Sensitivity Analysis of PHOENIX RapidFire

Technical Report

A report prepared by:
Derek Chong (University of Melbourne/Bushfire CRC)
Kevin Tolhurst (University of Melbourne)
Thomas Duff (Bushfire CRC)
Brett Cirulis (University of Melbourne/Bushfire CRC)

Theme:	Understanding Risk
Project:	Fire Impact and Risk Evaluation Decision Support Tool
Component:	
Milestone:	2.4.4
Date:	7 May 2013

This report was produced with financial support from the Bushfire CRC. This is not a published report and has had internal review, but not independent external peer review. Any opinions expressed in this report are those of the authors and not the University of Melbourne or the Bushfire CRC. Any use of original concepts, ideas or results in this report should be done in consultation with the authors.

Executive Summary

An analysis of the sensitivity of the outputs of PHOENIX Rapidfire (PHOENIX) to a range of inputs and simulation parameters was undertaken. This was done using two separate methods; assessment of model response in an artificially generated idealised landscape and assessment using case-studies of real fires.

The ideal landscape was used to evaluate model sensitivity in response to temperature, relative humidity, wind speed, fuel type and wind direction relative to slope. The model was evaluated under two sets of weather conditions, mild (representing moderate fire spread potential) and extreme (representing high fire spread potential). Each scenario was evaluated for each of two fuel types, forest and grass. Sensitivity was evaluated in terms of the gross area burnt when the input of interest was systematically changed while all other inputs were held constant. For all evaluations except relative wind direction, model sensitivities were compared to an equivalent area burnt using the corresponding McArthur Forest or grassland fire danger meter (assuming an elliptical fire shape). The combination of wind direction and slope resulted in simulated fires that were not elliptical, so comparisons with shapes generated with the fire danger meters were not valid. PHOENIX predictions differed from those generated using point estimates for some circumstances; however without further investigation it is unclear on what is causing these differences. Differences in predictive performance are not necessarily representative of model error, as there are a number of differing assumptions between the systems used. However, specific situations have been flagged for follow up work.

Two case study areas were used for PHOENIX model sensitivity evaluation; Wangary and Kilmore. Case study fires were simulated using observations from the day that the fires occurred with one input systematically varied. Three inputs were evaluated using the case studies; simulation resolution, start time (simulating fire ignition to occur earlier and later than observed) and start location (varying the ignition location in space). Sensitivity was evaluated by considering the change in the Area Difference Index (ADI, an index of the ratio between incorrectly predicted burnt area and the correctly predicted burnt area) from the baseline scenario (simulation resolution of 180m, ignition location and time as observed. Predictive performance varied wide with changing inputs. In general as the difference in input value to the 'best estimate' increased, predictive performance degraded.

Table of Contents

Artificial Landscape

An artificial 'Ideal' landscape was used to test the sensitivity of PHOENIX Rapidfire (PHOENIX) to changes in temperature, relative humidity (RH), wind speed, fuel load and wind direction relative to slope. The artificial landscape had homogenous fuel, constant weather, consistent topography (flat, or in the case of slope evaluation, consistently sloped but at a fixed elevation). As PHOENIX currently incorporates different spread engines based on fuel, two fuel types were evaluated; forest and grass. For sensitivity evaluations, the variable of interest was systematically changed while all other landscape attributes were held constant. For each evaluation, PHOENIX was run for four hours and the gross area burnt recorded.

Evaluations were undertaken using two different scenarios, mild (representing easily suppressible fire behaviour) and severe (representing fires occurring under conditions where suppression would not be effective). The parameters used for each scenario are summarised below.

<u>Mild Conditions</u>

Temperature: 20^{o}C
Relative Humidity: 50%
Wind Speed: 20 km/h
Drought Factor: 10
Curing: 100%
Cloud Cover: 0%
FFDI: 7 (Moderate)

<u>Severe Conditions</u>

Temperature: 35^{o}C
Relative Humidity: 15%
Wind Speed: 40 km/h
Drought Factor: 10
Curing: 100%
Cloud Cover: 0%
FFDI: 61 (Extreme)

Grass fuel loads were assumed to be 'grazed' with an effective load of 4 tonnes /ha. Forest fuels were assumed to have a surface fuel load of 15.8 tonnes/h, an elevated fuel load of 1.8 tonnes/h and a bark load of 5 tonnes/h. A wind reduction factor of 1.2 was used for grass fires and 3.5 for forest fires. Fires were assumed to ignite at 1300h and burn for 4 hours.

For comparative purposes, for each evaluation the corresponding one dimensional fire model (the McArthur Forest Fire Behaviour MK5 model for forest or the CSIRO grassland model for grass) was used to compute a head-fire travel distance for four hours equilibrium spread. Using an assumption of elliptical fire spread with zero wind speed conditions for lateral spread, the head-fire travel distance was converted to an estimate of burnt area. The use of Huygen's algorithm of fire spread in PHOENIX does not result in perfectly elliptical fires, so some differences in growth rates are expected between the elliptical conversion of the one dimensional model and PHOENIX.

Model sensitivities are presented in terms of the absolute area burnt and the rate of change in area burnt (change in area burnt per increase in the variable of interest).

Temperature

Grass Fuel type

Both the CSIRO Grass meter and PHOENIX showed consistent responses to changes in temperature. The area affected was found to be slightly higher using the CSIRO Meter, however the rate of change was identical for both models for both Mild and Severe conditions, with an exponent of 1.2 for an increment of 5° C (Figs 1 & 2). The difference in absolute area is likely to be a function of the 'build-up' phase and the influence of the time of day considered in PHOENIX.

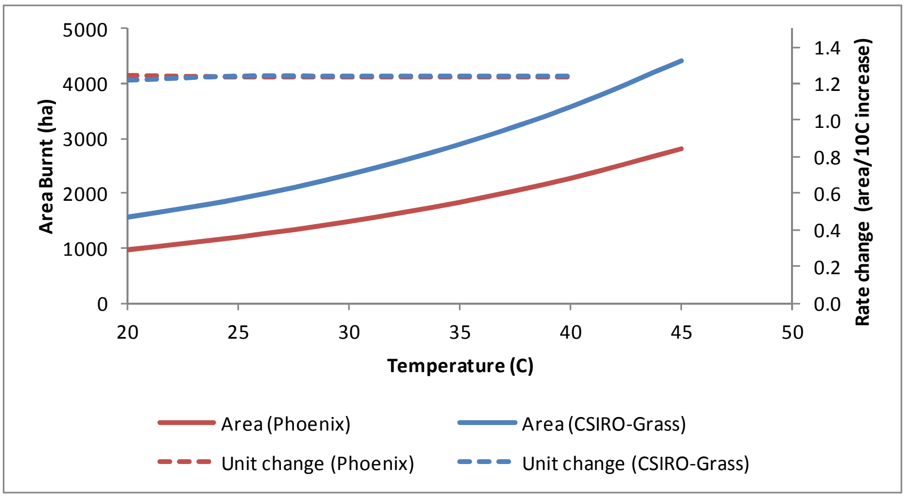

Figure 1 Grassland fuel type area burnt in response to systematically changing temperature under mild conditions.

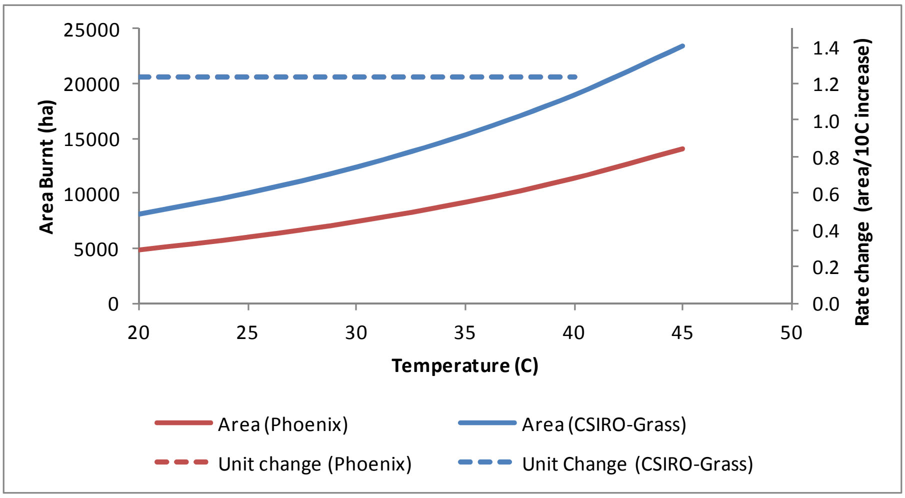

Figure 2 Grassland fuel type area burnt in response to systematically changing temperature under severe conditions.

