Australian Government
Bureau of Meteorology

bushfire&natural HAZARDSCRC

FINAL PROJECT REPORT
bnhcrc.com.au

IMPROVED PREDICTIONS OF SEVERE WEATHER TO REDUCE COMMUNITY IMPACT

Final project report

Jeffrey D. Kepert, Kevin J. Tory, Eng Ching, Robert J.B. Fawcett, Serena Schroeter, William Thurston, David Wilke, and Dragana Zovko-Rajak
Bureau of Meteorology and Bushfire and Natural Hazards CRC
Corresponding author: Jeff.Kepert@bom.gov.au

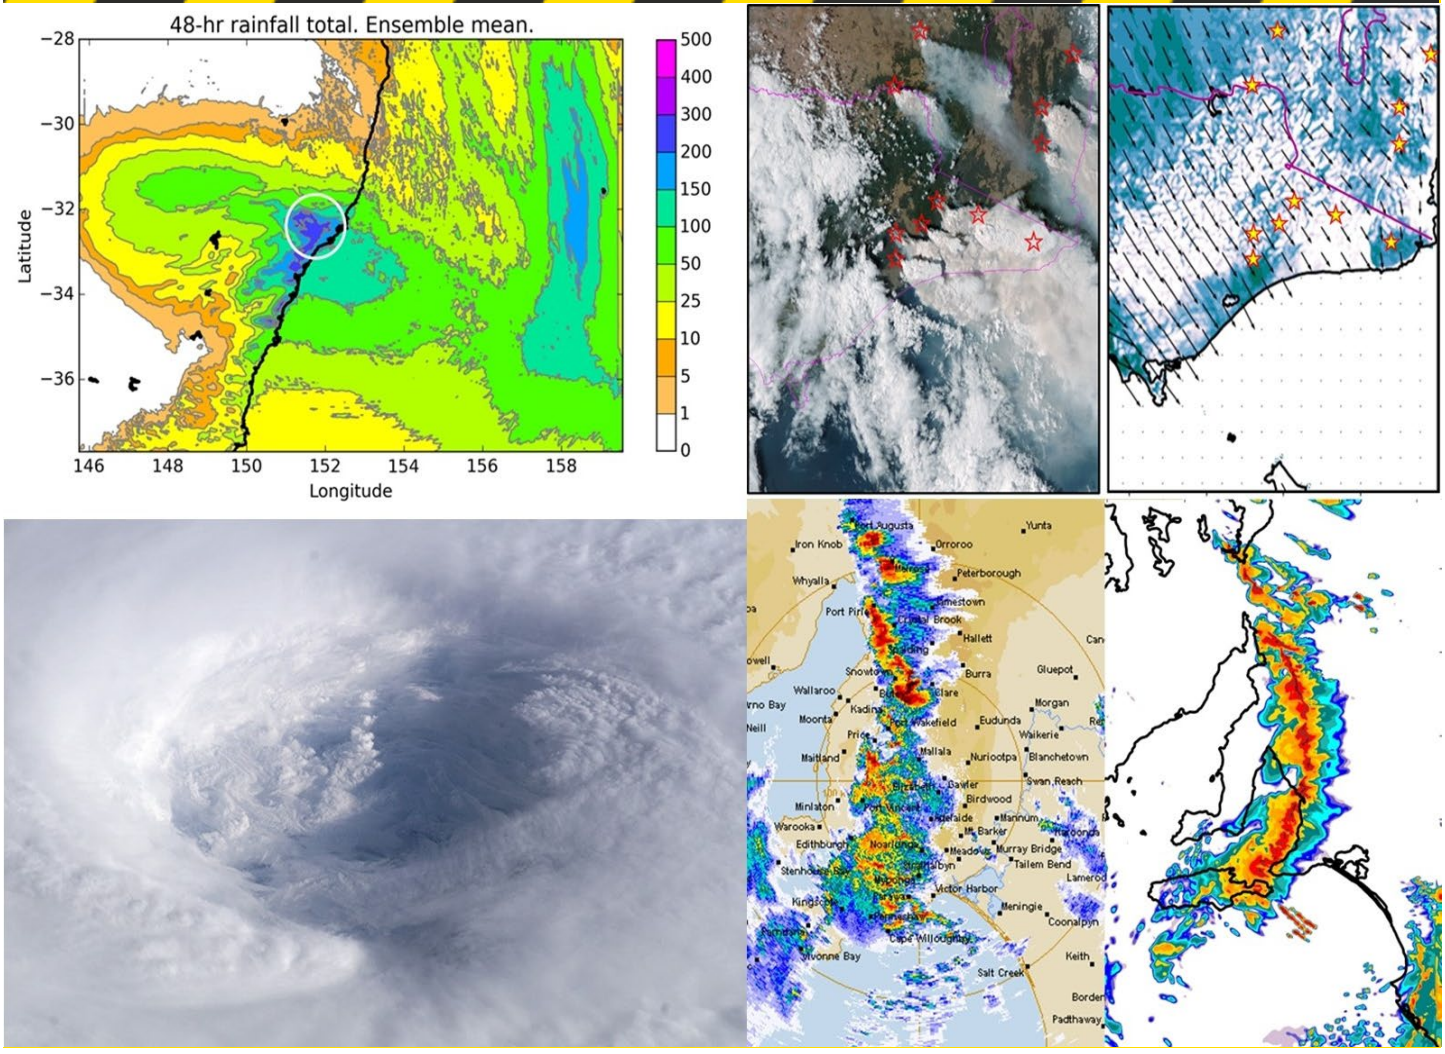

Version	Release history	Date
1.0	Initial release of document	15/02/2022

Publisher:
Bushfire and Natural Hazards CRC

February 2022

Citation: Kepert J, Tory K, Ching E, Fawcett R, Schroeter S, Thurston W, Wilke D, & Zovko-Rajak D (2022) Improved predictions of severe weather to reduce community impact – final project report, Bushfire and Natural Hazards CRC, Melbourne.

Cover: Four aspects of our research. The top left image is the simulated rainfall in the east coast low of April 2015, with the white circle indicating Dungog, where there was significant flooding. The top right pair of images shows observed pyroCb, together with diagnosed favourable areas. The lower left shows the eye of a tropical cyclone. The lower right pair are the observed and simulated rainfall rate in the SA tornado outbreak.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

We thank the UK Met Office for the provision of their Large-Eddy Model code and Adrian Hill and Ben Devenish for their assistance in setting up the model. We are also grateful to our colleagues at the Bureau of Meteorology's Earth System Modelling group for their assistance with the ACCESS system, and for their excellent work in implementing it. The analysis of the State Mine and Tathra fires was greatly assisted by advice and data from the NSW Rural Fire Service, particularly Simon Heemstra, Laurence McCoy, Mel O'Halloran and David Philp. We are also grateful to our colleagues in the Bureau's operational wing, particularly those based in Sydney and Adelaide, for helpful discussions on our case studies of severe weather events. Many international colleagues have helped us understand the latest overseas progress and incorporate that knowledge into our work, as well as encouraged us with the importance of our results.

We particularly acknowledge the helpful feedback from our end users on the practical implications of what we do, and for their thoughtful guidance on the direction of the project.

This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI), which is supported by the Australian Government.

EXECUTIVE SUMMARY

Extreme weather often occurs at relatively small scales. Accurate forecasts and understanding of such small-scale processes require high-resolution modelling. Forecasts are especially useful in severe weather events, since they play an essential role in allowing communities, industry and emergency services to prepare for and mitigate the impacts. Because forecasts are inherently uncertain in the severity, location and duration of an event, preparation needs to be more widespread than the eventual impact – but this over-preparation comes at a cost. Detailed prediction of the probabilities of severe impacts would avoid the risk of failing to alert areas with the chance of an impact, while minimising the cost of over-warning.

This project has studied the dynamics, predictability, and processes of severe weather, including fire weather, with the purpose of understanding phenomena with severe impact, improving forecasts of severe weather, and better depicting forecast uncertainty in these events. These goals help facilitate better risk management, improve user preparation, reduce adverse outcomes, and enable more effective mitigation.

The project has featured two main strands.

The first strand comprised case studies of severe weather events. For these, we have combined high-resolution numerical weather prediction (NWP) with a wide range of in situ and remotely sensed observations, to better understand both the high-impact event in question, and other events of that class. All except one of the studies used a version of the Bureau's operational NWP system, ACCESS. The events comprised two severe mid-latitude systems (an east coast low, and a severe thunderstorm and tornado outbreak), two severe fires, and a tropical cyclone, thereby covering the gamut of severe weather in Australia. Two of the studies featured ensemble simulation, with the east coast low case being the first time that ensemble ACCESS had been run at this resolution within the Bureau and foreshadowing the Bureau's new operational ensemble capability. Each case discovered important fine-scale features that contributed to the severity of the event, advancing our knowledge base and ability to respond.

The second strand studied two important phenomena associated with bushfire plumes: the formation of pyrocumulus clouds and ember transport. In each case, we began with idealised, high-resolution simulations of plumes using a large-eddy model. We used this technology, in which the model is run on a 50-m grid, to capture the most energetic size range of the turbulent eddies that are prominent in fire plumes. We found that the mean travel distance of firebrands depends mainly on wind speed and fire intensity, but the spread in the landing positions shifts from being substantially crosswind at light winds, to dominantly along-wind at high winds. This spread is greatly increased by the turbulence in the plume, and the maximum spotting distance can be more than double the mean for this reason.

We also used our plume modelling to study pyrocumulus clouds and analysed the processes that lead to pyrocumulus, with special attention on the relative importance of moisture from two sources, the atmosphere and combustion, and showed that the latter is negligible except in very dry environments. This

somewhat controversial result has been confirmed by a conceptual study of the thermodynamics of pyrocumulus formation.

These initial studies laid a firm theoretical framework for our subsequent development of tools to provide predictions in a form, and at a speed, that is suitable for operational use.

For moist pyroconvection, we developed a paradigm to combine the necessary meteorological and fire information, the pyrocumulonimbus firepower threshold, or PFT. The PFT is defined as the power output from a fire at which pyroCb will begin to form and depends solely on meteorological parameters. Fires hotter than the PFT will initiate pyroCb, according to this paradigm, and cooler fires do not. We also developed a way of computing an approximate PFT from either NWP or observed data and tested this in real time during the extraordinarily severe 2019-20 fire season. The results of that trial were outstandingly successful, with nearly all the 30-odd events being captured, and the forecast guidance for non-events was also reported to have been extremely valuable. The Bushfire and Natural Hazards CRC have provided us with utilisation funding to take this tool another step closer to operational use. The tool has also attracted significant international interest, particularly given the recent spate of events in the USA. We have begun a utilisation project to take the PFT work further, detailed in this report.

For ember transport, we developed a simple model to predict plume-dominated transport, incorporating well-known plume modelling concepts, a model of plume turbulence, and a probabilistic model of ember transport within and beneath the plume. The results agree well with our large eddy-based simulations. This model was coupled to CSIRO's Spark fire-spread model and tested on the Kilmore East fire of Black Saturday with excellent results. The combination with Spark will make this tool suitable for future operational use, and the ember transport model could also, of course, be coupled to other fire spread simulators, including the coupled fire-atmosphere model ACCESS-Fire. Working with CSIRO and AFAC, we are putting the necessary things in place to be able to convert the prototype implementation of the ember transport parameterisation in Spark, into an operationally robust and supported system. This will both facilitate further research, and provide the framework to allow users to learn, understand and apply the system.

The value of case studies has been reinforced by our being asked to undertake five case studies of severe fires from the 2019-20 summer. This work will be done jointly with our colleagues from the CRC *Coupled fire-atmosphere modelling project*, who further developed the ACCESS-Fire model. We expect that this work will provide an impetus to further develop and streamline that capability.

We are also very excited about the potential of the PFT and ember transport tools. While both are relatively early in their development, with the PFT being the more mature, each is presently producing very encouraging results. We are also very conscious that each is raising further scientific questions as they address others, and that both are at the beginning of a steep development and refinement curve that we expect will further improve their accuracy and utility. We are confident that further scientific investment will yield substantial dividends.

END-USER PROJECT IMPACT STATEMENT

Paul Fox-Hughes, *Bureau of Meteorology, Tasmania*

Lachie McCaw, *Department of Biodiversity, Conservation and Attractions, Western Australia*

John Bally, *Australasian Fire and Emergency Services Authorities Council, Victoria*

INTRODUCTION

The Severe Weather forecasting project has made an important contribution to understanding a variety of weather phenomena that can result in significant impacts to the community and the infrastructure on which communities depend. Severe weather events can also result in major impacts to the environment with consequences that last for decades, and in some cases even longer. The project team have brought an integrated focus to their work, combining physical principles, high resolution numerical weather prediction and observational data from case studies to provide insights into critical factors that drive severe weather events. The Bureau's operational numerical weather prediction system ACCESS has been central to this work, which is important because it provides a clear pathway for development and delivery of guidance products that can be used by forecasters and end-users in the emergency service and land management agencies.

The team has also ensured that awareness of their work has been disseminated widely within operational meteorologist and emergency management end-user groups through a number of presentations in a variety of different forums, through a very successful real-time trial of a new forecast technique, and through scientific and summary publications targeting not only the research community but also practitioners.

SEVERE MOIST PYROCONVECTION

The initial work on pyrocumulus resulted in the unexpected conclusion that pyrocumulus development is strongly dependent on moisture from the atmospheric environment, rather than from the underlying fire. This led to a closer examination by the research team of the thermodynamics of pyrocumulus formation, culminating in the development of a forecast tool.

The pyrocumulonimbus firepower threshold (PFT) forecast tool developed by the project has been very well received by both meteorologists and fire managers. It is focussed on calculating the firepower needed to trigger the formation of pyrocumulonimbus clouds (pyroCb), given the evolving atmospheric conditions. This is great for predicting pyroCb potential, independent of the presence or absence of fire. With recent developments in bushfire simulation, there is also the possibility of integrating the tool with a simulator, using the simulated firepower to signal the pyroCb risk at each time and section of a fire.

The PFT is a forecast product of great interest to fire management agencies. It enables a clear understanding of how the characteristics of the atmosphere and

a fire may combine to produce violent pyroconvection. From a relatively early stage in its development the PFT diagnostic has been made available for evaluation and operational testing and this has contributed to enthusiastic uptake by end-users, and to refinements in the diagnostic that have improved its performance and utility. During the 2019/20 bushfire season the PFT diagnostic was trialled extensively across a wide range of conditions and became a valuable and trusted guidance product for forecasters and fire behaviour analysts. We are pleased to see that it is now the subject of a utilisation project.

EMBER TRANSPORT

Forecasting the spread of fire is a powerful tool and as a result, bushfire simulators have become deeply integrated into the work of fire and land management agencies. There is strong recognition in the land and fire management sector of the need for a next-generation fire simulation capability. A core component of the next-generation bushfire simulator will be an ember transport and spotting sub-model that can easily be integrated into a practical simulation system. It should be able to be configured to work effectively in a full range of fuel types and a wide range of meteorological conditions and fire intensities.

The ember transport model developed in this project has a sound theoretical basis and the only input data it requires are fields that will be available in simulators currently under development. The parameterised version of the ember transport scheme captures the important features of the explicit ember transport modelling and implements those in a way that can be easily integrated and most importantly, can run very efficiently in practical operational systems. Early indications are that the parameterised ember transport scheme developed by this project does a great job reproducing spotting driven spread for recent major bushfires and looks to be a very promising method of predicting fire spread in an operational setting.

The ember transport parameterisation represents the culmination of a sustained effort over the full six years of the project. Initially, the team identified some key knowledge gaps in the problem, which they filled through detailed simulations using a large eddy model coupled to an explicit ember transport model. These results were scientifically valuable in their own right, but also provided the key understanding that was encapsulated in the computationally fast, final parameterisation. The team's experimentation with the parameterisation in the CSIRO fire spread simulator, Spark, provided insight into one of the most devastating fires in Australia's history, and was a great illustration of the potential value of this new tool. This work has received strong interest from within AFAC and the state fire agencies.

