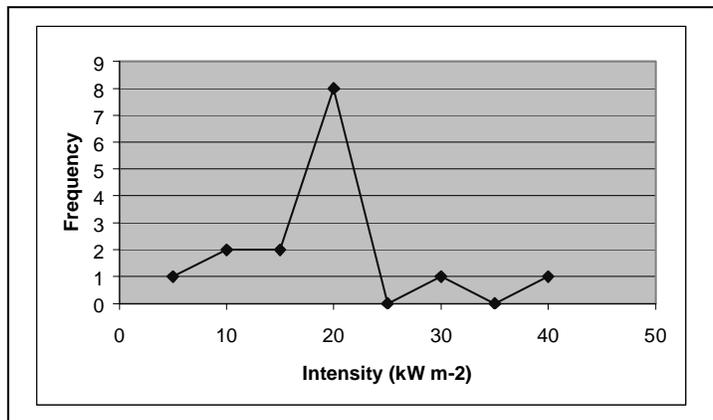
This process of simulating fire growth extends the existing models of fire behavior to a two-dimensional fire front but ignores any combustion activity inside the fire. It only propagates the edge of a fire polygon. No information is stored on the trajectories of the vertices as they progress to the new positions on the landscape. Substantial modification of the internal structure of *FARSITE* was required to store these trajectories, but also to apply a model of combustion history to them. This report summarizes work on incorporating a recently developed simulation model of woody fuel consumption (Burnup, Albini and Reinhardt 1996) into *FARSITE*, and on developing the methods for numerically integrating the combustion history (from Burnup) to represent spatial dynamics of post-frontal combustion.

Review of Models and Previous Work

The Burnup model (Albini and Reinhardt 1996) simulates the combustion history of a fuel complex composed of duff and woody fuels. The fuel complex is described as horizontal loading (kg m^{-2}) of any number or size distribution of woody fuels and their individual physical characteristics (*i.e.* density (kg m^{-3}), heat content (kJ kg^{-1}), etc.). The model requires an input of the initial fire conditions that ignite the elements of the fuel complex, namely fire intensity (kW m^{-2}) and the residence time (sec) of the ignition pulse. The environmental conditions are also required, including windspeed and the moisture contents of the woody fuels and duff. The Burnup model then simulates the burning of the fuel elements over time, but explicitly excludes effects of wind on burning rate of the woody fuels or duff. Timesteps are typically 5-30 seconds, short enough to capture the burning time of the finer fuel elements. The outputs from Burnup are intensity (kW m^{-2}) and fuel weight loss at each timestep. These apply to a unit-area of the specified fuel complex.

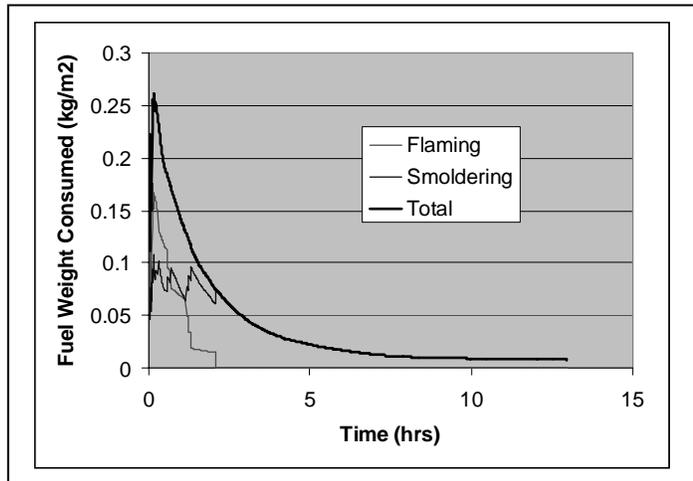
Prior to this effort, the Burnup model was modified to distinguish flaming from smoldering combustion. This distinction permitted emissions factors to be applied separately to each fraction of the fuel weight lost. Flaming was distinguished from smoldering using an intensity threshold (17 kW m^{-2}) for each size class of fuel particle in the fuel complex at each timestep. Below this threshold, smoldering combustion was assumed.

The threshold was the average of 15 observations of intensity at the end point of flaming combustion on slash fires (Sandberg et al. 1996, Figure X). Field data from the slash burns contained ocular estimates on the elapsed time that flaming combustion ceased. The fire intensity at those times was then obtained from the raw smoke sample data on those fires that were used to compute the intensity from the CO_2 output. Duff combustion was assumed to occur by smoldering only. Thus, the weight loss and intensity occurring at every timestep was divided into smoldering and flaming fractions. These can be graphed over time to show the decline in flaming combustion and increasing smoldering fraction (Figure X).