Forest Fuel type

In forest fuels there were substantial differences between the results of the McArthur MK5 and PHOENIX. The McArthur MK5 showed a consistent response to cross the entire range of temperatures assessed. Under severe conditions the absolute area burned increased however the rate of increase remained constant (Figs 3 & 4), with an exponent of 1.4.

In contrast, under mild conditions for PHOENIX simulations, the burnt area increased with increasing temperature, however each unit increase in temperature had a diminishing effect on the total area burnt. As temperatures increased above 35°, a unit increase in temperature had a limited effect on fire size (Fig 3). The rate of change of the exponent was -0.082 per 5° increase in temperature. The cause of the disparity between models is unclear at this stage, and has been flagged for additional analysis.

Under severe conditions, PHOENIX exhibited greater burned areas for equivalent conditions in comparison to the McArthur MK5 (Fig 4). This is somewhat expected, as substantial work has gone into the simulation of the spotting process in PHOENIX, which had long been a recognised limitation of predicting fire spread with McArthur. However, as with the analysis for mild conditions, the rate of increase of affected area per unit increase in temperature appears to decline above 35°. This is likely to be due to the same mechanism; however the cause will be investigated. The rate change in area burnt for changing temperature under severe condition was inconsistent, most likely due to the interaction between surface fire and spotting driven spread.

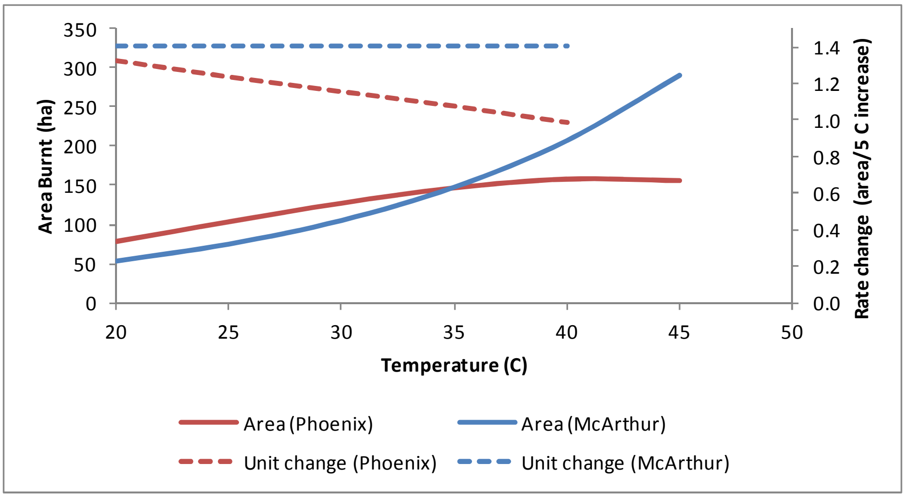

Figure 3 Forest fuel type area burnt in response to systematically changing temperature under mild conditions.

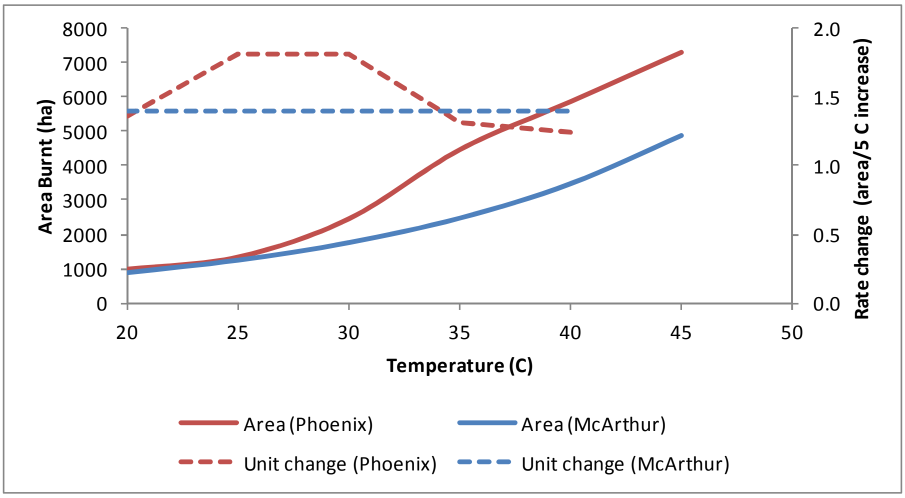

Figure 4 Forest fuel type area burnt in response to systematically changing temperature under severe conditions.

Relative Humidity

Grass Fuel type

In grass fuel, both the CSIRO Grass meter and PHOENIX showed consistent responses to changes in RH. The area affected was found to be slightly higher using the CSIRO Meter, however the rate of change was identical for both models for both Mild and Severe conditions, with an exponent of 0.863 for an increment of 5 % (Figs 5 & 6). The difference in absolute area is likely to be a function of the 'build-up' phase and the influence of the time of day considered in PHOENIX.

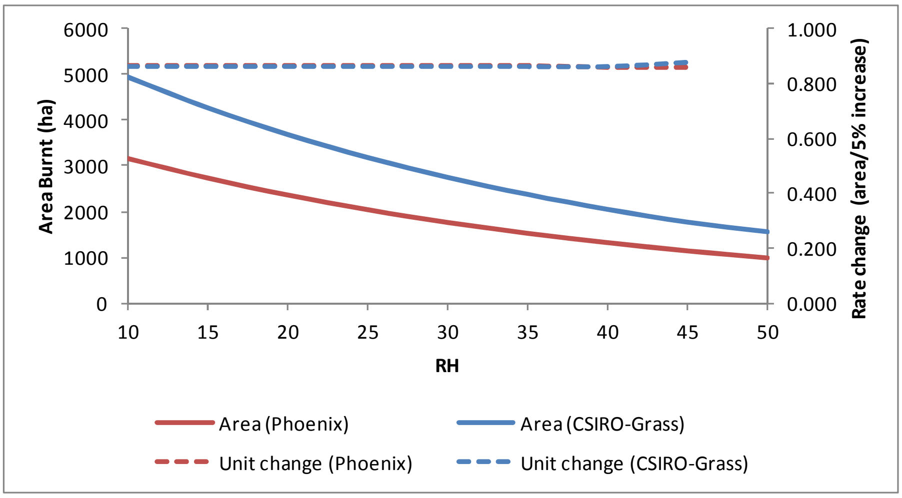

Figure 5 Grass fuel type area burnt in response to systematically changing RH under mild conditions.

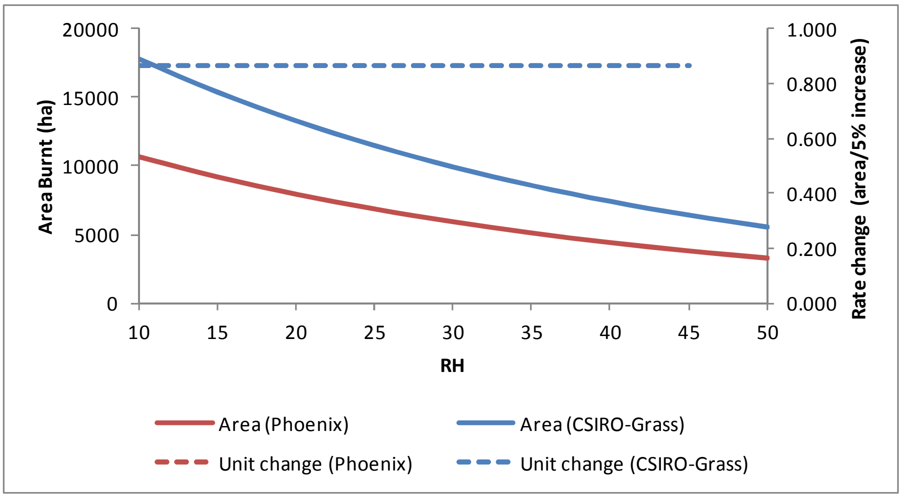

Figure 6 Grass fuel type area burnt in response to systematically changing RH under severe conditions.