CASE STUDIES AND HIGH-RESOLUTION ATMOSPHERIC MODELLING

Case studies were a significant feature of the project and provided critical insights into the processes and impacts of several "high-end" events. The use of high-resolution simulation, particularly in ensembles, also helps inform the next generation of operational forecasting, for which ensembles will become a routine tool. It also allows emergency managers to gain a feel for the capability

of this approach, and to consider how they might incorporate its insights into their decision-making processes.

The case study of the Dungog east coast low (ECL) of April 2015 has provided insights into ECL dynamics, as well as the potential for operational use of forecast ensembles. Low-level onshore jet intensification, for example, occurred as a result of the development of a strong pressure gradient between the low-pressure system and upper-level ridging descending to near-surface levels. The 24-member 1.3 km resolution modelling pointed to the extent of the predictability of the ECL development, and to ways that operational meteorologists and emergency managers can employ ensembles to better forecast severe weather events. Such knowledge is important with the increasing use of ensembles in operational forecasting.

The emergency management sector has been aware for some time that ensemble prediction delivers a much richer and more nuanced prediction of developing hazards and impacts, but there are few examples of practical methods to deliver these in an operational setting. The sheer volume of material produced by ensemble prediction systems presents challenges to the effective delivery of this powerful new information source. This project has demonstrated the utility of ensemble matched mean forecasts, probability of exceeding thresholds and of the use of ensembles to shed light on the confidence of prediction as well as develop alternative scenario forecasts. The presentation of these ensemble results delivered by Dr Kepert as part of this project should be shared as a master class for operational meteorologists delivering operational briefings to emergency managers using ensemble predictions.

The South Australian severe thunderstorm case study has helped understand the setting and triggers of that destructive event. The detail available from sub-kilometre resolution modelling on the structure of the storms is remarkable. The outcomes of this part of the project will inform future thunderstorm forecasting in all mid-latitude Australian regions. In particular, the project has demonstrated the value of calculating updraft helicity as a guide to the potential occurrence of tornadoes in severe thunderstorm environments, and provides a great illustration of the value of ensembles for understanding and quantifying uncertainty in prediction of these very high-end events.

The project has delivered an important advance in the understanding of the drivers of extreme fire behaviour through the analysis of the State Mine Fire event. The influence of topographically forced winds and connection to upper atmospheric processes in enhancing the fire danger on the day of rapid fire advance clearly demonstrates the potential influence of these factors in forecasting and managing prescribed burns or wildfires in mountainous areas. Both these features were unlike those seen in previous fires: the mountain wave activity was prominent during daylight hours while the band of dry air that impacted on the fire did not form by the usual dry slot mechanism.

The Tathra bushfire case study demonstrates the value of high resolution modelling, in teasing out the complex and varying interactions between weather, topography and fire in the lee of the Great Dividing Range in southern NSW. The complexity of the low-level flow before the wind change, combined with the very detailed structure during the passage of cool changes poses

considerable forecast difficulty, but likely had profound implications for the behaviour and impact of the fire. This study was cited by the New South Wales Rural Fire Service in their evidence to the Coronial Enquiry into the Tathra bushfire.

SUMMARY

In summary, the *Improved Predictions of Severe Weather to Reduce Community Impact* project has delivered outstanding value in enhanced understanding of severe weather events, and in developing better methods for their prediction. The team has provided some really useful insights into the weather phenomena under consideration, but also, by utilising modelling systems similar to those expected in the next generation of operations, provided an opportunity for forecasters and emergency managers to consider how future operational forecasting approaches might assist in their work.

The research has progressed our understanding of severe weather events and provided tools that address gaps in our understanding. The strong relationships between researchers and end-users, along with support from the CRC and the Bureau's focus on delivering value from research, will help to ensure that the work of the Severe Weather project is further utilised effectively to support decision making, and that such operational use will improve outcomes for the Australian community during future extreme weather events.

This project has advanced our collective knowledge of severe weather prediction across a wide range of severe weather hazards and phenomena. Just as importantly, the approaches taken by the project are well suited to operational application, and indeed, there have been or are in train a range of initiatives to deliver very tangible operational benefits arising directly out of this work. The research team should be congratulated for an outstanding scientific effort, coupled with great communication and an eye for translating science into practical benefits for the fire community and the public.

PRODUCT USER TESTIMONIALS

PYROCUMULONIMBUS PREDICTION

The pyrocumulonimbus firepower threshold (PFT) proved to be a very useful tool for identifying the threat of pyroCb during the most recent fire season. The real-time trial of the PFT was of great scientific value, highlighting strengths and weaknesses not just of the tool, but also of pyroCb forecast procedures more generally. Of equal value, was the strong connection made between researchers and end users, which has opened up communication channels that should see continued user engagement in future PFT development and usage.

End users were asked if they thought a PFT utilisation project would be valuable, including further refinement of the PFT. Their responses were overwhelmingly positive both in endorsing the utilisation and the value of the tool over the summer. The responses are included here:

Paul Fox-Hughes – Fire-weather forecaster, Bureau of Meteorology

From the perspective of a former operational fire weather forecaster, recent National Manager for Extreme Weather, and current AFDRS BoM Science lead, in my view the more refinement that can be carried out on the PFT, the better. PFT has already demonstrated very considerable value over the last fire season, and the extensions that you suggest will further increase its value. In the case of the entrainment modification, that will also be of substantial value to "conventional" thunderstorm forecasting.

Zach Porter – Fire-weather forecaster, Bureau of Meteorology

I am really supportive of and excited by your ideas to refine the PFT for real-time forecasting. By adjusting the PFT sensitivity to cold outbreaks, I think this will help to reduce the false alarm rate in post frontal scenarios. It's also very difficult for forecasters to account for dry-air entrainment when plotting trajectories on an Aerological Diagram, so this aspect will be a welcome addition to the PFT estimates. In addition, I believe your proposal for higher resolution plots and labelling will greatly assist with our operational briefings to Fire Behaviour Analysts at the RFS, by enabling clearer links between PFT features and fire locations of interest.

James Pescott – Fire-weather forecaster and trainer, Bureau of Meteorology

The PyroCb Firepower Threshold (PFT) is without question an important addition to the fire weather forecasters tool kit. Having used it operationally during the trial it would benefit from some improvements before full operational implementation. Minimising triggering during cold outbreaks and the incorporation of buoyancy losses due to entrainment of dry air are both logical improvements that should be made before the next summer fire season.

Andy Ackland – Fire-behaviour analyst, Department of Environment, Land, Water and Planning (DELWP)

I would endorse the project extension. We trialled your PFT and PFT-flag products on a few occasions in the SCC during this last fire season. I found them particularly

helpful and accurate on these days where we were a bit unsure about pyrocb potential. I seem to remember your products helped confirm both situations of high pyrocb potential but also situations where it was lower than people were thinking and we were able to translate this into pretty key messaging for senior staff in the SCC.

One particular day, the PFT flag was later on in the evening as the change came up over the divide. This one was particularly accurate as that was near when the Buffalo fire seemed to have gone pyrocu or even pyrocb over Mt Hotham – the one Andrew Tapper later confirmed.

Dean Sgarbossa – Fire-weather product developer, Extreme Weather Desk, Bureau of Meteorology

As we have discussed previously, I would highly recommend the extension of the PFT Utilisation project to incorporate your great work within operations. The PFT flag and its novel but physically consistent approach introduces a powerful means in which to more efficiently identify areas of potential pyroconvection that is important to communicate to fire fighting authorities to support not only fire behaviour analysis and strategic fire fighting planning activities, but also to protect life and firefighting assets.

During the 2019-2020 season, we, the EWD [Extreme Weather Desk], along with many other operational meteorologists within the State and Territory offices, provided numerous briefing activities with up-to-date information on atmospheric ingredients impacting the fire weather situation across eastern Australia. This included unprecedented detail at the request of government emergency response agencies about pyroconvection which as we can all attest to, significantly impacted the behaviour of these bushfires. Consequently, the availability of the PFT to operations to further assist Bureau operational meteorologists in communicating the risk of pyroconvection would be an asset and has the potential of streamlining the forecast process involved in diagnosing, forecasting and warning for the risk of pyroconvection, particularly if the PFT can be incorporated within ensemble prediction systems such as APS3 ACCESS-G ensemble.

Diana Eadie – Fire-weather forecaster, Extreme Weather Desk, Bureau of Meteorology

I wholeheartedly endorse this extension – any fine tuning that saves forecasters having to manually filter out unrealistic outputs would be of great value in particular on high pressure days (e.g. change days such as the example given), when time is already such a limited asset - with the caveat of Dean's point, over is better than underreporting in these instances. The second proposal to create a more realistic plume-rise procedure would also be of great benefit and would likely reap benefits beyond fire weather forecasting, feeding into thunderstorm forecasting more broadly.

David Philp – Fire-behaviour analyst, NSW Regional Fire Service

I did find the PFT very useful over the past fire season and have been meaning to put together some feedback for you but have got myself bound up in making a submission to Royal Commission in part seeking more support for the Bureau and

trying to get my head around what happened on New Years Eve on the Far South Coast.

Lachie McCaw – WA Department of Biodiversity, Conservation and Attractions

I strongly endorse this proposal based on our experience of evaluating the PFT as guidance product over the 2019/20 southern bushfire season in Western Australia.

Our evaluation of the PFT spanned fires in a variety of environments and fuel types (eucalypt forest, mallee-heath, semi-arid woodland and heath) and included weather conditions associated with bushfires and prescribed burns. I consider that prescribed burns are important in this context for the reason that managers benefit from clear guidance as to conditions that may result in a prescribed fire interacting with the atmosphere to generate blow up fire behaviour, thereby compromising the objectives of the burn.

Based on evaluations conducted in Western Australia to date the PFT provides useful guidance in a form that is readily interpreted by fire managers and is consistent with the ACCESS modelling framework used to generate routine and incident weather forecasts.

David Grant – Fire-weather forecaster, Bureau of Meteorology

This sounds like a body of work that would be beneficial for my role that engages with those that support planned burn and wildfire operations at Queensland Fire and Emergency Services.

EMBER TRANSPORT

The ember transport parameterisation has gained significant interest within the fire community, particularly within AFAC.

John Bally – Fire Prediction Business Manager, AFAC

Forecasting the spread of fire is a powerful tool and as a result, bushfire simulators have become deeply integrated into the work of fire and land management agencies. There is strong recognition in the land and fire management sector of the need for a next-generation fire simulation capability. A core component of the next-generation bushfire simulator will be an ember transport and spotting sub-model that can easily be integrated into a practical simulation system. It should be able to be configured to work effectively in a full range of fuel types and a wide range of meteorological conditions and fire intensities.

The ember transport model developed in this project has a sound theoretical basis and requires as input fields that will be available in simulators currently being developed. The parameterised version of the ember transport scheme captures the important features of the model and implements those in a way that can be easily integrated and most importantly, can run very efficiently in practical operational systems. Early indications are that the parameterised ember transport scheme developed by this project does a great job reproducing spotting driven spread for recent major bushfires and looks to be a very promising method of predicting fire spread in an operational setting.

INTRODUCTION

This project aims to study the dynamics, predictability and processes of severe weather, including fire weather. We seek also to improve forecasts of severe weather, and to better depict forecast uncertainty in these events, thereby facilitating better risk management and more cost-effective mitigation. Our studies span a range of time and space scales and require a range of different methods.

Extreme weather often occurs at relatively small scales – here, the devil really is in the detail. Even when the meteorology driving the event is not small scale, small-scale perturbations within the overall framework can have a significant effect. Accurate forecasts and understanding of such small-scale processes require high-resolution modelling. Developing and validating such modelling, and extending it to all hazards, is the first aim of this project.

Forecasts are never perfect, but they are nevertheless useful. Forecasts are especially useful in severe weather events, since they play an essential role in allowing communities, industry and emergency services to prepare for and mitigate the impacts. Forecasting therefore underpins the work of emergency services and related agencies, and makes the PPRR (Prevention, Preparedness, Response and Recovery) process more efficient and effective. Because forecasts are inherently uncertain in the severity, location and duration of an event, preparation needs to be more widespread than the eventual impact – but this over-preparation comes at a cost. Detailed prediction of the probabilities of severe impacts would avoid the risk of failing to alert areas with the chance of an impact, while minimising the cost of over-warning. Thus the second aim of this project was to provide pilot predictions of not just the most likely course of events, but also the level of uncertainty, by identifying plausible alternative scenarios and their likelihoods.

A key aim of this project is to develop scientific understanding and to assist with the "lessons learned" from severe events. Our research thus adds to the collective wisdom of fire fighters, emergency services personnel and weather forecasters, improving our ability to manage these events and reducing the risk of adverse outcomes in the future.

The principal numerical weather prediction (NWP) modelling system used in this project is ACCESS, the Australian Community Climate and Earth-System Simulator. ACCESS is used operationally within the Bureau of Meteorology and is based on the UK Met Office's NWP system. Several other overseas national weather services similarly use this system. ACCESS therefore benefits from a wide user base, and the discipline of operational use and continual verification. The Met Office system is presently the second-best performing operational NWP system in the world, with the Bureau's implementation coming in third. The only system to consistently outperform it for global prediction is a global-only model which cannot be run in the high-resolution limited-area mode necessary to simulate fine-scale meteorology.

For situations where extremely high resolution modelling with a grid spacing of tens of metres, and the capacity to explicitly resolve atmospheric turbulence is needed, we use the UK Met Office's Large-Eddy Model (LEM). This specialised

model is designed not for forecasting, but rather for understanding phenomena that are highly sensitive to turbulence, including boundary layers, fire plumes and convective clouds.

Our ember transport work confirmed that the mean travel distance of firebrands for a given fire intensity depends mainly on wind speed. However, the spread in the landing positions shifts from being substantially cross-wind at light winds, to dominantly along-wind at high winds. This spread is greatly increased by the turbulence in the plume, and the maximum spotting distance can be more than double the mean for this reason.

These sophisticated and computationally intensive ember transport calculations were subsequently used to inform the development of a physically realistic and computationally cheap parameterization of ember transport for use in fire models. This simple model adequately reproduces the LEM results, makes intriguing new predictions of factors that are necessary for extreme-range spotting, and has been coupled with the CSIRO Spark fire-spread model.

We also used our plume modelling to study pyrocumulus clouds. Intense fire plumes in suitably moist environments can lead to cloud development, with the possibility of strong downbursts. We have analysed the processes that lead to pyrocumulus formation, with particular attention to the relative importance of moisture from two sources, the atmosphere and combustion, and shown that the latter is negligible except in very dry environments. This result surprised us and remains controversial. We were prompted to better understand the thermodynamics – that is, the interaction of fire heat and moisture with that from the atmosphere – in pyrocumulus, resulting in a detailed study of the basic underpinning processes, which confirmed our large-eddy simulation results. We used the knowledge gained to develop a forecast tool for pyrocumulus formation. The lack of sensitivity to fire moisture has important practical implications – fire moisture is not well characterised and would otherwise present a barrier to a forecasting system.

The challenge in predicting pyrocumulonimbus (pyroCb) formation is that it depends on both the fire and the atmosphere, and information from both scientific disciplines needs to be combined. Based on the above research into pyroconvective clouds, we developed a paradigm, the PFT or pyrocumulonimbus firepower threshold, that framed the question as, "what fire properties will lead to pyroCb in any given atmospheric environment?". We proposed that the use of a fire power threshold was the answer. Fires hotter than the threshold would initiate pyroCb, but not fires that were colder. We also demonstrated how to use a simplified plume model to estimate an approximate PFT from meteorological data (either model or observed soundings). We developed techniques to present maps of these data to users, and after extensive testing on historical events, ran a real-time trial over the summer of 2019-20. The prototype system performed extremely well over that season and became widely used during the summer. Of the 30-odd pyroCb events, it gave warning of nearly all, and proved to be highly valuable in correctly identifying when pyroCb would not form despite the presence of extreme fires. Analysis of the exceptions led to further improvement of the tool, and it was awarded CRC utilisation funding to take it the next step closer to operations.