The use of the Burnup model in computing emissions of large fires was dependent on modifications to the *FARSITE* simulation model that preserved the spatial histories of fire

progress. *FARSITE* had been modified to allow the storage of spatial data on past fire progress. The data structures, called “rings”, essentially contain two consecutive fire perimeters that are each composed of vertex pairs. Each vertex pair represents the trajectory of a single vertex over a single time step (its position at the start end of the time step). The ring structure contained the vertices, the starting and ending times of each perimeter, and descriptions of the fire behavior and fuel types at each vertex. Thus, during a simulation, each ring



could be retrieved at an arbitrary time step to make calculations for the area bounded by the two fire perimeters. The storage of each ring is complicated by the fact that the number of vertices in the ring changes over time. As the fire get bigger, vertices are inserted along the longest segments. Also, vertices may be clipped out along concave portions of the perimeter or when a fire merges with other fires. The order of the vertices also changes depending on many vector operations within *FARSITE*. All of these complexities require corrections to preserve the geometry of each ring, the trajectories of its vertices, and the area burned.

Methods

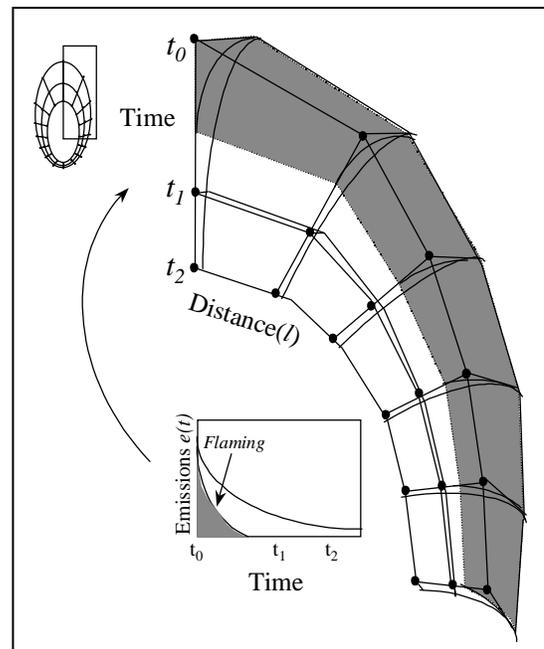
To incorporate Burnup into *FARSITE*, the following steps were required:

1. The Burnup code was modified to compile as a module within the *FARSITE* program,
2. *FARSITE* was modified to store the reaction intensity for each surface fire calculation at each vertex and to calculate the residence time of the surface fire. These values were needed for the ignition pulse to start the Burnup model. They were not needed previously in *FARSITE* because only the progress of the fire front was computed.
3. The *FARSITE* interface was modified to allow spatial data for coarse woody fuels and duff to be loaded and included into the spatial data file (landscape file) used for the simulation. The landscape file was modified to allow these layers to be added and accessed by the simulation. A translation utility was developed to allow backward compatibility with previous versions of *FARSITE*.
4. The process control portion of *FARSITE* was modified to store the fire behavior data from the surface fire and current fuel moistures for all woody fuel components into the data structures mentioned above – as the simulation progressed. This procedure preserves a complete picture of the conditions at each vertex for later running the Burnup model.
5. The process control portion of *FARSITE* was changed to call Burnup for each perimeter segment within each ring structure using the fuel, fire, and moisture data stored for that segment. This part of the program was modified to run in parallel (see below).
6. The combustion history results of Burnup were condensed to facilitate storage and retrieval by the integration routine (below). The raw output from Burnup is a time-series containing thousands of pairs of observations of weight loss or intensity by time step that demands an

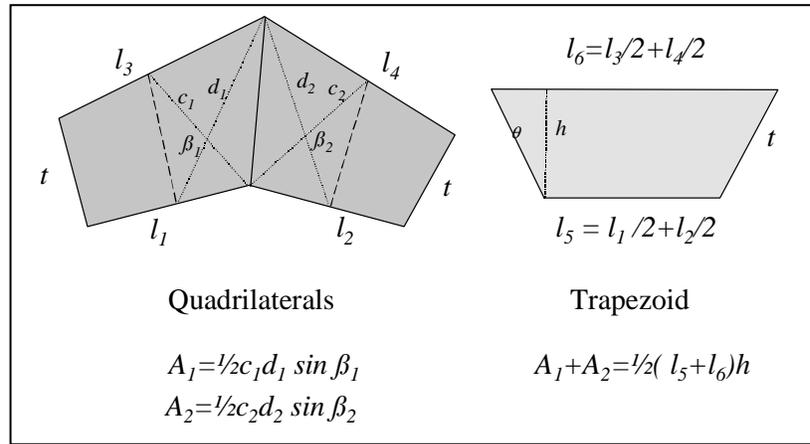
excessive amount of memory storage. Many techniques were tried and those that involved curve-fitting (by non-linear regression, polynomial regression etc) were rejected as too unreliable to operate automatically within the simulation program with no user input. The selected technique samples a few (10-20) pairs of observations from the raw combustion history data. These samples were taken at variable densities depending on curvature (higher densities along regions of higher curvature). The advantage is that it has very small storage requirements and can be accessed efficiently by an interpolation routine for later integration along arbitrary segments of the combustion history curve.