Forest Fuel type

In forest fuels there were differences evident between the results of the McArthur MK5 and PHOENIX. The McArthur MK5 provided consistent predictions for mild and severe conditions, with a constant rate of decrease for per unit increase of RH. The exponent for McArthur under both sets of conditions was 0.7 per 5% increase (Figs 7 & 8).

PHOENIX consistently affected a larger area, particularly at low RH values. It is likely that this is due to the enhanced consideration of spotting in PHOENIX; spotting contributes to fire spread and the algorithms for calculating spot-fire ignition use RH as a direct input. As RH increases, the rate of change in burnt area for phoenix predictions approaches that of the McArthur MK5 of around 0.7, indicating a similar sensitivity to unit changes.

Under severe conditions, there was noise evident in the PHOENIX predictions, with the rate of change varying by +/- 0.1 (Fig 8). This is likely to be due to the nature of the spotting algorithm, which can be sensitive to grid conditions. As the affected area was greater with PHOENIX under low RH values, the PHOENIX sensitivity to RH in forest fuels has been flagged for further investigation.

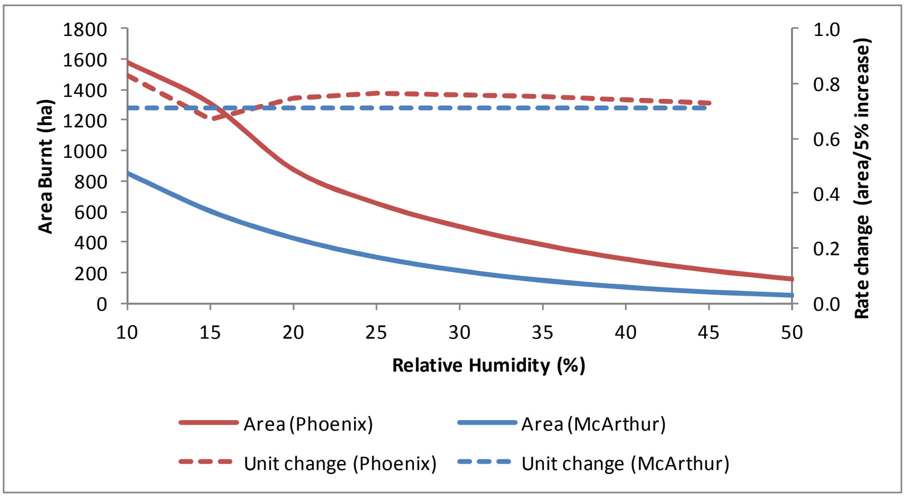

Figure 7 Forest fuel type area burnt in response to systematically changing RH under mild conditions.

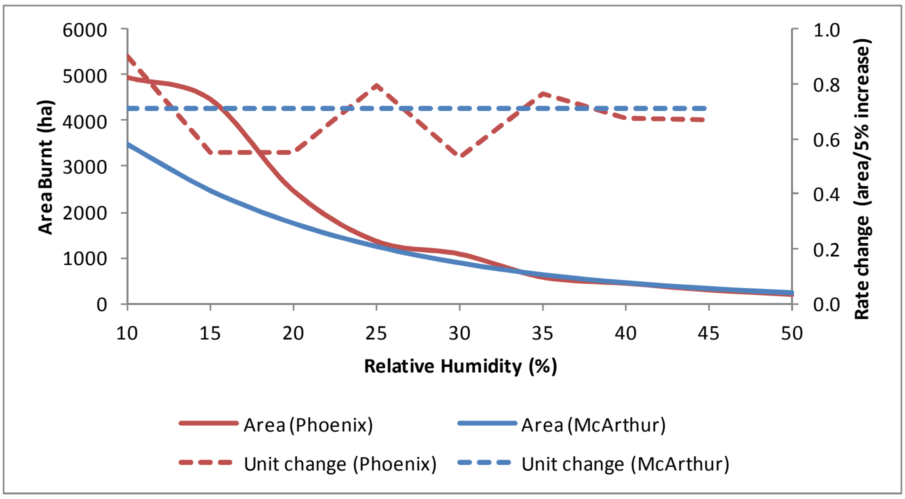

Figure 8 Forest fuel type area burnt in response to systematically changing RH under severe conditions.

Wind Speed

Grass Fuel type

In grass fuel, both the CSIRO Grass model and PHOENIX exhibited similar responses to increasing wind speeds, with the exception of zero wind (Figs 9 & 10). As expected, severe conditions resulted in a greater overall area burnt in both cases. As the algorithm is not designed for spreading fire in the absence of wind, this result is not unexpected. Robust predictions cannot be expected from any model when predicting outside its development range. Above wind speeds of 10 km / h, the rates of response for PHOENIX and the Grass model are identical (approaching 1.1 / 10km/h), however the absolute area burnt by PHOENIX is slightly lower. As with temperature, this is likely to be due to the 'build-up' phase and the influence of the time of day considered in PHOENIX. ,

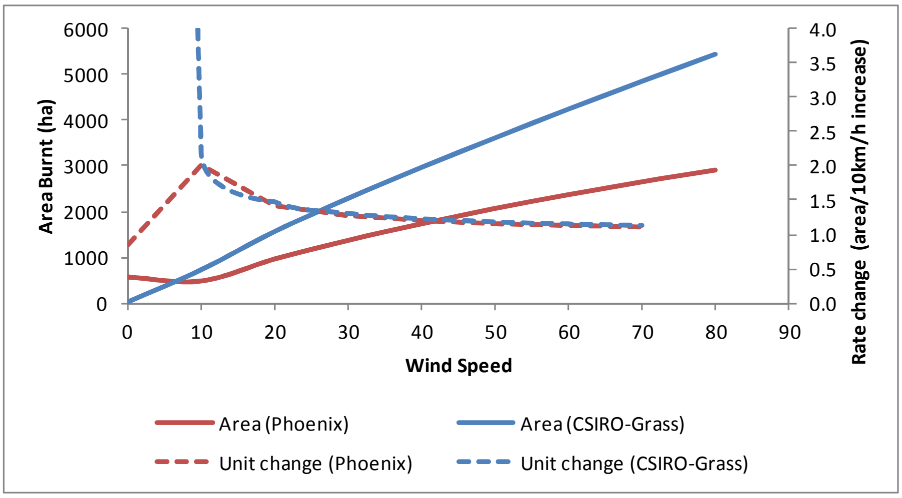

Figure 9 Grass fuel type area burnt in response to systematically changing wind speed under mild conditions.

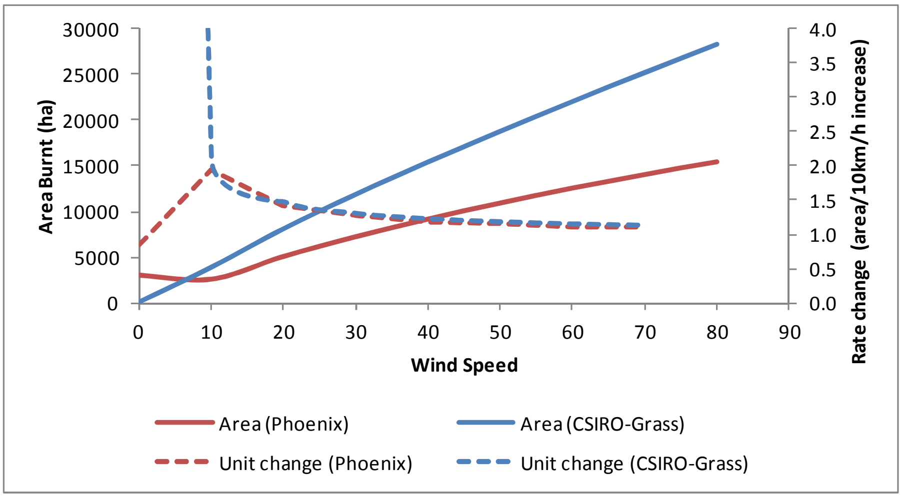

Figure 10: Grass fuel type area burnt in response to systematically changing wind speed under severe conditions.