We conducted detailed case studies of several severe weather events, based heavily on high-resolution ACCESS simulation, at times in ensemble mode.

We completed a detailed analysis of the meteorology of the Blue Mountains fires of October 2013, focusing on the most severe weather conditions and damaging fire activity of that event, the State Mine fire on October 17. We found two meteorological factors that likely contributed to the fire conditions on that day, mountain waves and a dry slot of unusual origin. The occurrence of mountain waves was interesting, since this was the third severe fire case in which we have made this link, and the first that occurred during the day rather than at night. Although this link has been suggested before, actual examples have been very thin on the ground, so we prepared a Hazard Note on mountain wave influence on bushfires, discussed further below.

The Tathra fire study combined observations with very high-resolution modelling to understand the mechanisms behind the disaster. At several points in the fire progression, atmospheric phenomena were active which greatly increased the risk. These were (a) mountain waves causing strong surface winds during the ignition, (b) the onset of boundary layer rolls in conjunction with weaker mountain waves interacting to intensify the fire and promote ember transport as it jumped the Bega river on its approach into Tathra, and (c) a complex interaction between the boundary layer rolls in the north-westerly flow and the wind change propagating up the NSW south coast leading to erratic conditions around Tathra in the late afternoon.

We conducted an ensemble case study of the April 2015 east coast low, which produced severe flooding, winds, and coastal erosion on the north coast of NSW. This study was the first we have undertaken using ensemble modelling, and the results of this study, as well as helping to understand the dynamics of this event, will also help guide users to make better use of this data source in the future. Indeed, that need is here now, for the Bureau now has two operational ensemble model systems: a coarse-resolution global system, and fine-resolution limited area systems covering large parts of eastern Australia. The ensemble dataset has demonstrated many of the benefits of ensemble modelling which are becoming well known – for instance, the fact that the ensemble mean tends to be a more accurate forecast, and that it provides the ability to compute probabilities and to consider a range of possible scenarios in responding to an event. It also enables a more rigorous approach to understanding the underlying meteorological processes. Often, we tentatively conclude that two factors which occur together also have some causal relationship, but since we have only one simulation cannot be certain. Now, we can examine the full ensemble for this association and potentially increase our confidence.

Our initial non-ensemble simulation of the South Australian tornado outbreak correctly captured the time and location of the thunderstorm mesocyclone that spawned one of the observed tornadoes. Events at such small scales as tornadoes are not normally regarded as being predictable in a deterministic framework, even though the broad area of vulnerability is, so we extended that simulation to be a 6-member ensemble. These showed a fascinating range of intensity and structure of the event, in which all had at least some tornado risk,

but the level and extent of the risk varied, highlighting the importance of ensemble modelling in providing guidance as to the extent to which we can be confident about model output in each particular situation.

Many potential users of ensemble or high-resolution simulations are aware that the computational demands are high, such that they are not yet available in real time. However, our studies also highlight the difficulty of extracting data in a timely fashion and putting it into a form that can be communicated to users and improve outcomes. Often, the conditions that lead to adverse outcomes are small in scale and locating them within the voluminous output takes time and skill. The Tathra fire was a good example of this.

However, some of this modelling capability is available in operations now. The Bureau has recently begun running operational limited-area ensembles on a 2.2 km grid in parts of eastern Australia. The configuration is similar in many respects to that which we used for the Dungog East Coast Low simulation (which was the first high-resolution ensemble simulation run in the Bureau). We expect that further growth in our analysis capabilities will continue to improve the ability to deliver excellent real-time outcomes from this resource.

About half of all intense tropical cyclones experience an eyewall replacement cycle, in which a new eyewall forms concentrically around the existing one and then contracts, eventually replacing the original eyewall. These events cause an initial weakening, followed by a re-intensification of the storm, and are associated with a marked expansion of the wind field, and hence broadening of the damage swath and increase in storm surge. We investigated a high-resolution simulation of an eyewall replacement cycle, with a particular focus on the role of boundary-layer convergence and its interaction with convection. Changes in the internal structure of tropical cyclones cause difficulties with their mitigation, particularly approaching landfall, so improved knowledge of a major cause of such change will lead to better forecasts and mitigation.

BACKGROUND AND MODELLING SYSTEMS

The principal numerical weather prediction (NWP) modelling system used in this project is ACCESS, the Australian Community Climate and Earth-System Simulator. ACCESS is used operationally within the Bureau of Meteorology and is based on the UK Met Office's NWP system. Several other overseas national weather services similarly use this system. ACCESS therefore benefits from a wide user base, and the discipline of operational use and continual verification. Our analyses of the East Coast Low of April 2015, which caused severe flooding and several deaths in the Hunter Valley and Dungog region, of the South Australian tornado outbreak, and of the Tathra fire event, have used ACCESS in very high resolution ensemble mode. The results of analysing these sets of ensemble simulations has shed new light on the dynamics and predictability of these systems, are the subject of published papers, and have been presented in several venues including as an invited speech at the Australian Academy of Sciences.

For situations where extremely high-resolution modelling with a grid spacing of tens of metres, and the capacity to explicitly resolve atmospheric turbulence is needed, we use the UK Met Office's Large-Eddy Model (LEM). This specialised model is designed not for forecasting, but rather for understanding phenomena that are highly sensitive to turbulence, including boundary layers, fire plumes and convective clouds. We initially used this model for ember transport, finding that the spread in landing position is greatly increased by the turbulence in the plume, and the maximum spotting distance can be more than double the mean for this reason. We then used these sophisticated and computationally intensive ember transport calculations to inform the development of a physically realistic and computationally cheap parameterisation of ember transport for use in fire models. This parameterisation requires negligible computer time to run and predicts the landing distribution of embers from a plume, providing a physically realistic prediction of the potential for spotfire development.

The final modelling system we use is CSIRO's Spark simulator for predicting fire spread. Spark is described in the user documentation as "a toolkit for simulating the spread of wildfires over terrain". The core of Spark is the solver, which is a computational engine for propagating a fire front. The rate at which the fire front spreads depends on factors including topography, fuel, and meteorology, is described by empirical equations (or fire spread models) such as those reviewed by (Cruz et al., 2015) and can be specified by the user. Spark reads necessary data (topography, fuels, and meteorology), interpolates them in space and time where necessary, and makes them available to the fire spread equations through the solver. Output data, including fire perimeters, are written to either shape files or as a grid. The complete configuration for a given simulation, including the above data, details of the ignition(s), and simulation start time and duration, is stored in an extended mark-up language (xml) file. This xml file is usually prepared using the Spark graphical user interface, which also displays the output data.

RESEARCH APPROACH

This chapter is by far the largest in the final report, since it presents the main scientific findings across the entire length of the project. This substantial volume of material is divided into two main topic areas. Hence, work is presented not in the order it was done but rather is organised to present not just one coherent story, but several running in parallel. Nevertheless, we have made individual sections reasonably self-contained to assist the reader interested in just one aspect.

The first main topic area covers our investigation of the science and applications of bushfire plume dynamics, which is further subdivided into two areas, ember transport and pyrocumulus formation. The discussion of each of these areas begins with a section where we explore the underpinning science, using a large-eddy model as our principal tool, followed by complementary sections where we use the knowledge developed in these computationally expensive calculations to develop much simpler analysis tools that are suitable for use in operational contexts.

The second main topic area is the application of high-resolution numerical weather prediction (NWP), including ensembles, to severe weather events. There are five sections in the category, all except the last in the form of case studies of severe events. The first two are ensemble studies of midlatitude severe weather events with strong dynamical forecasting: the east coast low of April 2015, and the South Australian tornado outbreak of September 2016. We follow these with two studies of severe fire events where the meteorology played a key and complex role: the State Mine Fire of October 2013 and the Tathra fire of March 2018. This topic area closes with a study of tropical cyclone secondary eyewall formation, using an idealised simulation of a North Atlantic hurricane.

Each section follows a similar format including a summary of conclusions on that aspect. Importantly, each section closes with a discussion of practical implications arising from that topic.

THE EFFECTS OF TURBULENT PLUME DYNAMICS ON LONG-RANGE SPOTTING

Introduction

Fire spotting is a hazardous phenomenon which leads to unpredictable fire behaviour and accelerated fire spread. Spot fires occur when embers are launched by bushfire plumes into the background wind, which then carries the embers a significant distance from the fire front. If the embers land in a suitable fuel bed and are still burning a spot fire is ignited. The magnitude of the problem is illustrated by Cruz et al. (2012), who provide evidence of long-range spotting in excess of 30 km during the Black Saturday bushfires of February 2009. Therefore, a better understanding of the processes that contribute to long-range spotting is essential for the prediction of fire spread. In this contribution we aim to assess the contribution of turbulent plume dynamics to the process of long-range spotting.

Methodology

We use a two-stage modelling approach to calculate the landing positions of potential firebrands launched by bushfire plumes. Firstly, we use the UK Met Office large-eddy model (LEM), described by Gray et al. (2001), to perform numerical simulations of idealised bushfire plumes. A number of plumes are simulated for background winds varying from 5 to 15 m s^{-1}. Secondly, the three-dimensional, time-varying velocity fields produced by the LEM are used to drive a simple Lagrangian particle-transport model. Potential firebrands are released near the base of the plume and then advected by the LEM velocity field with a constant fall velocity of 6 m s^{-1} applied, representative of jarrah and karri bark flakes (Ellis, 2010). The time-varying particle-transport calculations are then repeated using a steady plume velocity, calculated from the one-hour mean plume fields, in order to assess the contribution of the in-plume turbulence to the firebrand transport.

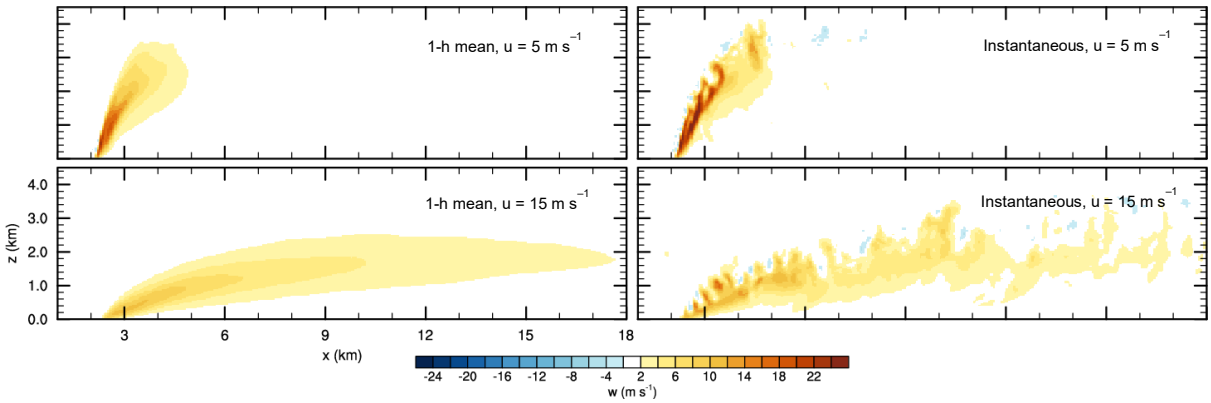

Figure 1. Vertical cross-sections of the mean (left) and instantaneous (right) vertical velocity, m s^{-1}, through the plume centre line, for background wind speeds of 5 (top) and 15 (bottom) m s^{-1}.

Results

Vertical cross sections of the instantaneous and 1-h mean updrafts for plumes in the weakest (5 m s^{-1}) and strongest (15 m s^{-1}) background winds are shown in Figure 75. The instantaneous plumes in strong wind have weaker updrafts, and are more bent over than the plumes in weak wind. The instantaneous strong-wind plume is turbulent over its whole height, whereas its weak-wind counterpart is only turbulent above a height of about 2 km. Plan views of the weak-wind plume, not shown here but seen in Thurston et al. (2014), reveal that the plume has two updraft cores that form a counter-rotating vortex pair. The 1-h mean plumes do not exhibit any of the turbulence that is visible in the instantaneous plume updrafts, and as a result have a weaker peak updraft.

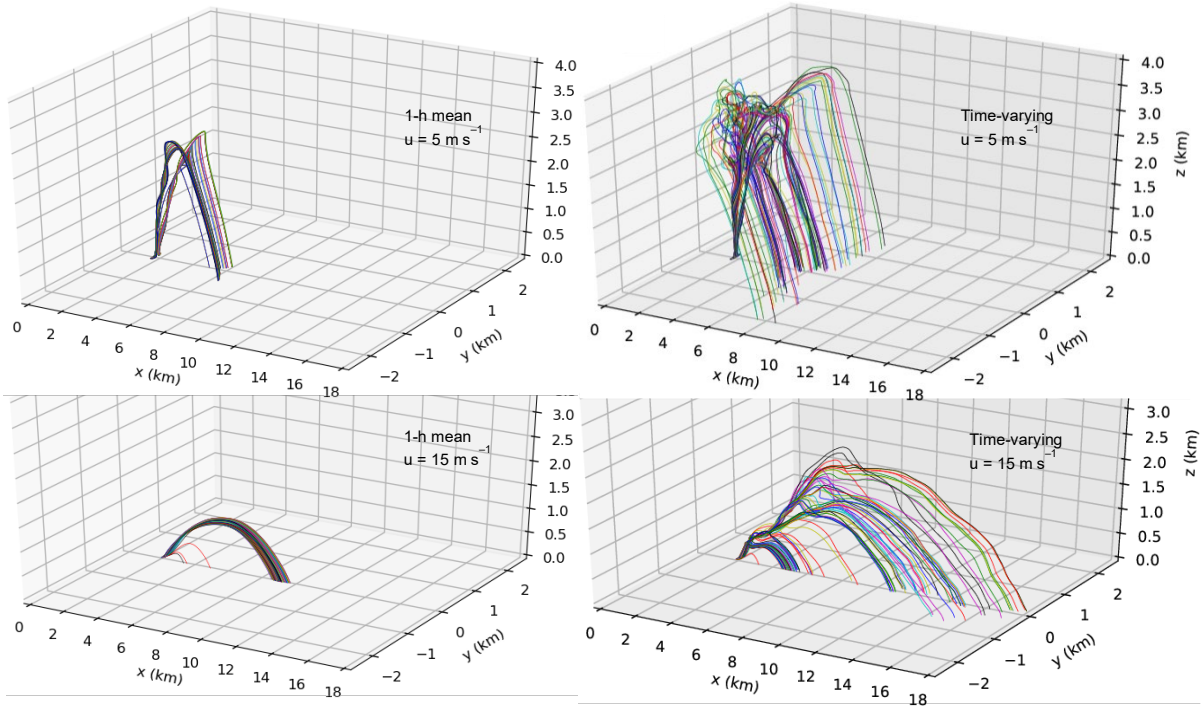

Figure 2. Trajectories of 100 firebrands lofted by the mean (left) and time-varying (right) plumes under background wind speeds of 5 (top) and 15 (bottom) m s–1.

The trajectories of 100 firebrands lofted by each of the plumes in Figure 1 are shown in Figure 2. Firebrands lofted by the time-varying weak-wind plume initially travel up the two branches of the counter-rotating vortex pair, and are then spread out further laterally as the reach the turbulent region of the plume above a height of 2 km. Firebrands lofted by the time-varying strong-wind plume do no exhibit any of this lateral spread, instead landing near the plume centre line. These firebrands appear to be lofted in clumps by the turbulent puffing of the plume, and hence tend to fall out in clusters. The trajectories of firebrands lofted by the 1-hr mean plumes highlight the importance of the in-plume turbulence. In the weak-wind case the firebrands still travel up the two branches of the counter-rotating vortex pair, but there is less lateral dispersion above 2 km. In the strong-wind case the effect of the in-plume turbulence is more pronounced, with most firebrands now having similar trajectories.