7. A polynomial interpolation routine was developed to access any part of the combustion history curve using the sampled points and provide a value along the curve at that point. The routine concentrates the sample points along areas of higher curvature which are typically located early in the combustion history.
8. A spatial integration was developed to compute the combustion products evolving over the area represented by each vertex. This routine was also adapted to run in parallel.

The task of computing post-frontal products of a spreading fire requires each pair of successive fire perimeters inside the fire front to be visited at each timestep until no more fire activity occurs (indicated by cessation of smoldering in the Burnup results). Each pair of perimeters is stored in a “ring” structure that contains pairs of perimeter vertices that represent the fire front as it moves over that original timestep. The timestep is internal to the *FARSITE* program and varies by how fast the fire is moving. The fire behavior history along each vertex trajectory (e.g. from the 1st to 2nd vertex in the pair) represents a finite area around that line (Figure X). This area is delimited by half the tangential distance to the vertices on each side of each vertex on each perimeter ring. Thus, the total area represented by each trajectory is the sum of the halves of polygons on each side of its path. At each visit, the weight-loss or intensity curve must be integrated over that area. This is done for all active fire rings on all fires at each timestep. The integration is made more efficient by performing it once instead of on each half-quadrilateral independently. Thus, the two quadrilaterals are combined into a trapezoid of equal area; its parallel sides are equivalent in length to the sum of half the distance between the neighboring vertices (Figure X). A trapezoid preserves the change in area over the length of the vertex trajectory that is present in the native quadrilaterals. This change in area is the critical property of the quadrilaterals that permits the combustion history at a given vertex to represent accurately a specific area on the ground and its position within the burning area. Combining the two quadrilaterals into a trapezoid abstracts the actual vertex trajectory into a virtual length equal to the height of the trapezoid. This is irrelevant to the combustion history or to the fidelity of the spatial representation because the time interval is constant across the trajectory and independent of distance the fire travels.



The development of this spatio-temporal integration scheme for the weight loss data from Burnup required two separate integration routines. The first one is needed only the first time an integration is performed for a given fire-ring structure. Each ring must be visited repeatedly during later timesteps until no further combustion activity occurs. All integrations for that ring structure at later timesteps



are accomplished using the second routine. An integration is first required at the end of the timestep in which the ring originated; essentially the fire has expanded from the first perimeter to the second perimeter over that timestep, with the fire edge represented exactly by the second perimeter.

1. The first integration proceeds *forward in time* from the first perimeter toward the second perimeter in steps dependent on the curvature of the combustion curve. This is necessary to capture the increase in burning area by combustion phase (flaming and smoldering) as the fire front moves from the first to second perimeters. The curves to be integrated represent the fraction of original fuel weight remaining and the fraction of flaming combustion. Both the total combustion and the combustion fraction curves are integrated simultaneously within the routine. At each step, the net weight consumed in flaming and smoldering phases is calculated by the difference between the total weight presently consumed over the area and the total weight consumed in the previous step. This is a convenient check on the values calculated for each step in that the total weight consumed cannot be greater than the total for the area.
2. The second integration is called at all later timesteps and proceeds *backward in time* in order to accurately capture the spatial location of the termination point of the combustion and flaming curves. *FARSITE* determines the internal timestep by the rate of fire movement—it is not constant and cannot be predicted. As later successive timesteps call this algorithm, the integration is only working on the dwindling tail of the combustion curves (either the total combustion or the flaming fraction curve). Eventually, the end-point of each curve is captured between the two fire perimeters in the ring. Then the integration becomes only concerned with the fraction of the area still showing combustion.

These integration routines were designed to preserve the spatial and temporal accuracy inherent in the *FARSITE* fire growth data. This ensures that the integration will have an absolute accuracy that is consistent with the space and time resolutions set by the user. In other words, no additional error is introduced to the integration when setting the space and time resolutions coarser. However the simulation will still represent less than the full spatial detail of the fuels and topography using a coarse distance resolution, or perimeter resolution, and be less sensitive to fuel moisture variation with longer timesteps.

Parallelization

The intensive nature of this simulation suggested that large gains in performance could be obtained through parallel processing. The most efficient way to implement parallel processing with a shared memory computer was to use "threads" because they require little operating system overhead compared to separate processes. Each thread would receive unique instructions on which parts of the fire to simulate. Two parts of the post-frontal combustion code were believed to have the most promise of increased performance: running the Burnup model for each vertex, and performing the spatial integration.