Forest Fuel type

In forests fuels, patterns of area burnt when wind speed was systematically changed were similar to those identified for grass fuels (Figs 11 & 12). Both models had issues when predicting for zero wind speed, however this wind speed is outside the range for which the models were designed. Above this, both models had very similar results, with unit change exponents of 1.2 per 10 km/h increase. PHOENIX predicted higher burnt areas, with a moderate degree of noise under severe conditions (fig 12). A key difference between PHOENIX and the McArthur Mk5 meter is the inclusion of the contribution of spotting to spread in PHOENIX. Spotting is an important mechanism of fire spread under more extreme conditions, so the effect would be expected to be more pronounced in the severe evaluation. Wind is a transport vector for embers in the PHOENIX spotting model, and so the increased area affected under severe condition is not unexpected. The variability in the prediction result is likely to be due to the interaction of the spotting model grid and the underlying source data. Methods of reducing noise are in the process of being investigated.

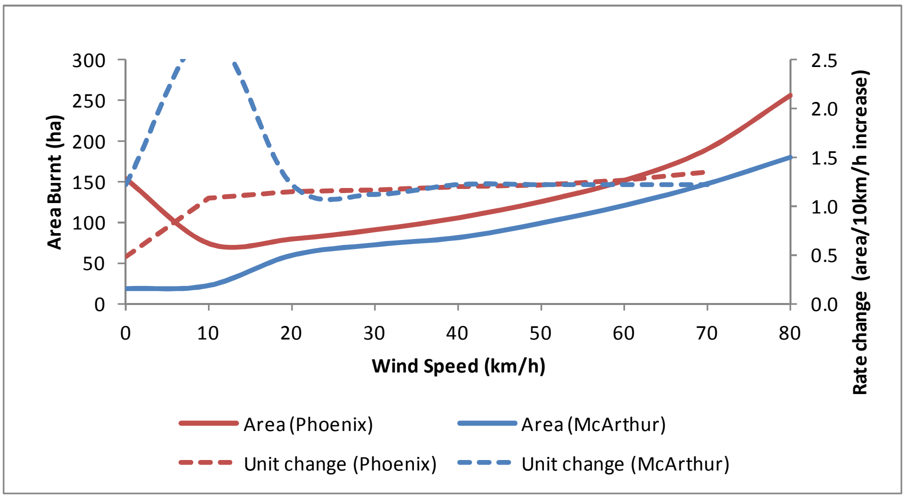

Figure 11: Forest fuel type area burnt in response to systematically changing wind speed under mild conditions.

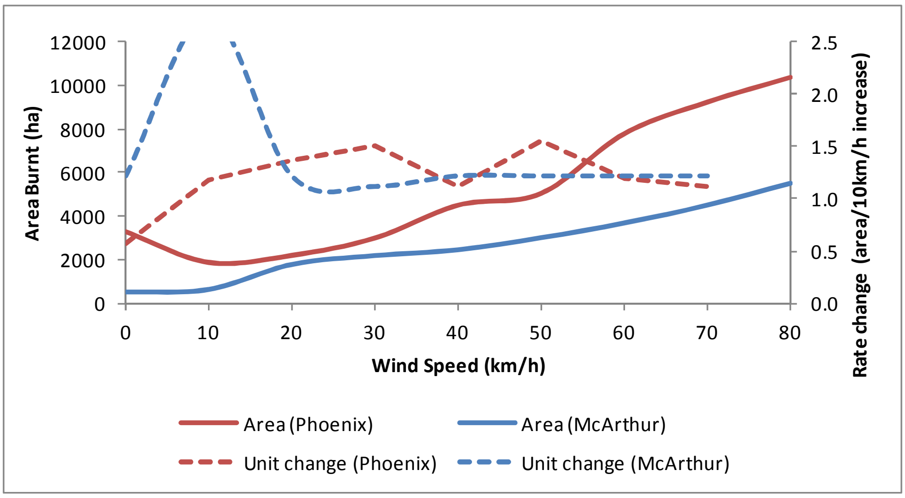

Figure 12: Forest fuel type area burnt in response to systematically changing wind speed under severe conditions.

Fuel Load

The influence of fuel load on fire affected areas was evaluated by systematically varying the total fine fuel load under a consistent set of environmental conditions. Fires were simulated in grass and forest fuels with the following input parameters:

Temperature: 35^{o}C
Relative Humidity: 15%
Wind Speed: 20 km/h
Drought Factor: 10
Curing: 100%
Cloud Cover: 0%
FFDI: 61
Wind reduction factor: 3.5(forest) and 1.2 (grass)

Grass fuels were simulated as loads of 2, 4 6, 8 and 10 tonnes per hectare. Forest fuels were simulated as loads of 10, 15, 20, 25 and 30 tonnes per hectare. For forest fuel loads, the fine fuel load was distributed between surface, elevated fuel and bark using a ratio of 9:1:3. A wind reduction factor of 1.2 was used for grass fires and 3.5 for forest fires. Fires were assumed to ignite at 1300 h and burn for 4 hours.

For comparative purposes, outputs of the CSIRO grassland fire danger meter and the McArthur MK5 meter were used to generate idealised ellipses for conditions corresponding to those used in the PHOENIX analysis. The CSIRO Grassland fire model does not formally incorporate fuel load, however for the purposes of comparison the three categories ‘natural’, ‘grazed’ and ‘eaten out’ were considered equivalent to fuel loads of 1.5, 4 and 6 tonnes per hectare. As only 3 values were used for the consideration of the CSIRO grassland meter, the rate of change was not computed.

Grass Fuel type

The CSIRO grassland model exhibited a linearly increasing affected area with increasing fuel load (Fig 13). This is in contrast to the PHOENIX output which showed a positive but diminishing effect with increasing fuel load. Observation of the PHOENIX output indicated that lateral spread was lower that was predicted with the CSIRO model, and is in indication that the fire shape assumptions (both of elliptical growth with the CSIRO model and the spread pattern of PHOENIX) may need more consideration. As the CSIRO model fuel categories are intended to take into account fuel structure as well as load, the CSIRO model outputs are a combination of both properties, which may be contributing to the differences to in the model results. Above 6 tonnes / ha fuel load, the area burned in the PHOENIX model was predicted to decline. This load is beyond the range for which the input modules within PHOENIX were developed so the result is extrapolation. The effective range of the PHOENIX model for grass fuel loads has been noted.

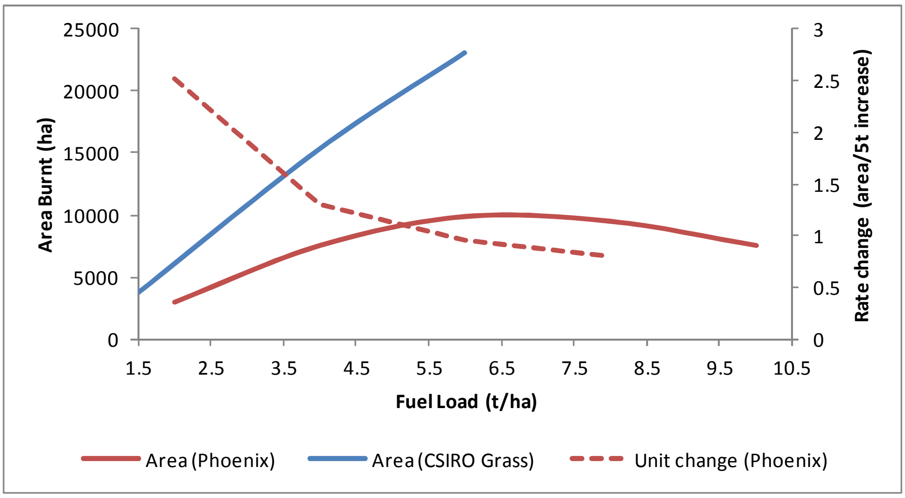

Figure 13: Grass fuel type area burnt in response to systematically changing fuel load.