Figure 3 shows the two-dimensional landing distributions for 1.5 million firebrands launched by each of the plumes in Figure 1. The counter-rotating vortex pair and upper-level turbulence of the time-varying weak-wind plume lead to the firebrands landing in a V-shaped pattern with considerable lateral spread. The landing positions of firebrands lofted by the 1-h mean plume in weak winds still land in a V-shaped pattern, but there is less lateral spread due to the lack of in-plume turbulence. Firebrands lofted by the time-varying strong-wind plume travel on average about twice as far as their weak-wind plume counterparts, have more longitudinal spread and less lateral spread in their landing distribution. The landing positions of firebrands lofted by the 1-h mean plume in strong winds show much less spread and crucially the maximum spotting distance is reduced by half from about 16.7 km to 8.4 km.

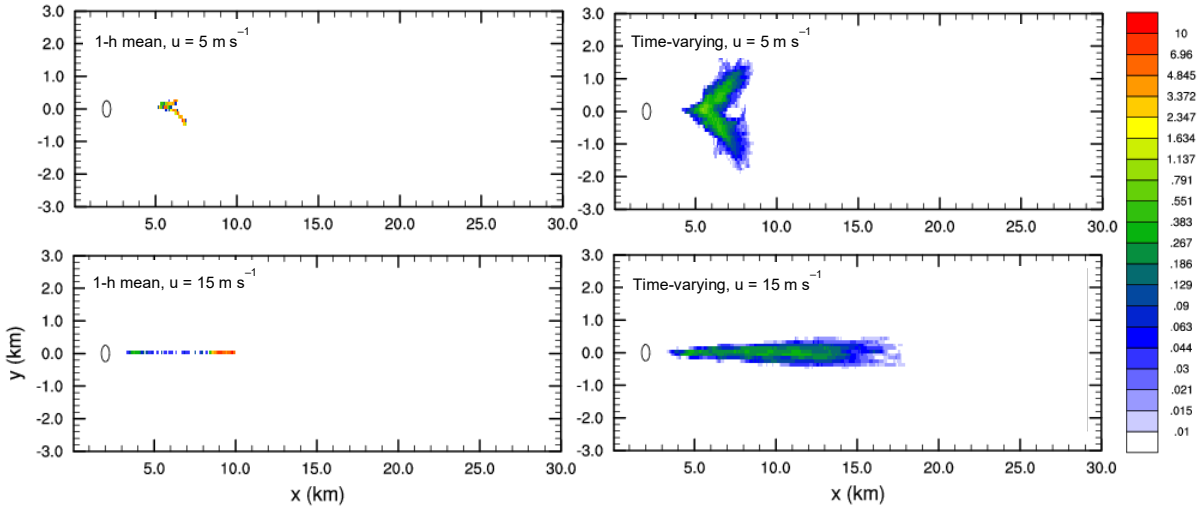

Figure 3. Spatial distributions of firebrand landing position (percent of particles launched per km2) for the mean (left) and time-varying (right) plumes under background wind speeds of 5 (top) and 15 (bottom) m s–1

A critical consideration in the potential for firebrands to start spot fires is whether they are still burning when they land. Therefore the flight times of the firebrands lofted by the time-varying weak-wind and strong-wind plumes are presented in Figure 4. Firebrands that are lofted by the weak-wind plume have a relatively long flight time, even if they do not travel a long distance. For example firebrands that are lofted by the weak-wind plume and subsequently travel only 0—2 km

are in the air for 7.5—12.5 minutes, whereas firebrands that are lofted by the strong-wind plume and travel only 0—2 km are in the air for 1.5—3.5 minutes. This caused by the plume dynamics seen in Figure 1: the weak-wind plume is more upright and has a stronger updraft, causing the firebrands to go almost straight up, reach a greater height and therefore be in the air for longer. This behaviour is confirmed by the trajectory plots of Figure 2. The firebrands that have travelled the furthest (16—18 km, in the strong-wind case) have a median flight time of 21.5 minutes and a 1st—99th percentile range of 19.3—23.4 minutes. This is of the order of the maximum burnout time of ribbon gum bark observed in the wind tunnel studies of Hall et al. (2015) and would suggest that firebrands that took these trajectories would still be capable of starting spot fires.

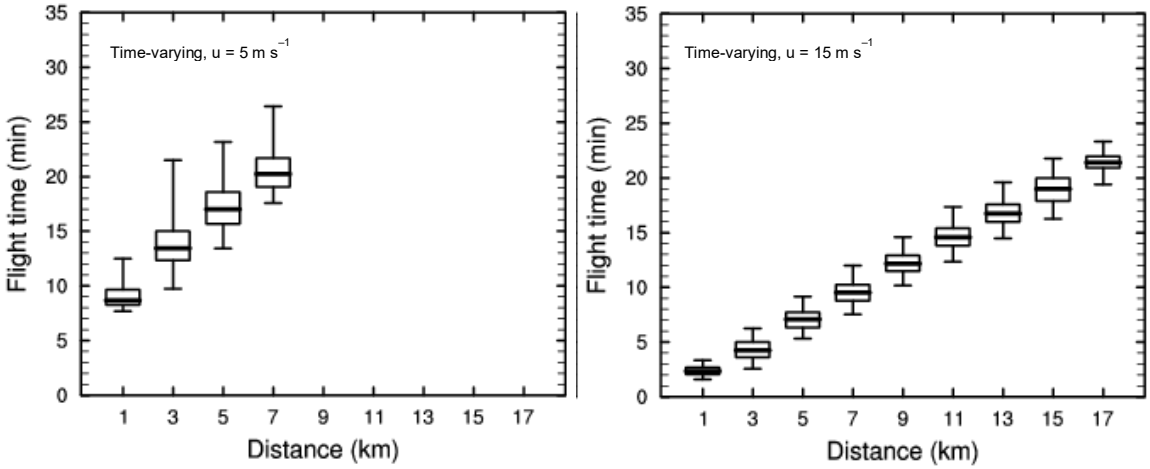

Figure 4. Box and whisker plots of flight times for firebrands lofted by the time-varying plumes under background wind speeds of 5 (left) and 15 (right) m s^{-1}. Flight times are binned according to the distance travelled by the firebrand, at 2 km intervals. The thick line shows the median flight time and the box spans the interquartile range. Whiskers represent the 1st and 99th percentile flight time.

Summary

We have combined large-eddy simulations of bushfire plumes with Lagrangian particle transport modelling to investigate how turbulent plume dynamics can affect long-range spotting. Plumes exhibited different dynamical and turbulent behaviour depending on the background wind conditions and this consequently leads to differences in firebrand transport. Plumes in weak winds contain a counter-rotating vortex pair which leads to large lateral spread in firebrand landing position. Plumes in strong winds are more turbulent and bent over, leading to more longitudinal spread in firebrand landing position and a greater maximum spotting distance. In-plume turbulence was shown to substantially reduce the lateral and longitudinal spread in firebrand landing position and in the case plumes in strong background winds, reduce the maximum spotting distance by a half. Systematic studies such as this will inform the development of better empirical spotting models.

Practical Implications

This work has highlighted the processes that affect long-range ember transport, and gave us the necessary underpinning knowledge to develop an operational

technique for calculating this important process. However, the present calculation is much too computationally expensive to be operationally feasible. Therefore, we developed a much simpler model, which nevertheless incorporated the key important physical processes which this work has revealed. This model, which is computationally cheap enough to be used in the field or as part of a fire spread simulator, is described in the next section.

A SIMPLE, PHYSICALLY BASED MODEL OF EMBER TRANSPORT WITHIN BUSHFIRE PLUMES

Introduction

Embers accelerate fire spread and make it less predictable. They contribute to fires breaking control lines and are a major factor in house loss (Blanchi and Leonard 2005). Therefore, methods are needed for predicting how far embers may travel from a fire and what the distribution of landing points and new ignitions will be.

Ember ignitions, or spotfires, are a particular problem in SE Australia, due to the prevalence of forest types that supply long-burning embers, easily detached from trees (Storey et al. 2020). Reliable observations exist of embers igniting new fires over 30 km from the front of the parent fire (Cruz et al. 2012 and references therein). Given that the median rate of spread for forest fires is of the order of 1 – 2 km/hr, this implies that such spotting may cause a fire to advance roughly a day's worth of spread in a few tens of minutes. However, while of particular concern in this region, other parts of the world and other forest types have also recorded spotfires in excess of 10 km ahead of the parent fire (refs).

Most attempts to model ember transport regard the phenomenon consists of three distinct processes: the generation of firebrands, their transport through the air, and the ignition of fuel at the landing position. Early studies (Tarifa 1965a,b, 1967; Albini 1979, 1981a,b, 1982, 1983a,b) presented models of a plume which lifted firebrands, and after the firebrands departed the plume they followed a slant trajectory to the ground, with the slope of the trajectory determined by their fall velocity and the environmental wind. Tarifa (1965a) introduced an important and accurate simplifying approximation, showing that the motion of the ember relative to the air was well approximated by its terminal fall velocity. The maximum spotting distance was an important prediction of these models, and was determined by the burnout time of the embers.

Most studies treat the ember trajectory as consisting of ascending and descending phases in this way. A few studies have proposed that fire-whirls are involved in the lofting process for spotting at longer ranges (Berlad and Lee 1968; Lee and Hellman 1969, 1970; Muraszew and Fedele 1977), but most regard the plume, or convection column, as the lifting mechanism. More recent studies have featured more complex, but still non-turbulent, models of the plume updraft (Wang 2011; Woycheese et al. 1997, 1999), often that of Baum and McCaffrey (1989). A more detailed history of transport modelling up to about this point can be found in Koo et al. (2010).

More recently, the availability of high-powered computing has enabled the application of large-eddy simulation (LES) to the problem. Large-eddy simulation involves numerical modelling of a fluid flow, where the larger turbulent eddies – those that contain most of the energy – are explicitly depicted, allowing the unsteady three-dimensional flow associated with the bushfire plume to be created and analysed. Bhutia et al. (2010) compared the results of applying the Baum and McCaffrey plume model to LES for ember transport, and noted that "in every case, a significant number of particles released from the coupled UU-LES dynamically–modelled fire plume interacting with a convectively–active

atmospheric boundary layer surpasses easily the spotting distance traversed by the particles released from a steady–state thermal plume superimposed on a constant mean wind domain." They commented also on the wavy trajectories of these embers. Koo et al. (2012) similarly used used the HIRGRAD/FIRETEC model with a horizontal grid spacing of 2m, on a 640 x 320 x 615 m domain, to simulate the transport and combustion of disc-shaped and cylindrical firebrands. Pereira et al. (2015) compared the Albini model to LEM-based simulations on a similarly small domain.

These two studies explored relatively short-range spotting, but Thurston et al. (2017) applied LES of bushfire-like plumes in a deep atmospheric boundary layer over a much larger domain to study ember transport at ranges of up to 20 km. They demonstrated that turbulence within the plume could more than double the transport distance of some embers, compared to identical transport calculations using the time-averaged (and therefore artificially non-turbulent) plume. Embers fell out of the plume episodically due to turbulence within the plume, leading to "substantial clumps" of embers impacting the ground. Flight times were up to 30 minutes, which they note would allow many real embers to burn out. However, they note also that their maximum spotting distance is still less than observed on Black Saturday (Cruz et al., 2012), probably because their fire intensity was substantially lower.

Thompson et al. (2019) presented similar calculations to Thurston et al. (2017), but with two refinements, both related to the terminal velocity assumption introduced by Tarifa (1965a). The fire intensity and boundary layer structure were essentially the same, and essentially confirm the earlier results on the role of turbulence in distributing the embers widely. They further showed that, for plumes of substantial vertical extent, that accounting for the density-dependence of fall velocity tended to reduce transport distance, presumably because with a faster fall velocity at higher altitude, embers are more likely to depart a plume earlier. Secondly, they showed some sensitivity to ember inertia; that is, the assumption that the relative motion of the ember in the air is just the fall velocity, further reduces the density of ember landing at the extreme of the range. In particular, they state that "the density of embers travelling 15,000 m is approximately 14 times higher under the constant terminal-velocity assumption, and approximately 5 times higher under the variable terminal-velocity assumption". We note, however, that these large differences occur at the downwind tail, where small reductions in transport distance result in large changes in density due to the super-exponential decay at this tail. An alternative interpretation of their results is in terms of the distance at which a given ember density is obtained. For instance, in their Fig 5, the distance at which the landing density falls to 1% of its peak varies by between about 5 and 25% under the different assumptions they explore, with the largest difference due to the variation of fall velocity with air density rather than the neglect of firebrand inertia.

A key conclusion from these LES-based studies is that ember transport within fire plumes is fundamentally a stochastic process. The plume is intensely turbulent, so where embers land is impossible to predict with precision. Moreover, whether the ember burns for long enough, and whether the landing point has suitable fuels for an ignition, likewise contain substantial random elements. This stochasticity implies a degree of uncertainty in the expected fire propagation, complicating

risk management. However, problems of risk management are eased when forecasts of expected behaviour (or of the range of expected behaviour) are available. In the case of ember transport, only the LES-based models presently predict the distribution of landing position. However, these are computationally much too expensive for real-time use, and are likely to remain so for many years. Therefore, there is a strong need for a computationally affordable model of ember transport that incorporates the effects of turbulence. The principal aim of this paper is to present such a model.

One of key findings of Thurston et al. (2017) was that embers could be transported more than twice as far when turbulence is included in the plume data, compared to simulations using the time-mean plume. The reason for this difference is that, when turbulence is omitted, embers mainly fall from the plume at the point where its mean updraft declines below the ember fall velocity, typically around 6 m s^{-1}. The plume updraft reaches a maximum a short distance above the fire and thereafter decreases due to the entrainment of air that is both cooler and without vertical momentum, diluting both the buoyancy and the vertical momentum of the plume air. Assuming the plume is strong enough to support some embers, most are dropped near the point in this decelerating part of the plume where the updraft drops below the fall velocity.

In contrast, in a turbulent plume, the time-mean updraft has superimposed on it a series of evolving, spatially and temporally coherent eddies. The additional updrafts and downdrafts due to these eddies are an appreciable fraction of the plume's time-mean updraft. Hence, some embers encounter a part of the plume that is stronger than average and will therefore be transported further. Others experience a weak part of the plume and travel a shorter distance.

The review of Koo et al. (2010) notes that many spotfire models treat the problem as three, sequential processes: ember generation, transport, and ignition. Each of these is complex, and dependent on multiple other factors. Our focus in this paper is on the second of these three processes, ember transport. Some models of maximum spotting distance based this constraint on ember burnout time, but the existence of spotting over 30 km from the fire suggests that, at times, burnout time may not be a strong limiting factor. Storey et al. (2020 IJWF) note that "Current spotting models are generally not suitable for more extreme fires because highly convective plumes or large aerodynamic firebrands are not accounted for", and we note that Thurston et al. (2017) and Thompson et al. (2019) each used a 19.2 GW fire.

The purpose of this section, then, is to describe a computationally cheap model, or parameterisation, which calculates the ember transport landing distribution knowing only large-scale properties of the atmosphere and fire but incorporates sufficiently accurate representations of plume dynamics to be accurate enough for operational use, and to apply it to a simplified simulation of the devastating Kilmore East fire in southeastern Australia on February 7, 2009.

Formulation

Parameterisations are simplified models of a particular process. Often, they are intended as components of a more complete model. For instance, in numerical weather prediction, processes that operate at scales below the model

resolution, such as turbulent mixing and moist convection, are typically parameterised since explicitly resolving them would be prohibitively expensive computationally. Being simplified models, some sacrifice of accuracy is inevitable – indeed, much of the art of parameterisation development is that of achieving a suitable compromise between computational cost and accuracy. Parameterisations can be viewed as theories and tend towards the lofty goals expressed by Einstein (1933): "It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience." Parameterisation developers may knowingly accept some inaccuracies where the consequences are modest, as indicated by the word "adequate" in this quotation.