The parallelization was performed on a sub-fire basis. That is, each fire was divided into parts based on the number of vertices it contained. Each part was processed separately for running the Burnup model and for the integration. For example, four threads were created for running Burnup, and a different set of four threads were created for doing the integration. An interesting aspect of parallelizing with threads was that the threads only needed to be *created* once -- at the time each parallel routine was first called. The control section for each thread then contained a "do-loop" with an exit condition that could only be met if externally set (to prevent exiting until the user intentionally terminates the simulation, at which time this condition flag is set to exit the loop). The last statement in the do-loop (after all the work is done) is a call that suspends each thread. Subsequent calls to each thread merely resume the threads after having specified new data to work on. This repeats the procedure inside the do-loop without the additional overhead of repeatedly creating and destroying threads. The completion of all threads was synchronized with a "semaphore" so that no new work could be assigned to any thread until all were finished with the previous task.

Simulation Results

Some simple simulations of post-frontal combustion were performed for idealized conditions. These suggest the importance of fuel moisture and diurnal changes in fuel moisture to fuel consumption. Even with constant windspeed, terrain, and fuel type, the changing fuel moistures increase the spread rate, heat output, proportion of flaming and smoldering combustion. Although only "fine" dead woody fuels are sensitive to changes in atmospheric humidity, temperature, and solar radiation on a daily basis, the Burnup model suggests that drier fine fuels will change the ignition, consumption, and consumption rate of larger woody fuels. These fuel moisture effects are probably responsible for the slope changes visible for an otherwise steadily growing fire (Figure X). At about 10AM the drying of fine fuels in response to decreasing humidity and increasing temperature begins increasing the total fuel consumption and spread rate. This causes a noticeable increase in slope, first on the flaming curve and shortly thereafter on the smoldering curve. Smoldering lags the flaming response because the forward edge of the fire is expanding rapidly compared to the reaction time of the flaming phase of combustion.

Recommendations for Further Work

The results of this work highlight the need for improved modeling of moisture content by size class for all woody fuel and duff components of fuels used in the Burnup model. In simulations of large fires over complex terrain and for long time periods, the fuel moistures will likely be variable, spatially, temporally, and within the size-class distribution, dependent on the seasonal timing of the simulation. A fuel moisture model will need to be added to the simulation that is capable of simulating fuel moisture changes for round wood fuel elements of all size classes. Also, the simulations will require a simple set of initializing conditions for the suite of fuel moistures, particularly the large size classes that have long time-lags. Users will likely not be able to estimate moisture contents of all particles for all conditions that vary spatially and temporally. For example, higher moisture contents would be expected under a heavy forest canopy, on a steep north-facing slope at high elevation compared to an open south slope at low elevation. The difference between these conditions would likely change by season, where early season differences could be less than by mid summer.

Emission factors will have to be set for each phase of combustion from the surface fuel complex. Also, emission factors and weight loss will have to be added for crown fire.

There is a huge range of outputs possible from these simulations because of the spatial and temporal nature of post-frontal combustion products. It is likely that many output formats will need to be produced given the different applications for heat and emission data. Possible formats for these products include:

Product	Advantages/Disadvantages	Uses
Summary strip-charts showing intensity or emissions for the fire as a whole (as displayed above)	<ul style="list-style-type: none"> ↑ Easy, fast, available at runtime ↓ No spatial detail, not much info on large fires 	Monitoring of overall fire emissions
Summary charts for arbitrary portions of a landscape	<ul style="list-style-type: none"> ↑ Spatially explicit, user designated at run time ↓ Complex interface design, difficulty in development 	Separate watersheds, fuel types, elevations
Vector data stream on intensity and emissions for individual vertices over time	<ul style="list-style-type: none"> ↑ Easy to output, doesn't take much disk space, good representation of actual data ↓ Irregular point locations, more difficult to input as source to other models, irregular time intervals 	Input to other models, rendering of data, storage of data
Gridded maps at arbitrary time steps showing concentrations of intensity or emissions	<ul style="list-style-type: none"> ↑ Fixed spatial frame of reference for outputs, works well in GIS, can make arbitrary grid resolution ↓ Post-processing of data 	Useful for animation, data storage, smoke dispersion and transport modeling, modeling fire-interactions

	required to produce grid, data requirements are huge (1 grid at each timestep)	
Gridded maps at regular intervals	<ul style="list-style-type: none"> ↑ Fixed spatial and temporal frame of reference, ↓ Can't be shorter than a timestep, loose temporal/spatial detail 	May be required for some kinds of interaction between simulation and wind models, or for animation
Animation of fire characteristics	<ul style="list-style-type: none"> ↑ See the whole picture on fire behavior over space and time ↓ Long rendering times, specialized software probably required 	Good for presentations, technical and non-technical audiences