Forest Fuel type

In forest fuels, both the McArthur MK5 model and PHOENIX exhibited increasing burnt area in response to increasing fuel loads (Fig 14). At low fuel loads, the absolute predicted areas between the two models were similar, however PHOENIX exhibited a greater rate of increase with increasing fuel load, resulting in approximately double the affected area at a fuel load of 30 t/ha. An important feature of PHOENIX is the ability to discriminate between different sources of fuel. At low calculated fire intensities (such as when fuel load is low), only the litter fuel will be consumed as a fire burns. As the intensity increases, the elevated and bark fuels will also be burned, accelerating spread rates. An important mechanism of fire spread at high intensities is the initiation of new fires through the transport of embers. As PHOENIX spreads fires spatially, this ember transfer can be simulated and can materially contribute to increasing fire affected area. It is likely that the increase in difference affected area above 25 t/ha is due to the simulation of spotting due to ember transport.

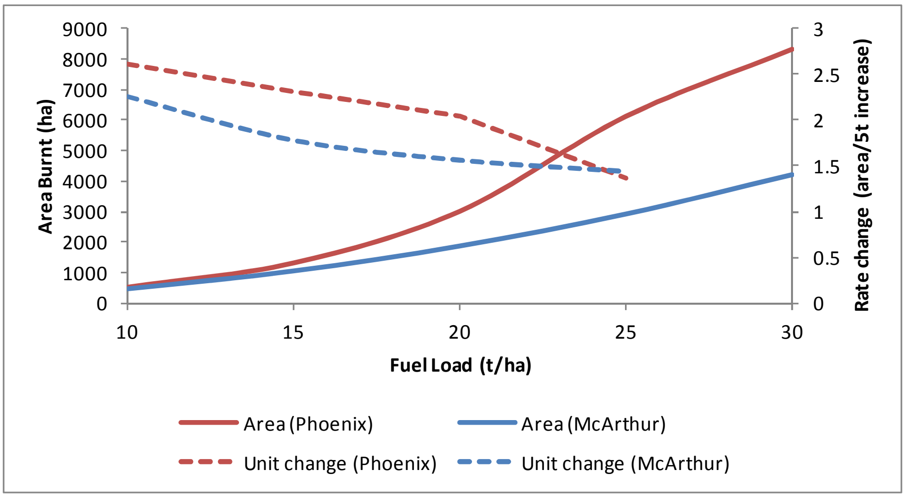

Figure 14: Forest fuel type area burnt in response to systematically changing fuel load.

Wind Direction and Slope

The sensitivity of PHOENIX outputs of wind direction to slope was evaluated by creating an artificial landscape and varying wind direction and slope in a factorial manner. As PHOENIX recognises elevation in calculations, a code modification was made to allow a sloped landscape that was had constrained at a constant elevation. As the elliptical transformations of the CSIRO Grassfire meter and the McArthur meters are unable to recognise changes in fire shape, their outputs in this situation would be unrealistic and a comparison with PHOENIX would be inappropriate. The parameters used for assessing the interaction between wind direction and slope are summarised below:

Temperature: 35°C
Relative Humidity: 15%
Wind Speed: 20 km/h
Drought Factor: 10
Curing: 100%
Cloud Cover: 0%
Fuel: Grassland (5 tonnes/ha) and Forest (23 tonnes/ha)
FFDI: 61
Aspect: Northerly
Altitude: 0 (PHOENIX RapidFire modification required)
Wind Reduction Factor: 1.2 (Grassland), 3.5 (Forest)
Burn time: 1300-1700 h

Evaluations were undertaken in both forest and grassland fuels. Slope was varied been 0 and 30° in 10° increments. Wind direction was varied from 0 to 90° in 30° increments and combine with slope in a factorial manner.

Fires burn most rapidly in the direction they area pushed by the wind and when travelling up slope. In grass, the combination of these two influences results in the greatest area burnt (Fig 15). However there is some interaction between the wind direction and slope; fires are a similar size when winds are blowing up to 30° from the predominant slope direction. When the predominant wind direction is offset from the slope direction, the head-fire is likely to spread more slowly. However the affected area remains similar due to the fire becoming wider due to the effect of slope. This is particularly evident in Fig 15 when the slope is 30°, as it has a proportionally greater influence.

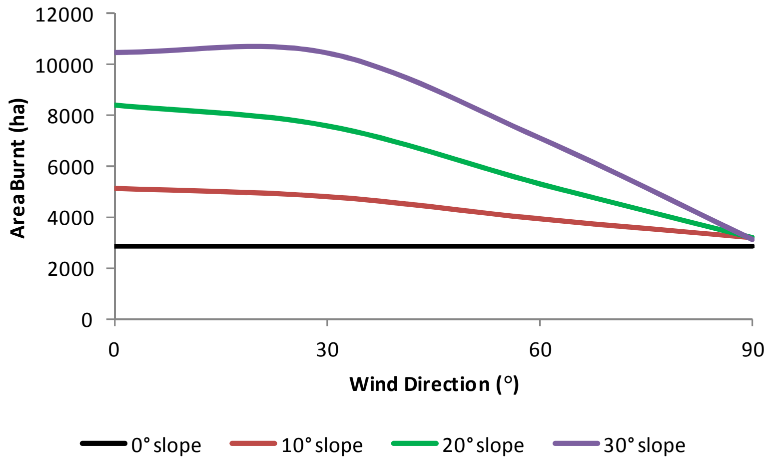

Figure 15: Grassland fuel type area burnt in response to systematically changing wind direction and slope.

Forests have a higher level of wind interception than grasslands, so beneath a forest canopy, fires spread proportionally more slowly. Consequently, the affected areas are much lower than fires simulated in grass. For slopes of 10 and 20°, the forest affected area followed similar patterns to the grass, however at 30° slope there was a substantial departure. At 30° slope when crosswinds (90°) were present, the affected area was greatly increased (Fig 16).

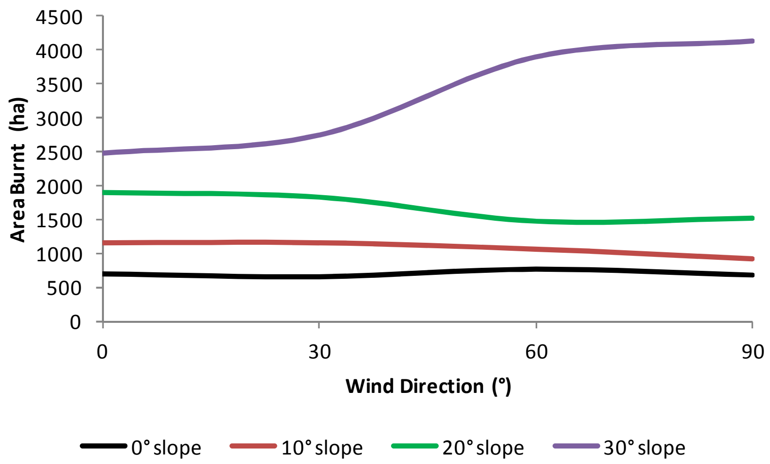

Figure 16: Forest fuel type area burnt in response to systematically changing wind direction and slope.