Based on the discussion in the introduction, we choose to incorporate four main components into the parameterisation: a bulk plume model, a plume turbulence model, an ember transport model within the plume, and an ember transport model below the plume.

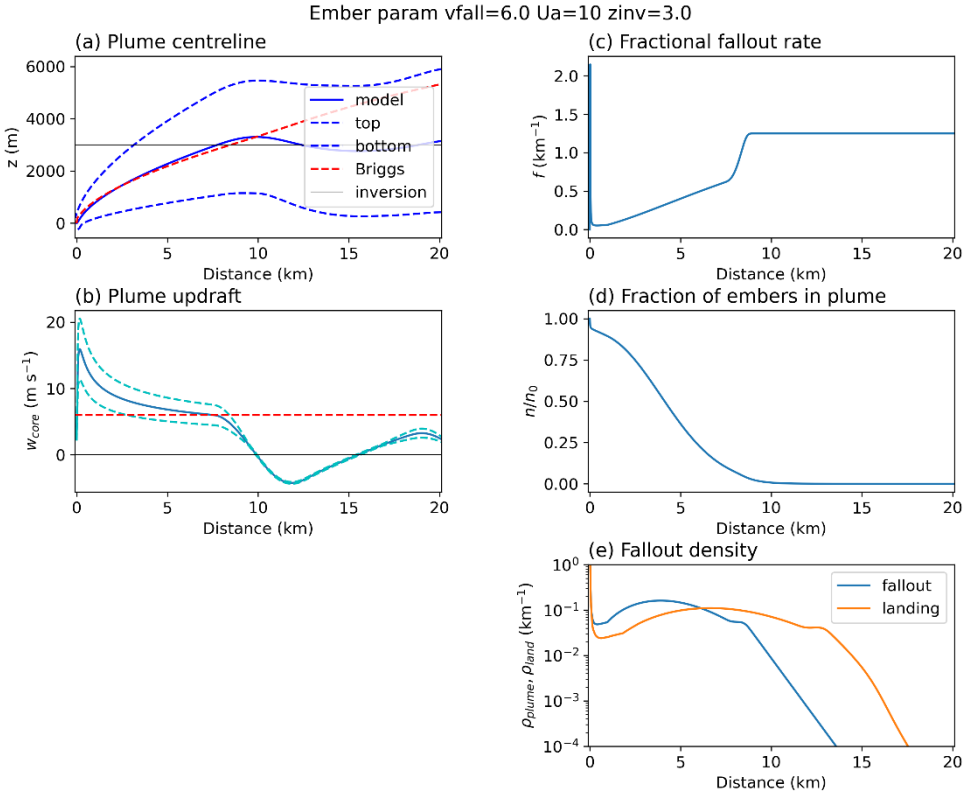

Figure 5: Components of the ember transport parameterisation. The left column shows the output of the bulk plume model and the right column the ember data. (a) shows the plume shape. Blue solid is the centreline, and the dashed blue mark the top and bottom of the plume. The red dashed line is the Briggs solution, which closely agrees with the numerical solution until they reach the inversion level (thin black). (b) shows the updraft strength (blue) plus and minus its standard deviation (cyan dashed) plotted against downwind range. The ember fall velocity is red dashed. (c) is the fractional fallout rate, with (d) showing the normalised concentration of embers in the plume. (e) shows the density of falling embers immediately below the plume (blue) and on landing (orange).

The Bulk Plume Model

Bulk, or integral, models of plume rise date back to Morton et al. (1957) and have been used for decades in applications such as pollution dispersion. They are based on the volume-averaged equations of buoyant plume rise, and Weil (1998) and Briggs (1984) give useful introductions. With further simplifying assumptions, analytic solutions are possible. The Briggs model, which assumes wind speed constant with height and neutrally stable stratification, is a particularly well-known example of an analytic solution. It has recently been applied to plumes from bushfire-scale fires for the purposes of forecasting the formation of pyrocumulonimbus clouds (Tory et al. 2018, Tory and Kepert 2020) and used to estimate the surface heat flux from lidar observations of a wildfire plume (Lareau and Clements 2017). The bulk model of Baum and McCaffrey (1989) was used in several of the ember transport models reviewed in the introduction.

The bulk equations are derived from the fluid dynamical equations by averaging across cross-sections of the plume, normal to the plume axis (e.g., Weil 1998, Tory and Kepert 2020). There are four equations describing the conservation of heat, momentum (vertical and horizontal components) and mass, and two describing a kinematic condition that determines the shape of the plume. These six coupled differential equations are integrated from the fire upwards and along the plume's axis. If the wind is constant with height, the stability is neutral, and for a simple entrainment closure, the integration can be done analytically and the Briggs model is obtained. However, we use a numerical solution of the bulk equations here, because we found the additional assumptions in the Briggs solution to be unduly restrictive for this application. In particular, Thurston et al. (2017) and Thompson et al. (2019) both show plumes transporting embers through the full depth of a deep boundary layer, suggesting that atmospheric factors affecting plume rise, including wind shear and stability, need to be included.

The inputs to the solution are vertical profiles of wind speed and static stability, and the flux of heat from the fire into the plume. Only part of the total fire power H_{fire} is carried away from the fire by convection in the plume H_{plume}. The remainder is lost by radiation, heating the ground, and so on.

The key part of the solution for our purposes are the shape, updraft strength, and width of the plume, but the temperature and along-stream velocity perturbation are also obtained. The updraft so obtained is the average across the plume cross-section, but we require the core updraft, since this describes the ember-lofting potential of the plume more accurately. Therefore, we apply an assumed Gaussian velocity profile across the plume, consistent with the diagnosed width and mean updraft, to obtain the core updraft strength. Several authors (e.g. Cunningham et al 2005; Bhutia et al 2010) note that the Gaussian profile is a poor approximation to observed and simulated plumes; however, we use this assumption only to set the ratio of the time-mean, cross-plume mean updraft to the peak time-mean updraft, which we calibrate against the LES data.

The left panels of Figure 5 show output of the bulk plume model, including the plume shape (panel a) and updraft strength (panel b). Note the oscillatory

behaviour as the plume core penetrates the inversion, damped due to continuing entrainment of environmental air.

The Turbulence Model

We can regard the ember motion as the result of two influences: the motion of the air in which it is carried, and its downwards fall velocity. Thurston et al. (2017) approximated the motion of the ember to be equal to the motion of the air plus its fall velocity, while Thomson et al. (2019) showed that accounting for the inertia of the ember led to modest changes in the predictions. Here, we simplify further, by accounting for only the vertical component of the turbulent part of the motion, consistent with usual practice in parameterising atmospheric boundary layers and that whether embers stay in the plume or fall from it will depend, to first order, on the relative magnitude of the plume updraft and their fall velocity.

How we characterise turbulence

Two properties of the turbulent eddies are important. The first of these is the magnitude of the updraft fluctuations. The second is their spatial scale, for this determines how long an ember will be under the influence of a particular updraft or downdraft within the plume, and hence how far it will travel. We will express the first quantity by the standard deviation of the updrafts within the plume, and the second by the autocorrelation length scale of velocity fluctuations along the plume axis. That is, by a scale distance within which air parcels within the plume have statistically well-correlated motion. This treatment is the simplest consistent with the simulated ember trajectories of Thurston et al. (2017) and follows the standard Lagrangian stochastic technique of use a Lagrangian decorrelation timescale to give a memory to the random component of velocity.

We will first derive these statistics from a large body of turbulent plume data as simulated in a large-eddy model, and then develop statistical models for these parameters in terms of quantities which will be available from a bulk plume model.

The large-eddy model

The plume simulations analysed here are performed using the same model configuration as Thurston et al (2017). Specifically, we use version 2.4 of the UK Met Office Large-Eddy Model (LEM). The LEM is a high-resolution cloud-resolving model that solves the Boussinesq equations of atmospheric motion on a three-dimensional Cartesian grid. The subgrid stresses are parametrised using a stability-dependent version of the Smagorinsky scheme, described by Brown et al. (1994). Further details of the LEM are presented by Gray et al. (2001) and references therein.

All simulations presented here are performed on a model domain of 768 by 384 grid points in the horizontal, with a uniform horizontal grid length of 50 m, resulting in a domain size of 38.4 by 19.2 km. The domain has 256 grid points in the vertical, in a stretched grid with a vertical grid length of 10 m near the surface, increasing smoothly to 50m at a height of 3 km, and constant above this up to the model top at 10 km. The lateral boundary conditions of the model are doubly periodic, and a no-slip condition is applied at the lower boundary, with a specified

roughness length of 0.05 m. The domain size is large enough for the periodic lateral boundary conditions not to affect the simulated plumes. A free-slip condition is applied at the upper boundary of the model domain in conjunction with a Newtonian damping layer in the top 2 km of the domain, to prevent the reflection of vertically propagating gravity waves. Our simulations omit parametrisations for radiation, moisture or microphysics and neglect the Coriolis force, all reasonable assumptions given the scale of the plumes simulated here.

In all simulations, we first spin up a fully developed boundary layer. The "fire" is then ramped up over a 5-minute period to its final strength. The model is then run for 30 minutes to enable the plume to reach a statistically steady state, followed by a further 60 minutes of simulation of the mature plume, which are analysed.

As in Thurston et al (2017), we represent the fire by a circular region of radius 250 m with a prescribed heat flux. While they considered only a heat flux into the plume of Q_{plume} = 100 kW m^{-2}, or a total fire plume heat flux of H_{plume} = 19.6 GW, we consider heat fluxes of Q_{plume} = 1, 2, 5, 10, 20, 50, 100 and 200 kW m^{-2}, with total fire power into the plume H_{plume} ranging from 196 MW to 39.2 GW. The Q_{plume} = 100 kW m^{-2} simulations here are the same as the ones they used. In addition, they found that to obtain converged ember transport calculations, it was necessary to output winds from the LEM every 5 s, leading to a very large data volume. For our purposes, such high-frequency sampling is unnecessary, so we instead utilise 30-s output, except in the Q_{plume} = 100 kW m^{-2} cases.

The model output is on a staggered grid, in the sense that the velocity data points are displaced half a grid length from the pressure and temperature data. The output files are post-processed to interpolate all the variables to a common grid before further analysis is performed.

Diagnosis of plume axis and extraction of data

The plume axis is diagnosed from time-mean temperature and updraft data. Starting at the heat source, the strongest updraft and hottest point are found in each grid slice normal to the wind direction, giving two estimates of the plume axis, based on updraft and temperature. The updraft and temperature perturbation are normalised by their maximum values for that plume, and then a consensus plume shape calculated as the weighted mean of the updraft-based and temperature-based positions, with the normalised perturbations providing the weights. We used a weighted mean of both quantities because in some cases, the updraft criterion gave ambiguous results close to the heat source, where the updraft was relatively weak but the temperature perturbation strong, and because the temperature perturbation became too weak to reliably detect far downwind. Taking the absolute maximum updraft occasionally led to apparent jumps in plume height far downwind due to the plume generating gravity waves in the stable layer above 3 km, which were detected and removed. The resulting plume shape was then lightly smoothed using a spline filter.

The LEM output data were interpolated to the location of the time-mean plume axis, as found by this technique, for each simulation and stored for subsequent analysis.

Updraft perturbation strength

The strength of the perturbations is defined by their standard deviation, as a function of distance along the plume. Figure 2 shows this quantity plotted against along-plume distance and mean updraft strength. It is apparent that there is a near-linear relationship between the updraft strength and its variability. Some outlier points are apparent; these are in the accelerating part of the plume near the source region. As will be further discussed in the next section, there are limitations in the LEM's ability to resolve the turbulence in this part of the plume.

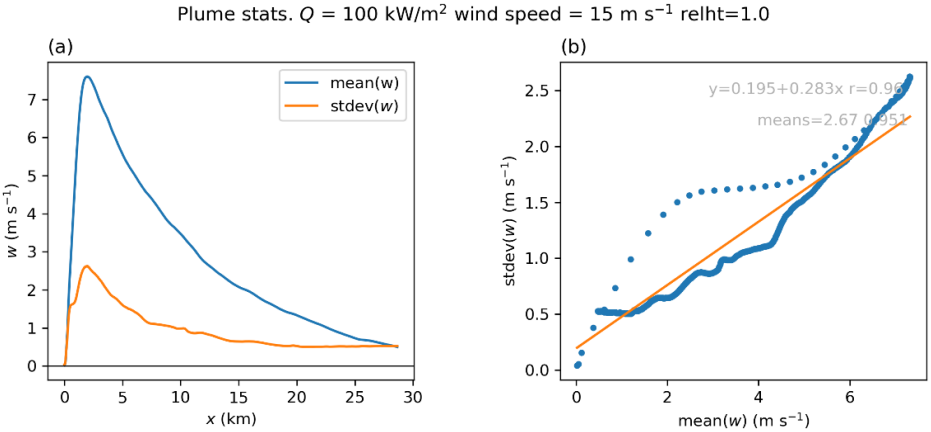

Figure 6: (a) Updraft strength (blue) and standard deviation (orange) along the plume axis for the same case as in Figure 1. (b) Scatter plot of mean against standard deviation (blue dots), together with the least-squares line of best fit (orange line). The blue dots include both the vertically accelerating (spaced dots) and decelerating (overlapping dots) parts of the rising plume.

Similar data were extracted for all the plume simulations and plotted in this manner. Nearly all of them displayed similar behaviour, in that a strong near-linear relationship between the updraft time-mean strength and its perturbations was apparent. One mechanism for the generation of turbulence is wind shear, and the shear between the plume and environment is dominated by the vertical wind component, so this correlation is physically reasonable. It is also consistent with the common practice for the parameterisation of entrainment rate in bulk plume models, which popularly depends linearly on the velocity difference between the plume and the environment (Webster and Thomson 2002).

However, while all cases exhibited an approximately linear relationship, the slope of the line showed a modest level of variation between cases, implying a dependence on other parameters.

Multiple linear regression was used to derive a turbulence intensity model that was applicable across the parameter space. Data deemed unreliable due to the limitations of the LEM, as discussed earlier and in the next section, were excluded. Candidate predictors were chosen based on the known physical behaviour of plumes and their availability from the bulk plume model. Although some evidence of modest nonlinear behaviour between predictors was apparent, products and powers of predictors were not included to avoid unphysical behaviour in extrapolations outside of the data region. All possible linear models using selections from this set of candidate predictors were fitted, and various statistics of the goodness of fit calculated. We chose to select the model with the best (lowest) Bayesian information criterion score, which included

all considered predictors. The in-sample root-mean-square error of the fitted values was 0.21 m s^{-1}.

The fitted equation for the standard deviation of the updraft in units of m/s was

$$\sigma_w = 2.15 - 0.254\ \log_{10} H_{plume} + 0.556\ U + 0.305\ \max(w_{core}, 1) + 0.00336\Delta\theta_{core}$$

where H_{plume} is the fire power into the plume in W, U is the wind speed in m/s, and w_{core} and $\Delta\theta_{core}$ are the time-mean updraft and potential temperature perturbations on the plume axis in units of m/s and K, respectively.

Turbulent length scale

The length scale is important because it describes how long an ember experiencing especially favourable conditions for transport – i.e., a stronger than average updraft – will continue to do so, and conversely for weaker updrafts or downdrafts.

Figure 7 shows the time-variation of plume updraft, along the plume axis, for the case with H_{plume} = 19.2 GW, U = 15 m s^{-1}. The updrafts and downdrafts in this figure occur as diagonally aligned streaks, with clearly much stronger coherence along these space-time diagonals than in the purely spatial direction (i.e. parallel to the ordinate). The slope of these diagonal streaks is close to 1/(15 m/s); that is, the reciprocal of the environmental wind speed.