Follow up investigation found this to be due to ember driven spotting. Spotting is proportional to fire intensity and the amount of bark fuel. As there is no bark in the grass fuel type, spotting is not simulated to occur. In the forest fuels, when wind was blowing upslope, the resultant fires were very long and narrow (Fig 17). When the wind was blowing across slope, fires were shorter and wider. However, as slope contributes to fire spread, fire intensity was higher when slope was greater. When slopes were at 30 °, it appears a threshold was reached where the calculated fire intensity was great enough to incorporate all fuels, including bark. The consequent inclusion of ember driven spot-fires greatly increased the width of the fire, resulting in a much greater burnt area (Fig 18).

There was some noise in the result of the 0° slope result for forests. This is potentially a result of interaction between the spotting simulation grid and the fire spread simulation grid. This has been flagged for follow-up.

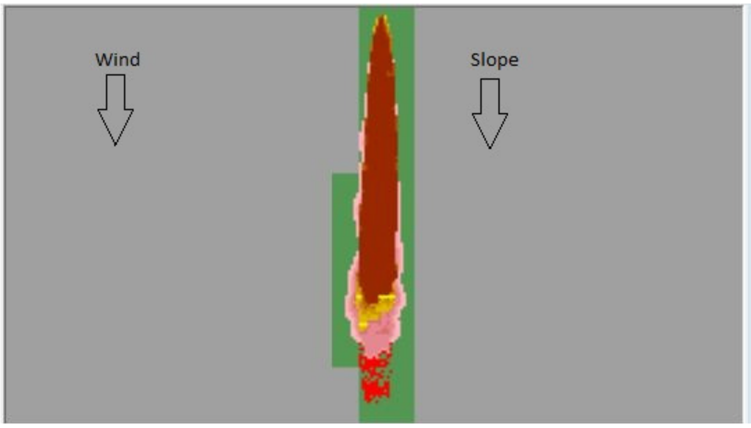

Figure 17: Fire modelled in PHOENIX under consistent slope and wind conditions. Slope is 30°, wind offset is 0°. Embers are shown in red.

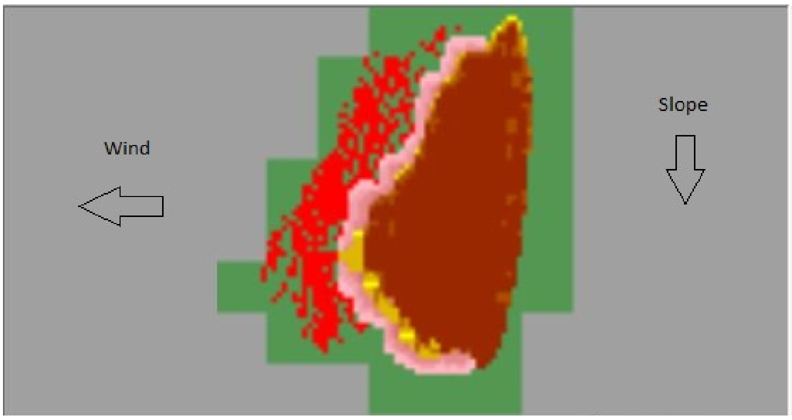

Figure 18: Fire modelled in PHOENIX under consistent slope and wind conditions. Slope is 30°, wind offset is 90°. Embers are shown in red.

Summary of Artificial Landscape Results

The sensitivity of PHOENIX was evaluated for a range of conditions in an artificial, homogenous landscape by systematically varying inputs. Some departures from predictions made using the CSIRO grassland fire danger meter and the McArthur MK5 meter were evident. This is not unexpected as there are a number of design feature of PHOENIX that make it differ from static, point based spread models. Included in these are:

- The consideration of time-of-day within PHOENIX. PHOENIX recognises time of day and compensates for diurnal temperature, relative humidity changes and solar radiation.
- The inclusion of spotting within PHOENIX. PHOENIX includes a simulation process for emulating the creation, transport and ignition of embers. The generation and transport of embers is a function of fire intensity, wind speed and bark fuel load.
- Dynamic fuel incorporation within PHOENIX. PHOENIX recognises different fuel strata and will incorporate different fuels dynamically depending on fire intensity. Under mild conditions, not all fuel will necessarily be burnt by a fire.
- The consideration of multiple fire models within PHOENIX. Under mild conditions (FFDI<12), PHOENIX will use the McArthur model for prescribed burning conditions rather than the McArthur MK5 model. These two models were designed for different purposes and are not perfectly matched.
- The incorporation of a ‘build-up’ phase within PHOENIX. PHOENIX simulates the development of a fire from a point. These results in slower growth for the first 30 minutes of spread, and may consequently result in smaller overall fire sizes. The CSIRO grassland and the McArthur MK5 models provide equilibrium spread rates and do not recognise changes in fire properties through time.

However, In the process of sensitivity analysis, there were a number of unexpected outcomes that have been flagged for follow up. Many of these were outside the range of normal model use, however in the interest of developing robust predictions, these will be investigated.

Case-study Fires

The sensitivity of PHOENIX predictions to changes in inputs at real fires was evaluated. Datasets were compiled for three fire complexes; Wangary (SA), Warragamba (NSW) and Kilmore (VIC) as part of the Bushfire CRC fire DST project. Due to the data limitations with the Warragamba fire, in particular uncertainty regarding ignition time, ignition location and progression pattern, evaluations were done only for the Kilmore and Wangary fires.

For sensitivity testing, predicted fire areas generated using PHOENIX were compared to the corresponding observed fire areas. Comparisons were made using the Area Difference Index (ADI), which is an index of the incorrectly predicted area divided by the correctly predicted area. Accurate predictions will have an ADI close to 0, as predictive performance decreases, the ADI will increase. An ADI of 1 indicates that the incorrectly predicted area (combining both under prediction and over prediction) is equivalent to the correctly predicted area (the area where the prediction and the observed burnt area intersect). Comparisons are presented in terms of the ADI percentage change relative to the 'best estimate' case. Three properties were evaluated using the case-study fires; start time, start location and simulation grid size.

The Wangary fires used for model evaluation stemmed from a wildfire that occurred on 10 January, 2005. The fire was not fully contained and on 11 January, four ignitions occurred outside containment lines, resulting in a number of fast moving fires that eventually coalesced into a large fire complex. For the purpose of evaluation, the two most southerly fires (described here as fire one, which started at 9:51 am and fire two, which started at 10:00 am) are simulated and evaluated until the point where they join. In addition, the entire fire complex (with all four ignitions) was modelled and evaluated. For the entire complex, the timing of each ignition that contributed to final fire was varied. The weather stream used for evaluation was constructed from AWS data and field observations. The baseline ADI for fire 1 was 0.85, fire 2 was 1.90 and for the entire complex was 0.45.

The Kilmore fire occurred on February 7, 2009 near the township of Kilmore, Victoria. The fire was started by an electrical fault, with ignition estimated to have occurred at 11:47 am. The fire spread rapidly, driven by strong northerly winds. At approximately 5 pm, a wind change occurred, driving the fire to the east. As this resulted in a substantial shape change, evaluation was done on the spread of the fire up until the time of the wind change. PHOENIX simulations are based on observations recorded at the Kilmore gap AWS. The baseline ADI for the Kilmore fire was 0.77.

Start Time

A total of 11 different scenarios were run for each fire being considered. With all other inputs held constant, start time was varied earlier and later by 5, 10, 20, 40 and 60 minutes. An exception was made for the Kilmore fire, which was run 50 minutes rather than 60 minutes earlier due to constraints in the weather dataset. For the Wangary complex fire, the timing of each ignition that contributed to final fire was varied. The ADI was calculated at the same point in time for each simulation.

The results for fire 1 are presented in Fig 19, for fire 2 are presented in Fig 20 and the Wangary complex in Fig 21. The results show that in general, where start time is varied from the observed value, the final prediction will be much poorer in quality, with ADI values increasing greatly. As the ignition time moves farther from the observed time, predictive power decreases. Fire 2 was an exception to this, with predictions that used ignition times up to 40 minutes early exhibiting better performance than the simulation that used the actual ignition time. Delayed ignition times had a particularly adverse affect on predictive performance, with a large proportion of error relative to the correctly predicted area. It is likely that sensitivity to start time found for fire 1 and fire 2 is partially due to the short burn times allowed before the ADI was sampled. The ADI was calculated based on the observed fire perimeter at a point in time before the fires joined into a large complex. As a result the fires had been burning for less than four hours, and consequently the time values used for sensitivity testing were a relatively high proportion of burn time.

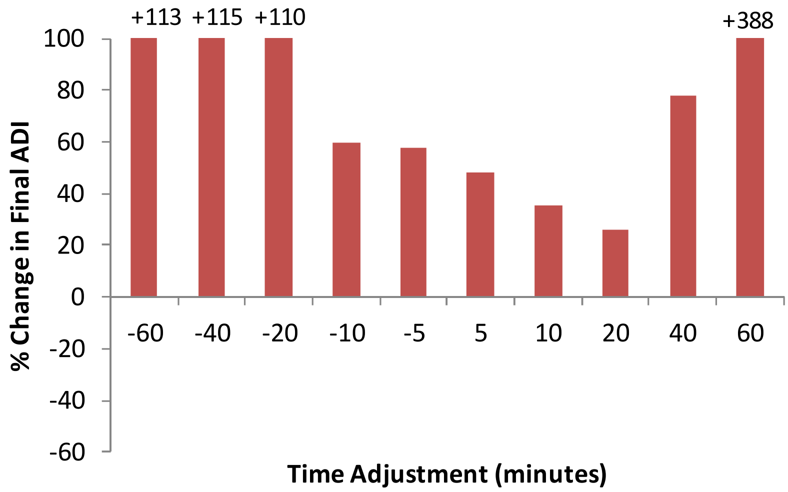

Figure 19: Wangary fire one percentage change in ADI due to start time alteration (Baseline ADI =0.85).

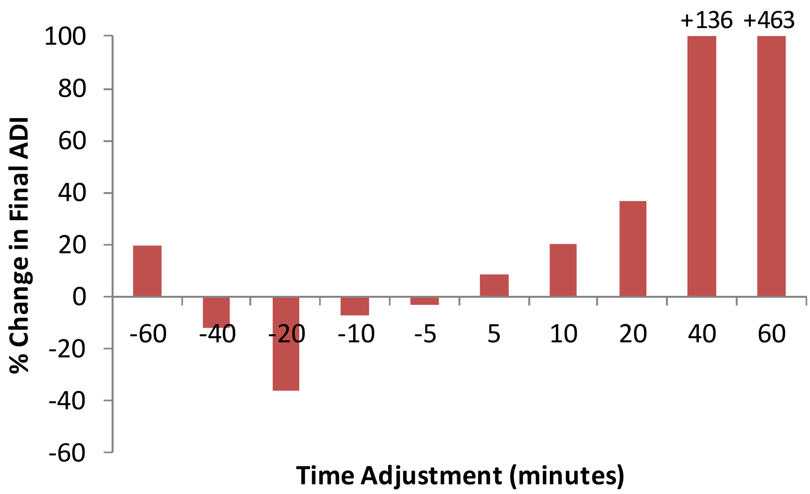

Figure 20: Wangary fire two percentage change in ADI due to start time alteration (Baseline ADI =1.90).

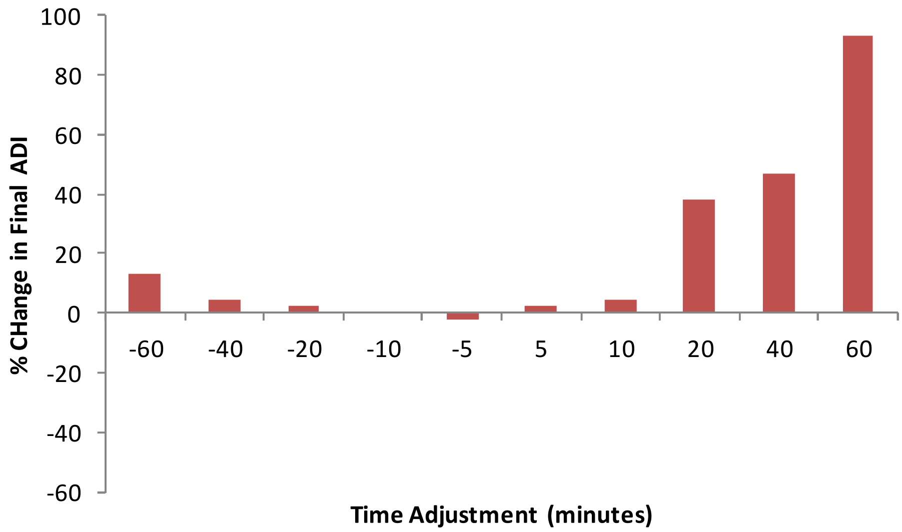

Figure 21: Wangary fire complex percentage change in ADI due to start time alteration (Baseline ADI =0.45).

In the Kilmore study, start time seemed to have a much smaller effect on the overall area affected, with the final ADI robust across ignition time variation values (Fig 22). In all cases the change in predictive performance was less than 40%, even when the start time was changed by an hour. This may be due to overall duration, size and complexity of the fire. In contrast to the Wangary fires, the length of time that the Kilmore fire burnt before having the ADI calculated resulted in a change of an hour being relatively smaller in proportion. Interestingly, most of the predictions that used different start time than that observed achieved a better match to the observed fire shape. The contrasting result between the Kilmore and Wangary fires is indicative of the complex nature of real fires and the need to consider site-specific influences.

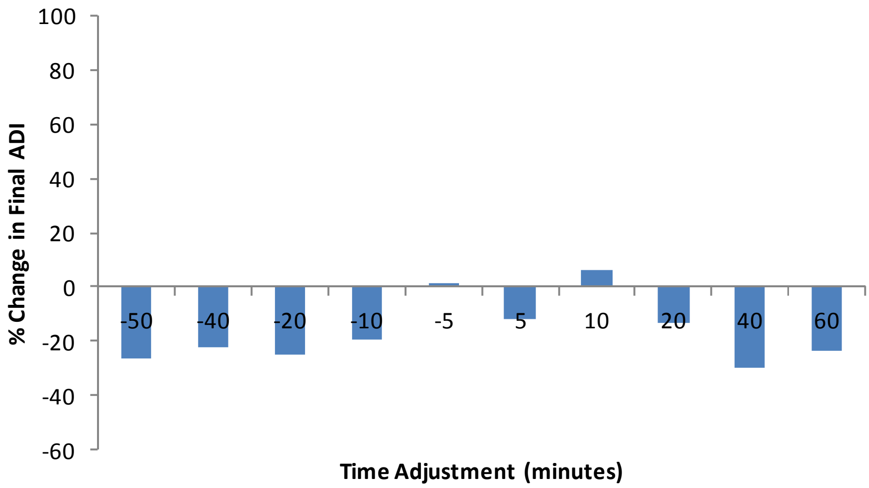

Figure 2: Kilmore fire complex percentage change in ADI due to start time alteration (Baseline ADI =0.77).

Ignition Location

For the Kilmore and Wangary case-study fires, start location was altered by 100m, 500m and 1000m in four cardinal directions (N, S, E, W). The Wangary fires were simulated as a single complex with all ignition locations moved together for each simulation.

The results for the Wangary fire complex are presented in figure 22. The impact of moving the ignition location varied with both scale and direction. Interestingly, for simulations where the ignition locations were moved in any direction a distance of 100m there was an improvement of fit. As the burnt area is the result of a spatially explicit process that travels through a complex environment driven by highly variable weather, it is likely that this improvement is the result of coincidence rather than any process inherent within PHOENIX. The fact that the Wangary fire complex is an aggregation of four different ignitions may make the final result more robust to changing ignition location. As expected, the result becomes more variable as the distance the ignition location is moved increases.