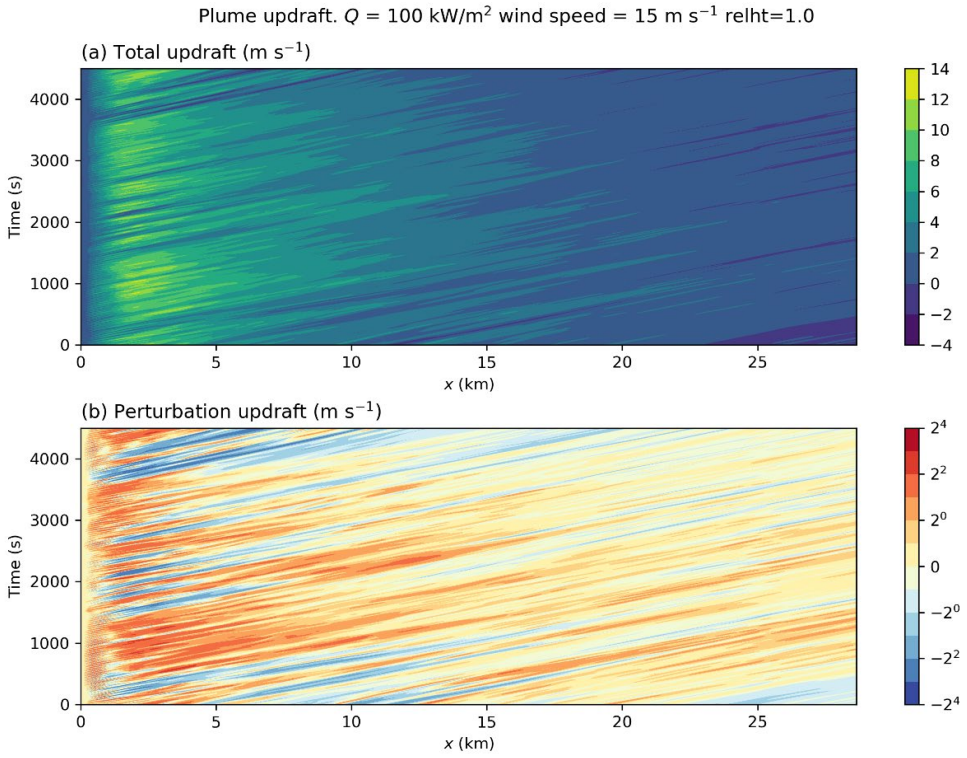

Figure 7: Contour plots of (a) total updraft and (b) updraft perturbation (using a double-ended log-like colour scale) for the same simulation as in Figure 5 and Figure 6. The diagonally slanting lines indicate the propagation of coherent regions of enhanced updraft along the plume and have slope approximately equal to the horizontal wind speed at the plume height.

The mean horizontal motion of embers is close to the environmental wind (Thurston et al. 2017), so embers will tend to follow trajectories nearly parallel to

these streaks. They will not be exactly parallel, because the embers are also affected by the horizontal velocity fluctuations, the horizontal velocity within the plume can differ from that in the undisturbed flow, and the boundary layer possesses vertical shear. However, none of these effects has a more than second-order effect on the ember trajectories.

At each x point in the data plotted in Figure 7, we interpolate data at five points (x-2u dt,t-2dt), (x-u dt,t-dt) … (x+2u dt,t+2dt) for each time point. We then compute the correlation matrix across all the time values at that spatial point, which describes the updraft fluctuations within the vicinity of that point. The length scale is defined by

$$L = \left(\frac{\partial^2 C}{\partial x^2} \right)^{-1/2}$$

where C is the autocorrelation (Gaspari and Cohn 1999). The derivative is evaluated by a fourth-order accurate finite difference method across the five-point stencil defined above.

Figure 8 shows the length scale so defined for the same case (i.e. Q_{plume} = 100 kW m^{-2}, u = 15 m s^{-1}). In this case, it is close to zero at the plume source, and increases rapidly to about 800 m, then fluctuates around this value in the remainder of the plume. The fluctuations represent statistical noise, to which the second derivative in the definition is quite sensitive. For comparison, the instantaneous correlation length scale is also shown, which has a similar structure but fluctuates around 200 m in the bulk of the plume. Examination of panel b in this figure, where the length scale is plotted against the plume mean updraft strength, reveals that in this case the linear growth part of the length scale curves corresponds approximately to where the updraft is accelerating to its maximum value, and the near-constant part to where the updraft is slowly weakening as the plume updraft decays downwind. However, this neat relationship is not exhibited in all cases.

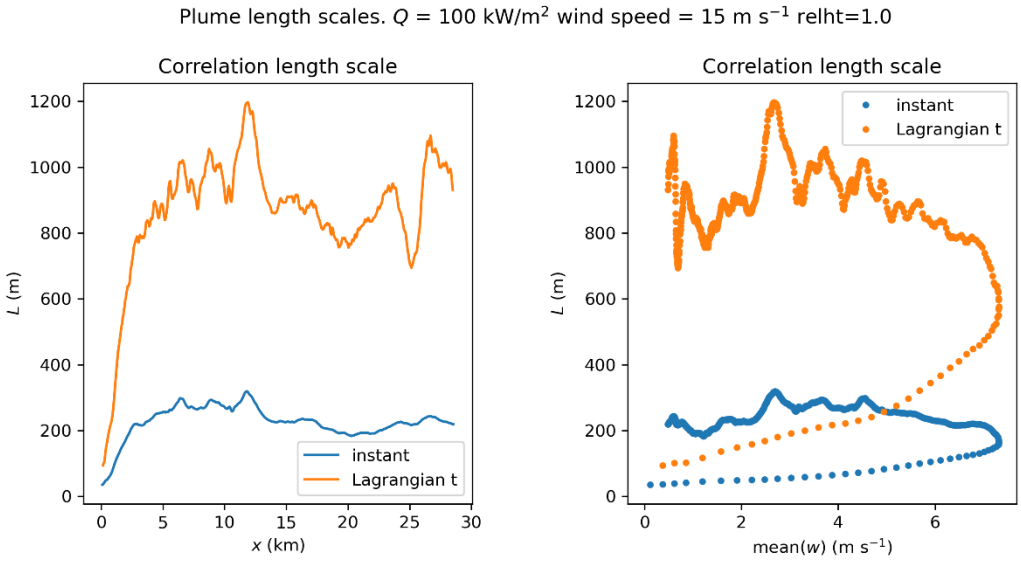

Figure 8: Plots of motion-following correlation length scale (orange) against (a) distance along the plume and (b) plume updraft. The blue curves show the instantaneous length scale (i.e. not motion-following).

Similar calculations were performed for all combinations of heat flux and wind speed and are plotted in Figure 9. While nearly all displayed the downwind near-constant length scale seen here, there was less consistent behaviour in the near-source, upwardly accelerating part of the plume. Some cases had near-constant *L* throughout, and some began with near-zero values and increased towards the near-constant *L*. A few cases depart from *L* = 800 m in the far downwind part of the plume; these we ignore since the updraft (mean and standard deviation) are too weak to suspend more than a very few embers, if that, in this portion of the plume.

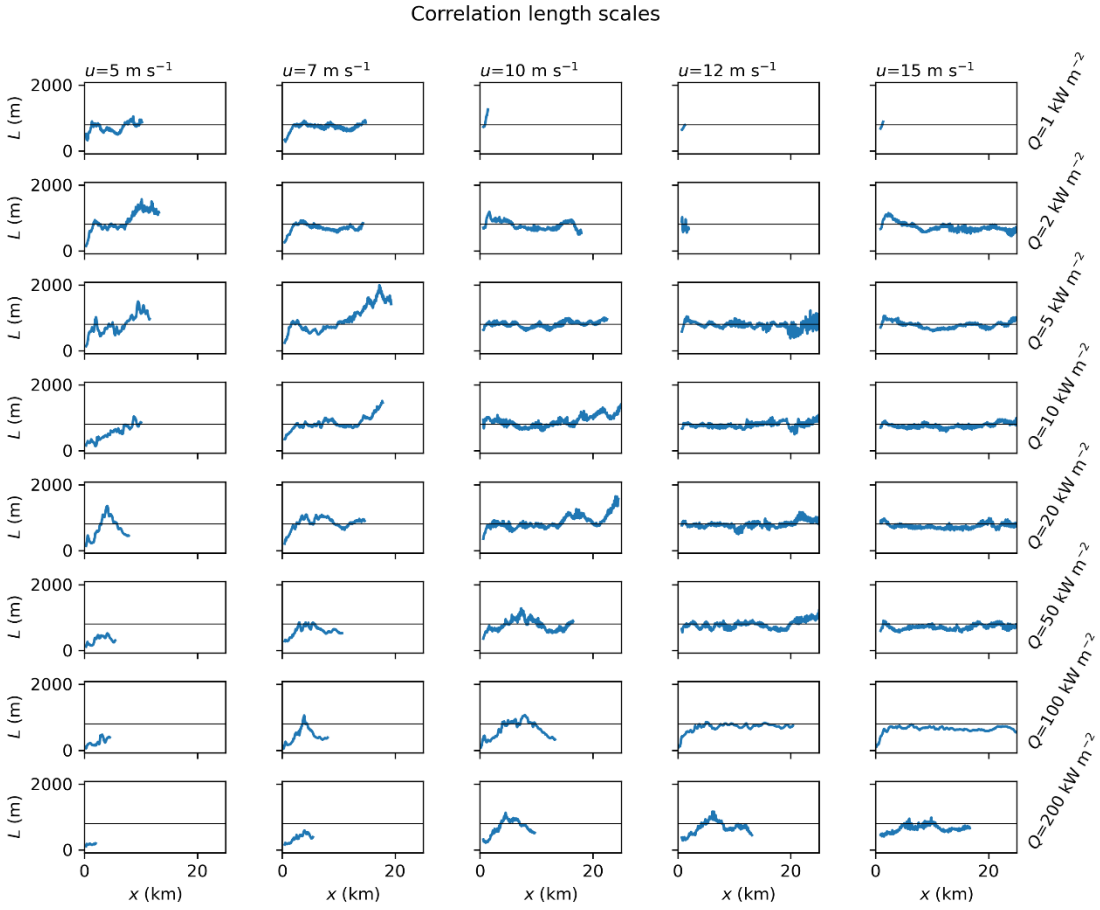

Figure 9: Plots of the correlation length scale against along-plume distance for all cases. The data are arranged by wind speed (columns, values indicated at the top in m s^{-1}) and heat flux (rows, values indicated to the right in kW m^{-2}). The black horizontal line in each panel is at L = 800 m.

The reasons for this divergent behaviour near the source are believed to relate to inherent properties of large-eddy models and the nature of turbulence. Large-eddy models aim to represent the turbulence above a certain scale and to parameterise its effects below that scale. Their formulation includes an explicit filter at this scale. How much of the turbulence is explicitly represented, compared to how much is parameterised, is flow and location dependent. For example, turbulent eddy size is constrained near a surface, so less of the turbulence is represented explicitly and more is parameterised near the Earth's surface. In addition, a few of the simulations exhibited regular oscillations of the near-surface plume behaviour, rather than more irregular fluctuations. These

may be due to insufficient resolution to fully represent the plume dynamics, possibly where the plume forms a near-surface neck.

The downwind distance at which the length scale first reached 700 m is shown in Table 1 – a threshold slightly lower than 800 m was chosen to reduce contamination by statistical noise. Roughly the half of the table above the diagonal in the upper right (and similarly the corresponding half of Figure 9) have only a small region of ramp-up (i.e., they reach the length scale in about a km or less). The lower left corner of the table and figure have some cases with high heat flux and low wind speed which never reach L=700 m, these correspond to near-vertical plumes. Examination of these simulations shows that the plume appears to rise too quickly for the model to be able to adequately depict the development of turbulence therein, so these are deemed to be unrealistic. Between these two groups is a band of cases which take several kilometres to develop the quasi-asymptotic length scale. It is unclear whether this behaviour is completely physical. In the prototype ember transport model, including a ramp-up region of a few kilometres drops more embers near the source, than when it is omitted. Accordingly, we take a ramp-up width of 3 km for plumes in the lower-left diagonal half of the table, and 1 km in the upper-right diagonal half.

Table 1: Downwind distance (m) where the length scale first reached 700 m, for the data in Figure 9. Blank squares correspond to cases where that length was never reached.

Q/u	**5**	**7.5**	**10**	**12.5**	**15**
1	1150	1800	600	1000	1000
2	1450	1750	600	750	1100
5	1350	1250	800	950	900
10	6950	2050	600	950	950
20	2700	1850	1200	750	900
50		2750	1600	1000	1150
100		3400	3900	3050	2500
200			3450	3600	4400

Turbulence Parameterisation Summary

We have analysed a large volume of data from large-eddy simulations of turbulent plumes. The turbulent vertical velocity fluctuations, relative to the time-mean plume, were characterised by their standard deviation and along-plume correlation length scale, calculated in a flow-following (quasi-Lagrangian), rather than stationary (Eulerian), sense. Statistical models of these two turbulence quantities were generated. For the standard deviation of the updraft, this was a multivariate linear regression model. For the correlation length, behaviour was reproduced satisfactorily by two piecewise linear models. In both cases, the

predictors are either bulk plume quantities, environmental quantities, or combinations thereof.

The Ember Transport Model

Ember transport within the plume

At each point along the plume axis, we have a probability distribution of the vertical velocity, assumed to be Gaussian, and therefore completely described by its mean and standard deviation. The mean is given by the bulk plume model, and the standard deviation by the linear regression model just described. The probability that an ember falls from a unit length of plume is taken to be

$$f = \frac{p(w + v_{fall} < 0)}{L}$$

The numerator of this formula represents the probability that the plume updraft is locally too weak to suspend an ember, while the denominator recognises that the turbulence is spatially correlated, and that the particular random sample from the updraft distribution that might apply at this point and time, also applies nearby. The correlation length scale is the appropriate metric of that tendency. We note that it is a scale, rather than an absolute length, so in principle it may become necessary to multiply it by a tuning coefficient of order 1, although our experience indicates that such tuning is unnecessary.

This equation contains an implicit assumption, discussed in the introduction, that the embers are falling at their terminal velocity. Thompson et al. (2019) have questioned this assumption when turbulence is present, but, as discussed earlier, the effect on the range at which a given ember density occurs is modest and we therefore regard it as acceptable at the present stage of development.

The fractional fall-out rate is plotted in Figure 5c. Note that it has a maximum value of 1/*L*, since the probability in the numerator is at most 1. A sudden increase in *f* is apparent near 8 km, where the plume encounters the inversion causing the mean updraft to drop dramatically.

The term *f* thus represents the average fractional rate at which embers fall out of the plume. If the density of embers in the plume is *n*, then the along-plume variation of *n* is

$$\frac{\partial n}{\partial x} = -fn$$

The solution to this equation is

$$n(x) = n_0 exp\left(-\int_0^x f\right)$$

where n_0 is the density of embers at the fire and x is the horizontal projection of distance along the plume.

The density of embers within the plume is plotted in Figure 5d. This quantity has units of reciprocal length; it is the number of embers per unit length of plume.

The density of embers departing the plume is then

$$\rho_{out} = -\frac{\partial n}{\partial x} = fn$$

and describes the density of embers immediately below the plume. It is plotted by the blue curve in Figure 5e.

In our example case, these embers fall from heights of up to about 3 km. Their fall velocity is 6 m/s, while the boundary layer wind is 15 m/s near that height, but less near the surface due to friction. Hence, they will travel up to 2.5 m horizontally for every metre they fall, or up to about 5 km as they fall from plume axis to ground.

The relationship between plume departure point and landing point of an ember is given by

$$x_{land} = x_{plume} + z_{plume}\frac{\overline{U_a}}{v_{fall}}$$

It is necessary also to adjust the landing density because the plume is sloped but the ground assumed horizontal, so embers falling from high in the plume travel further than those from nearer the ground. This adjustment is given by

$$\rho_{land} = \rho_{out}\frac{\partial x_{plume}}{\partial x_{land}}$$

and ensures that the integral along the plume of falling embers equals the integral of landing embers along the ground. The landing density is given by the orange curve in Figure 5e.