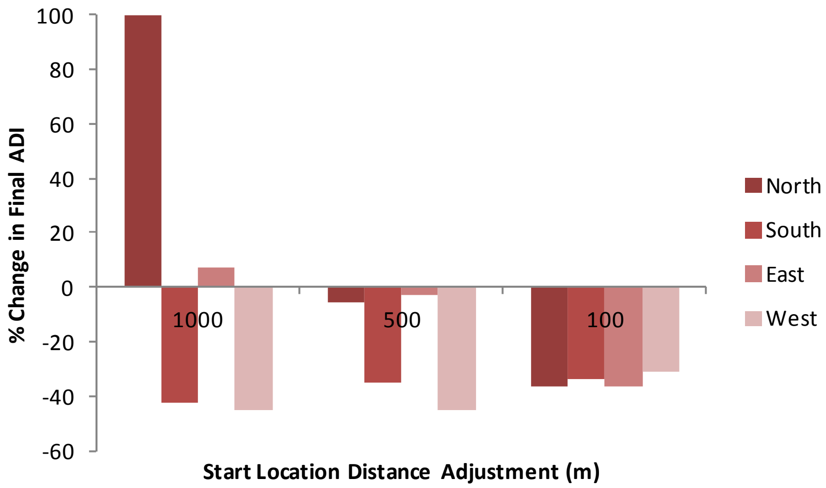

Figure 22: Wangary complex change in predictive performance due to variation in ignition location

The Kilmore fire exhibited substantial variation in ADI, generally showing poorer performance as the movement distance of the ignition point increased (Fig 23). However here were some small improvements when the ignition point was moved 100m. As PHOENIX uses data grids of 180m for simulation, any change resulting from a move smaller than this is likely to be grid / data interactions rather than any overarching difference in simulation properties. However, as some landscape properties, including linear barriers and spotting, are considered at absolute scales, this will be further investigated.

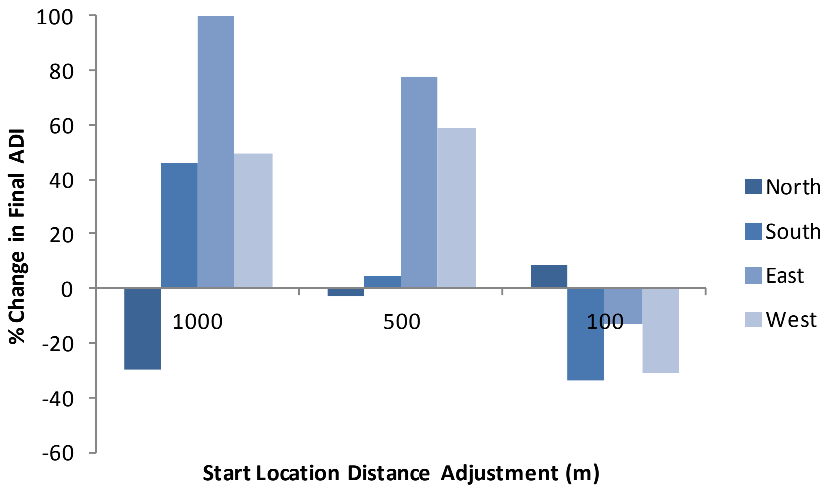

Figure 22: Kilmore change in predictive performance due to variation in ignition location

Cell-Size

PHOENIX processes landscape data as grids. Inputs are typically raster grids representing environmental properties including fuel and topography. Regardless of the input data scale, PHOENIX resamples the input grids and converts the data to consistent scales. Standard simulation grid resolution is 180m. The importance of the PHOENIX grid scales was evaluated by systematically varying the analysis scale. Analysis cell resolutions of 30, 60, 90, 180, 360, 540 and 720 metres were processed. Fires were compared to the standard (180m) resolution as a percentage change in ADI value (comparing the simulated result with the observed perimeter).

The Wangary complex was evaluated as a single fire event (Fig 23). Any departure from the standard cell resolution resulted in a poorer fit and the performance was degraded as the difference to the standard resolution increased.

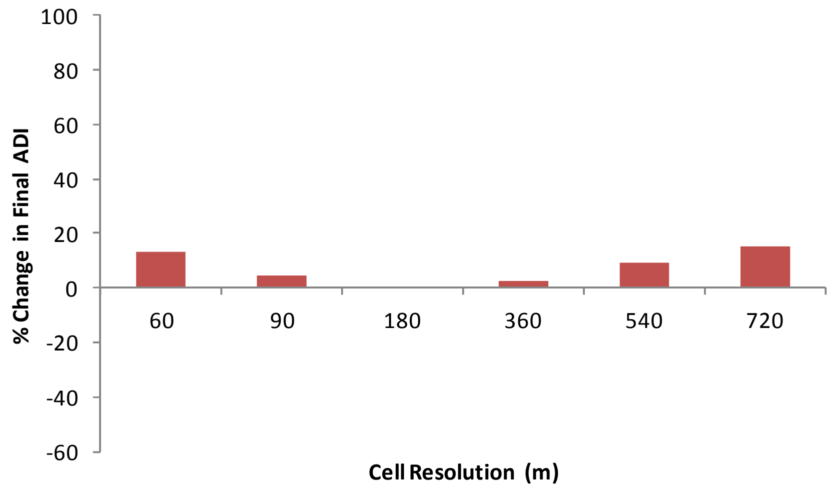

Figure 3: Wangary complex change in predictive performance due to variation in analysis cell size

The Kilmore case-study showed similar patterns to the Wangary complex; however the magnitude of the performance error was greatly increased. In both cases, as the resolution departed further from the standard 180m, performance declined. Interestingly, in both simulations, the 360m cell resolution showed only a moderate decline in predictive performance. As PHOENIX processing time is proportional to the inverse square of underlying grid cell resolution, this indicates that there is potential for processing performance gain with limited degradation of prediction quality. This effect has been flagged for follow-up.

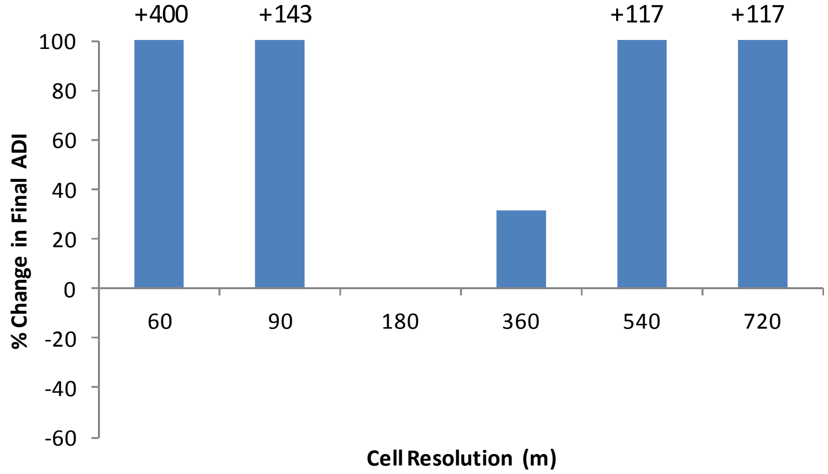

Figure 24: Kilmore complex change in predictive performance due to variation in analysis cell size

Summary of Case-study Results

The sensitivity of PHOENIX was evaluating using real case study data. Evaluations were carried out based on departure from the 'best estimate' outcome, using the Area Difference Estimate as a metric. There was substantial variation evident for when evaluating model sensitivity under real world conditions.

The simulation of the spread of fire through time is a complex process, as models must take into account changing fuel, topography, and weather. As demonstrated by the interaction between wind and slope illustrated in this report, the simulation of fire spread is highly sensitive to context and prediction uncertainty can result from a wide variety of sources. The different case-studies evaluated here exhibited sometimes contrasting responses to systematic variation of inputs. Due to the wide range of properties that can potentially affect outcomes, robust real-world sensitivity testing requires the evaluation of a large number of independent fires.

In general, the best simulation results were obtained when using the best available information available. As the deviation of the input increased, the predictive performance of PHOENIX decreased. However, the input data used for the case studies, particularly progression lines, is based on post-fire reconstruction. This information consists of the 'best available estimate' and cannot be considered a perfect 'truth'. Consequently, without further replication, the relative contributions of input quality and model function to prediction error cannot be apportioned.