We now compare these predictions to explicit calculations from the data used in Thurston et al. (2017).

Comparison to Thurston et al. (2017)

Thurston et al. (2017) present the results of numerous simulations using an explicit ember transport model and analyse the landing distribution and flight times of the embers. They used a large-eddy model to simulate a buoyant plume as described above and used the time-varying turbulent wind fields so generated to transport embers with a prescribed fall velocity. Similar simulations, with the refinements of accounting for the finite inertia of the firebrands and for the altitude dependence of fall velocity, were presented by Thomas et al (2019) with broadly similar results.

Figure 10 compares the longitudinal distribution of landing distribution with a fall velocity of 6 m/s and three different wind speeds, from Thurston et al's (2017) simulations to the ember transport parameterisation. Note that the total number of embers depends on the source strength, which is not well known, so the vertical position of the curves in these plots is arbitrary. However, the shape of the curves, and the relative numbers landing at a given distance, is robust. The parameterisation tends to transport embers up to about 20% further than the explicit calculation at all wind speeds. The shape of the tail of the distribution at long distance is well captured at the lower wind speeds, but too extended at high. All three pairs of simulations exhibit a relatively flat distribution, with little decay in distance of the ember concentration, in the first 10 km or so, although some differences in shape are apparent. One curious feature is the small bump

apparent in the parameterisation near x = 8 km in the 7.5 m/s run, and near x = 13 km at 10 m/s, which is due to the core of the plume encountering the inversion, accompanied by an abrupt decrease in updraft and hence dropping of embers. This is likely a limitation of the bulk plume model, which we aim to improve in the future.

The parameterisation is similarly reasonably successful for embers with higher or lower fall velocities. For example, Figure 11 shows results for wind speeds of 10 and 15 m/s, and fall velocities of 4 and 8 m/s. As expected, the higher (lower) fall velocity leads to shorter (longer) travel distances, with transport well in excess of 20 km predicted for the fall velocity of 4 m/s – sufficient to carry the embers out of the large-eddy model domain! While there are some differences between explicit calculation and parameterisation apparent, the trends with wind speed and fall velocity are reasonably faithfully captured.

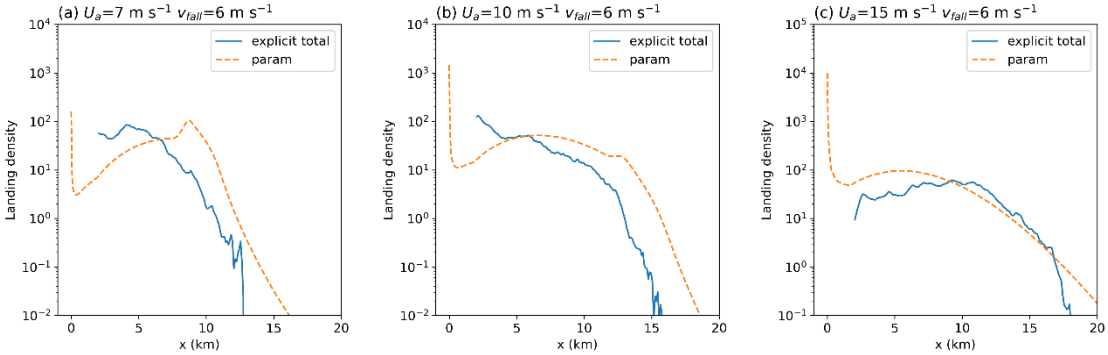

Figure 10: Longitudinal distribution of landing positions for embers with fall velocity of 6 m/s and three boundary layer wind speeds, measured at the top of the boundary layer, of (a) 7.5 m/s, (b) 10 m/s and (c) 15 m/s. The blue curve shows the calculations of Thurston et al. (2017) and the dashed orange curve is the parameterisation.

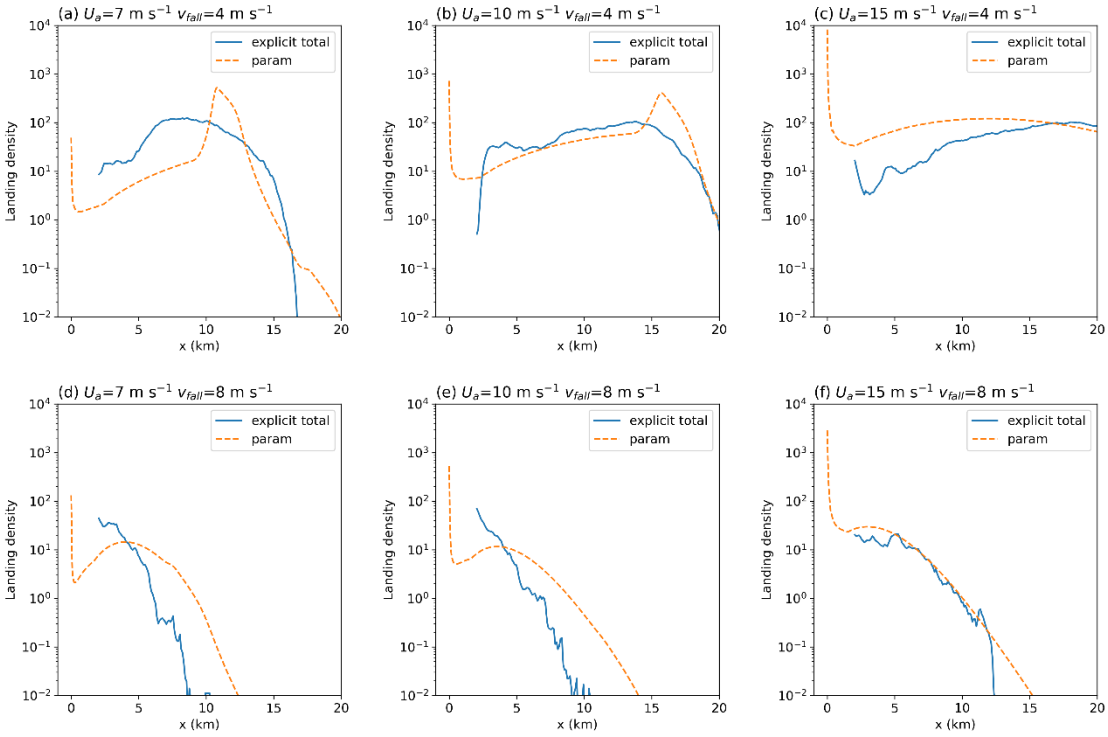

Figure 11: The same as Figure 10, except for different fall velocities. Panels (a) to (c) show data for a fall velocity of 4 m/s and panels (d) to (f) for 8 m/s. Panels (a,d) are for a nominal wind speed of 7.5 m/s, (b,e) for 10 m/s and (c,f) for 15 m/s.

Some properties of the model

Maximum ember transport distance

We consider fire characteristics fire radius r_0 and firepower Q_{plume}, meteorological characteristics wind speed and inversion height, and ember characteristic fall velocity. In each case, the wind profile is set to be a unidirectional logarithmic profile and are labelled by the free wind speed. The atmospheric temperature profiles consist of a neutral-stability layer up to the inversion, and potential temperature increasing at a rate of 1K/km above. The inversion layer is 500 m thick, and the stability is linearly interpolated between the above values within the inversion. The results are presented in Figure 12. Each panel considers the sensitivity to firepower and one other parameter: either wind speed, inversion height, fall velocity or fire radius. The fixed parameters in each plot have the values r_0 = 500m, wspd = 10 m/s, z_{inv} = 3 km, v_{fall} = 6 m/s as appropriate.

Maximum transport distance is insensitive to inversion height for low fire powers (Figure 12a, blue and orange curves) but increases approximately linearly with inversion height for the hottest fire (red curve). The 30 GW fire (green curve) exhibits both these behaviours, depending on inversion height. The reason for the

insensitivity to inversion height for cooler fires (and higher inversion for the 30GW fire) is that the plumes in these cases retain little vertical velocity when they reach the inversion, and hence nearly all the embers have already fallen out. In contrast, the Q_{plume} = 100 GW fire always reaches the inversion with a vigorous updraft, loses a substantial portion of its updraft and drops a substantial portion of its embers then. For higher inversions, this point is further downwind, leading to increased maximum transport distance.

The dependence on wind speed (Figure 12b) shows that maximum transport distance can either increase or decrease with increased wind speed. These contrary tendencies are because wind speed has a dual role – increasing wind speed transports the embers downwind faster, but also causes the plume to entrain more vigorously, lie flatter and have a weaker updraft, causing the embers to fall out sooner. The decreasing trend is found for weak fires or strong winds, consistent with this analysis.

The dependence on fall velocity (Figure 12c) is, in contrast, always monotonic. Embers with a greater terminal velocity tend to exit the plume sooner, leading to shorter transport distances.

Lastly, we consider the sensitivity to fire radius (Figure 12d). This case also exhibits monotonic behaviour, but with varying slopes of the curves. For two plumes with the same heat flux, but different radii, the narrower one will have the stronger updraft because the same heat flux must be carried aloft by a narrower column of air, requiring a stronger updraft. The narrower updraft will therefore transport embers further. In the case of the 100 GW fire, the slope is small because all cases shown reach the inversion at 3-km height with strong updrafts and drop the bulk of their embers there. The weakest case shown has insufficient updraft to carry embers except when very narrow, also leading to a relatively flat curve. In the intermediate firepower cases, the vertical growth of the plume and the strength of its updraft steadily diminish as the fire radius grows, reducing the maximum transport distance.

In all cases, maximum transport distance increases monotonically with firepower, because hotter fires lead to more vertical plumes with stronger updrafts, other things being equal. However, the effect of a hotter fire can range from very little, to a dramatic increase in transport distance, depending on the other parameters.

Several of these sensitivities have not, to our knowledge, been previously proposed for ember transport. We intend to further examine them using a LEM in the future.

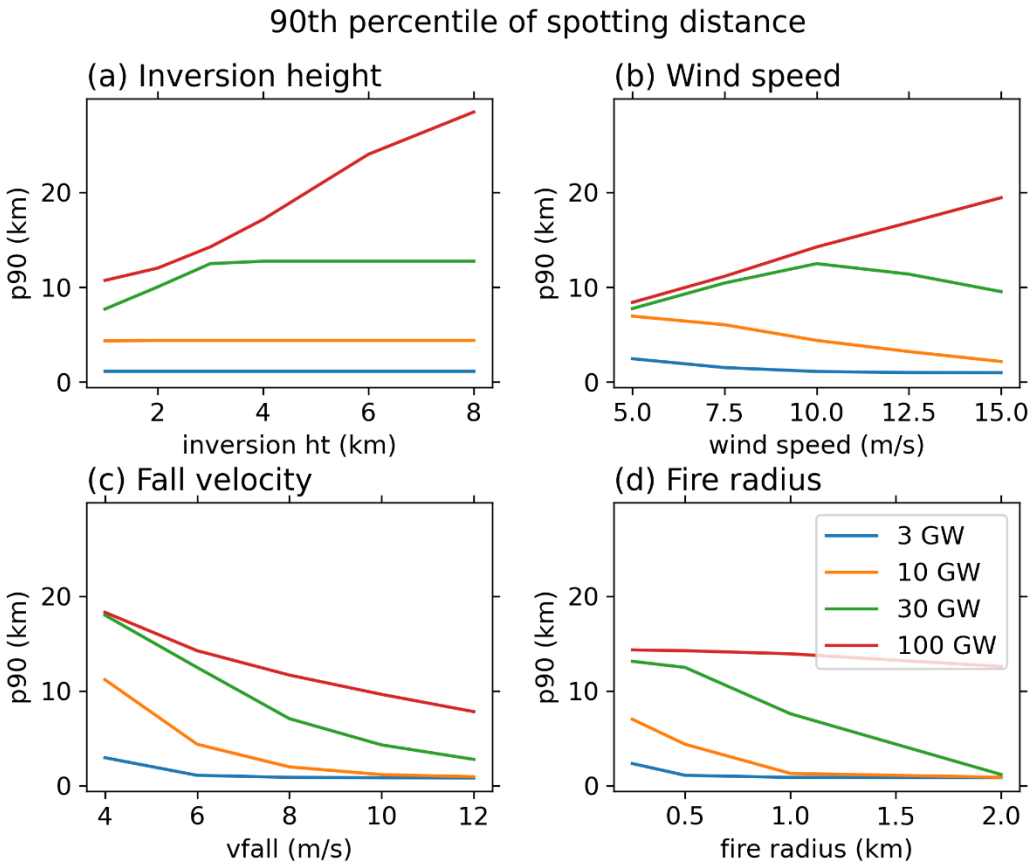

Figure 12: Sensitivity of 90th percentile spotting distance to atmospheric, fire and firebrand properties. Each panel considers 4 fire powers into the plume, H_{plume} = 3, 10, 30 and 100 GW, according to the legend in (d). Sensitivity to (a) inversion height, (b) wind speed, (c) fall velocity and (d) fire radius.

Does long-range ember transport possess multiple regimes?

Figure 11 tends towards an exponential decay (i.e., a straight line on these semilog plots), in contrast to the flatter distribution in the other cases.

A recent observational study of spotfire distribution (Storey et al 2020a) divides the distribution of the number of spotfires with distance from the fire front into three categories:

- Exponential: most spot fires within the first 250 m, decreasing numbers of spot fires in each bin out to the farthest spot fire and no gap > 1 km between any two spot fires.
- Multi-modal: Like type 1, except with a gap >= 1 km between any two spot fires.
- Other distribution: No prominent peak in the frequency plot, or the peak in spot fires beyond the first 250 m bin.

We choose to combine the first two categories for our purposes, noting that changes in the probability of an ember producing a new ignition or of the number of embers may be sufficient to transform one case into the other. We note also that their study is not exactly comparable with ours, for we are

considering all embers, whereas they consider only embers that produce ignitions. Embers with long flight times may have burnt out and be incapable of producing ignitions, while some may not fall into suitable fuel beds. Nevertheless, their first two cases appear like the exponential distribution already noted here with the less intense fires and larger fall velocities. The third case is plausibly the result of applying finite sampling to the distributions found here with intense fires and low-to-moderate fall velocities, that are either flatter or have a maximum some distance from the fire.

Given this apparent consistency between observations, the parameterisation and the explicit simulations, we suggest that these differences are likely real, and that plume-driven spotting leads to (at least) two regimes at the extremes: either an approximately exponential decay or a relatively uniform distribution of spots, although finite sampling of individual ignitions may obscure this picture. Between these extremes there may be a transition zone.

Obviously, it would be useful to be able to predict where on this range of behaviour a given fire falls. Key parameters will include the fire characteristics, meteorology and ember fall velocity. Regarding the fire characteristics, we see consistency between the sensitivity to fire intensity apparent in the parameterisation, and Storey et al's (2020b) identification of fire area as important.

Case Study of the Kilmore East Fire, 7 February 2009 ("Black Saturday")

Overview

The fires in southeast Australia on Black Saturday, 7 February 2009, are the worst on record in terms of lives lost, with 173 fatalities due to the direct effects (i.e. excluding the effects of smoke inhalation) of the fires. The deadliest of these fires ignited at Kilmore East (see map in Figure 14 for locations) at about 11:45 in the morning (all times in this section are given in Australian Eastern Summer Time, 11 hours ahead of UTC) and was responsible for about 70% of the fatalities on the day. It spread rapidly to the southeast under the influence of a hot, dry, vigorous and deep airstream, covering about 55 km in about 6 hours, before a marked wind change pushed the fire to the northeast then north. In its first 12 hours, it consumed over 100 000 ha of mostly forest and grassland. Some days later, it merged with another major fire (the Murrindindi fire) and was not eventually extinguished until after several weeks.

Cruz et al. (2012) provide a detailed reconstruction of the fire. Their analysis notes the significant role that spotting played in the rapid spread to the southeast and includes the times and locations of numerous spotfires. The longest spotting distance analysed was 34 km, which they note is similar to maximum spotting distances previously reported in similar forest types ((Hodgson, 1967; Cheney and Bary, 1969; McArthur, 1969). Spotting distances of this magnitude require not just long-burning embers, but also favourable atmospheric and fire conditions. The Kilmore East fire was certainly extremely powerful, with the estimated power exceeding 8 TW following the change when the north-eastern flank became the front. Although some of the spotfires were extinguished by fire crews and therefore did not contribute to fire spread, a substantial number did.

A detailed discussion of the meteorology, based on very high-resolution numerical weather prediction and Bureau of Meteorology observations, was given by Engel et al. (2012). They emphasise the deep boundary layer and strong winds ahead of the change, conditions which we have already seen favour extremely long-range ember transport in the ember transport parameterisation (Figure 12).

Toivenan et al. (2019) present a simulation of this day of the fire, using the coupled fire-atmosphere model, ACCESS-Fire. They were unable to obtain realistic fire spread using only the empirical fire-spread equation of McArthur (1966, 1967), with the simulated fire-run covering only about half the observed distance. When they incorporated additional ignitions based on the observed spotfire locations in Cruz et al. (2012), much better agreement was obtained.

As a notorious example of a fire where spotting appears to have made a substantial contribution to the spread, the Kilmore East fire is a suitable initial test for the ember transport parameterisation presented here when coupled to a fire spread model. Our goal is to examine the sensitivity of the simulated fire spread to embers, to compare the results to the observed spread, and to begin to discover the issues associated with coupling the ember parameterisation to a fire spread model, rather than to produce the most realistic possible simulation. Therefore, we shall use time-series of surface and column meteorology at a point rather than spatially varying fields, and a simplified representation of the fuels.

We couple the ember transport parameterisation to the Spark fire spread simulation system (Hilton et al. 2019). Spark is described in the user documentation as "a toolkit for simulating the spread of wildfires over terrain". The core of Spark is the solver, which is a computational engine for propagating a fire front. The rate at which the fire front spreads depends on factors including topography, fuel, and meteorology, is described by empirical equations (or fire spread models) such as those reviewed by (Cruz et al., 2015) and can be specified by the user. Spark reads necessary data (topography, fuels and meteorology), interpolates them in space and time where necessary, and makes them available to the fire spread equations through the solver. Output data, including fire perimeters, are written to either shape files or as a grid. The complete configuration for a given simulation, including the above data, details of the ignition(s), and simulation start time and duration, is stored in an extended mark-up language (xml) file. This xml file is usually prepared using the Spark graphical user interface, which also displays the output data. However, the xml file can also be prepared or modified by other means, and the shapefile and gridded outputs use standard file formats which can be read by other software. Indeed, there is also a command-line version of Spark, which operates off pre-prepared input files and completely bypasses the graphical interface.

Details of how we coupled the ember transport parameterisation to Spark are given below.

Model Configuration

We used v1.1.2 of Spark. The software package includes three test cases designed to demonstrate the simulator's capabilities. These test cases include implementations of the CSIRO grass and forest fire spread models (refs), which

we use here. It also includes code to diagnose fuel type from the ABARES Catchment-scale Land Use Model (CLUM), which we adopt with slight modification. The CLUM data gives land use on tiles which are roughly 50 m square for the whole of Australia. Several hundred land use categories are available, which are simplified to four fuel types: non-flammable, grass, forest and urban. The fuel classification is shown in Figure 14. We further adopt the methodology in the Spark test cases by using remotely sensed topography data at 1 arc-second (approx. 30 m) resolution.

While Spark allows some degree of user customisation (for instance, different empirical fire spread models can readily be coded), not all the data necessary to run the ember transport parameterisation was available within the simulator, with the most important omission being the fire power. Hence, the ember transport parameterisation was not implemented "inside of" Spark. Rather, we developed an external code wrapper to provide the interface. This driver analyses a Spark configuration xml file and breaks the single simulation into shorter steps of (typically) 5 to 30 minutes duration. For each of these steps, Spark is run to propagate the fire front. The fire perimeters are extracted from the output files and matched to those at the start of that simulation step to determine the increase in area of each fire patch, taking account of any mergers. That change in area is projected onto the fuel map to determine the fuel consumption and hence the fire power. In particular, the fire power H_{fire} is

$$H_{fire} = \frac{C \int_{\Delta A} m \, dA}{\Delta t}$$

where C = 1.86e7 J kg^{-1} is the heat of combustion of cellulose, m is the fuel load in kg m^{-2}, and ΔA is the area burnt within the time interval Δt. We take the heat input to the plume to be $H_{plume} = 0.7H_{fire}$ to allow for the loss of some energy through radiation. Sullivan et al. (2001) consider this question, noting the constant of proportionality as being between 0.6 and 0.83; we therefore fall in the middle of their range.

The fire patch is also analysed to determine the width of the head fire, which provides the source width of the plume. Those two fire parameters, together with the necessary meteorological data, are passed to the plume model, and the results of the plume model to the ember transport calculation. Recall that the ember transport calculation predicts the spatial density of the landing distribution. A random number generator is used to draw the desired number of ignitions from this distribution, and their travel time calculated. The fire perimeters at the end of the previous Spark run, together with these new future ignition locations and times, are written into the xml control file and other necessary details updated. Spark is then run to propagate the fire front further.

The ember transport parameterisation predicts ember landing positions, which is not the same as spotfires. Whether an ember produces a continuing spotfire depends on multiple factors, including whether the ember burns out before it lands, whether it lands in contact with sufficient fuel, and whether suppression resources control the new ignition. Rather than introducing a further probabilistic accounting for these factors, we instead choose to model only "viable" embers – that is, those that produce unsuppressed ignitions. We assume that the rate of production of such embers is proportional to the rate at which forest fuels are burnt and adjust the constant of proportionality to give roughly the same number

of spotfires as was documented by Cruz et al. (2012). The value used is 0.002 per hectare, which we note is surprisingly small. For simplicity, and because of insufficient knowledge, we do not impose a burnout time, either as a single fixed value or as a probabilistic distribution, although we acknowledge that this would clearly be more physically realistic. Cruz et al. (2012) lists the dominant forest tree species, which include several ribbon-barked types that are known to produce embers that can burn for over twenty minutes (Hall et al. 2015), so while this assumption may be tolerable here, it could be severely unrealistic in other forest types.

Testing the simulation framework with the ember processes turned off showed that this framework led to slightly slower fire propagation (by about 5%) than using Spark in its usual fashion. This slowing of fire spread appears to be due to the conversion of Spark's internal model state variable, a two-dimensional grid of the level set variable, to fire perimeter shapefiles, then back to the grid, at each step of the driver routine. As the purpose of this section of the paper is to demonstrate the impact of coupling, rather than preserve absolute accuracy, we have not attempted to adjust for this slowing. A better solution would have been to use the level-set grid, rather than the perimeters, as the interface, but unfortunately this version of Spark does not provide user access to the level set variable.

The system requires several times the computation time than native Spark. A small part (under 10%) of the slow-down is accounted for by the plume and ember calculations, and somewhat more by the housekeeping associated with fire power calculations. However, the bulk of the extra time is spent within Spark and is presumably due to the overheads of starting and stopping the simulator each step, including reading and writing data files and conversions between the level set grid and shapefile perimeter formats. For research purposes, this computational slowdown is acceptable. A forthcoming version of Spark will permit a more efficient interface (James Hilton, pers. comm., 2020)

Meteorology for the simulations is taken from the Bureau of Meteorology Australian Region Reanalysis, BARRA. Meteorological reanalyses use state-of-the-art numerical weather prediction systems to analyse historical observations and produce a dynamically consistent, four-dimensional, gridded record of the atmosphere. BARRA covers the Australian continent and surrounds on a spatial grid of about 12-km spacing, from the surface to over 40 km altitude and from 1990 to 2019, and was prepared using a version of the Bureau of Meteorology's operational assimilation and prediction system, ACCESS. Further details are in Su et al. (2019). An important advantage of BARRA is that it provides high quality upper air data needed for the plume calculation. A disadvantage relative to automatic weather station data is that the time frequency is relatively coarse, one hour instead of (up to) one minute. Thus, features such as wind changes are smoothed out in time. We choose to use just the time-series of a single column of data, rather than the full four-dimensional data, for simplicity. For the simulation presented here, this column is taken from the closest grid point to the ignition (see map in Figure 14 for locations). The meteorology is summarised in Figure 13, which show that the temperature increased gradually during the afternoon, before dropping rapidly in the lower approximately 1 km as the wind change arrived. The humidity was very low pre-change, increasing afterwards. Winds were strong

and from the north, trending gradually towards the northwest, before turning abruptly from the southwest, tending southerly, with the change. There was a short-lived period of weak winds, followed by a pulse of strong winds, about the time of the change. The boundary layer, indicated by the layer of very low static stability, was initially about 3 km deep increasing gradually to almost 5 km during the afternoon, but a marked near-surface stable layer was prominent after the change arrived.

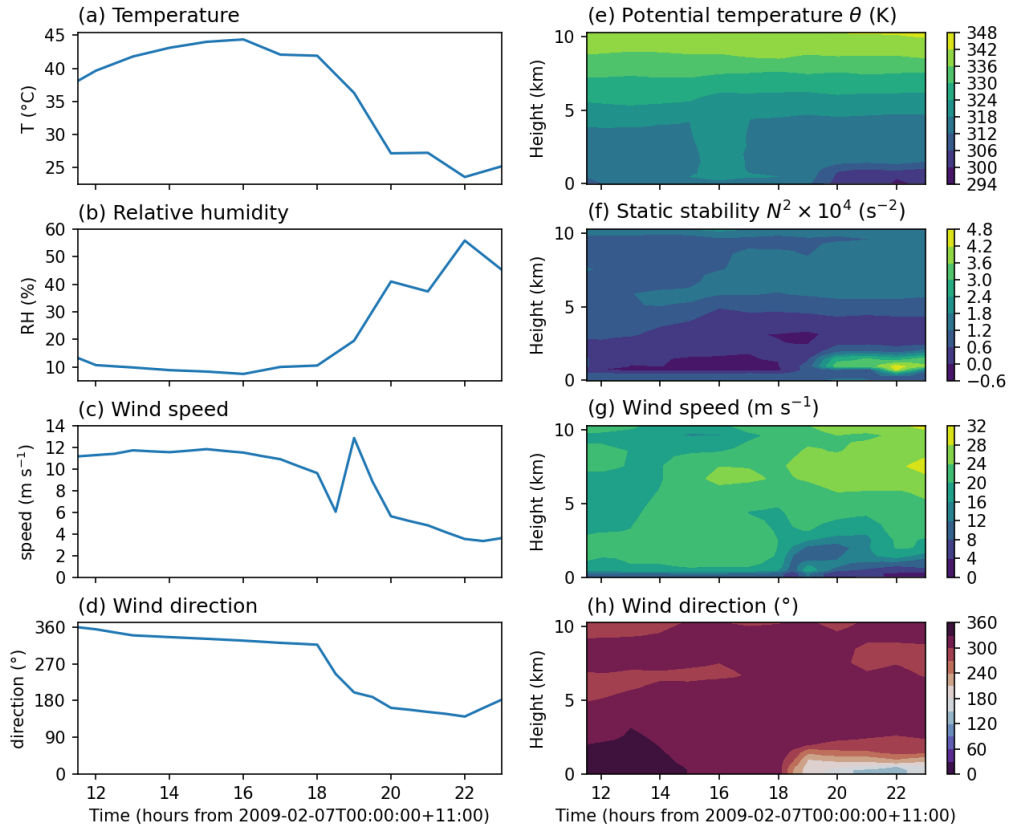

Figure 13: (a – d) Time series of surface meteorology for the simulations of the Kilmore East fire according to the BARRA reanalysis, showing respectively air temperature, relative humidity, wind speed and wind direction. (e – h) Time-height series of upper meteorology for the same simulations, showing respectively potential temperature, static stability, wind speed and wind direction.

Results

Without spotfires

We begin by presenting results for a simulation with no ember transport. The fire spreads to the southeast, reaching 36 km from the ignition point (roughly midway between Whittlesea and Kinglake) when the wind change arrived, which although relatively rapid was nevertheless much less than observed. The area burnt reached about 500 km^2 at midnight, about half of Cruz et al's (2012) estimate, and peak power was about 850 GW on the wind change, roughly one

tenth of their estimate. The simulated progression post-change is also markedly slower than observed.

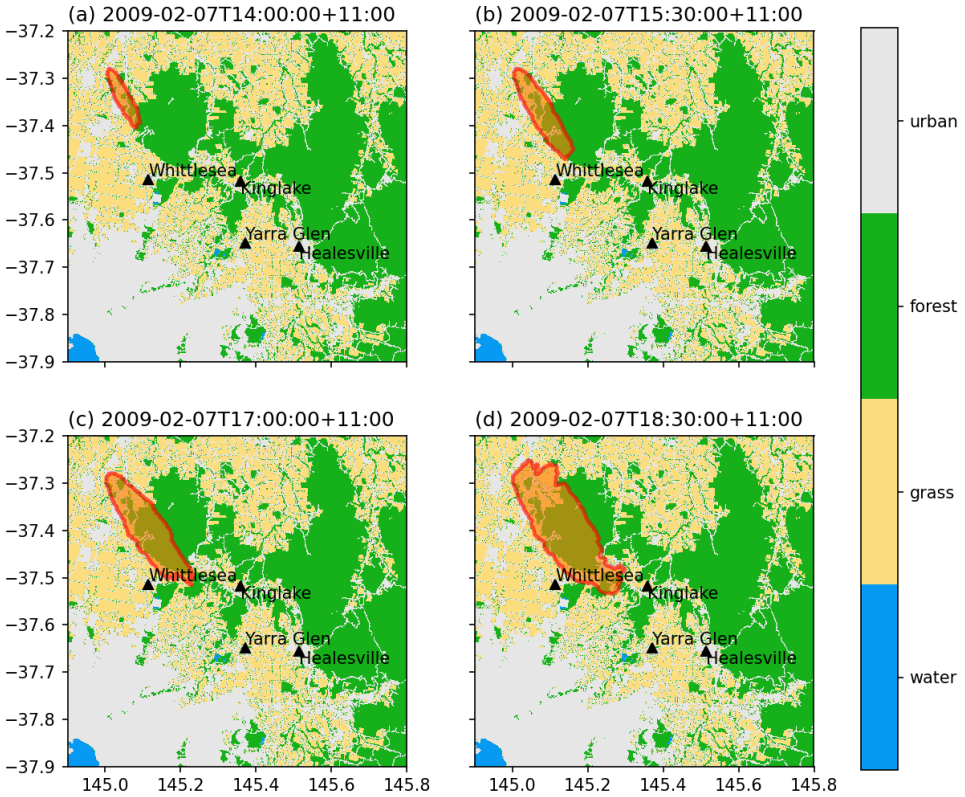

Figure 14: Simulated fire perimeters, without ember transport, at 90-minute intervals through the afternoon. The fire perimeters are shown by the red contours and orange shading, and fuel types according to the colour bar.

Deterministic

Figure 15 presents a simulation with spotting active, showing fire perimeters at 90-minute intervals through the afternoon. Spotfire ignition points are indicated by magenta dots. The spread to the southeast, defined by the furthest spot, is 63.8 km, about double that of the no-ember simulation. While the fire perimeter in the first half of the spread is reasonably continuous, spotting accelerates the spread there and causes the extension of the fire into the Yarra Glen to Healesville vicinity. The fire perimeter in this region is quite irregular, especially prior to the change, emphasising the role of spotting and consistent with Cruz et al's (2012) reconstruction. Following the change, these spots consolidate into a more contiguous area, albeit still with an irregular southwestern edge. Embers thrown during the wind change incorrectly travel to the east, due to time-interpolation of hourly wind data and our not accounting for changes of the wind while the embers are in flight.

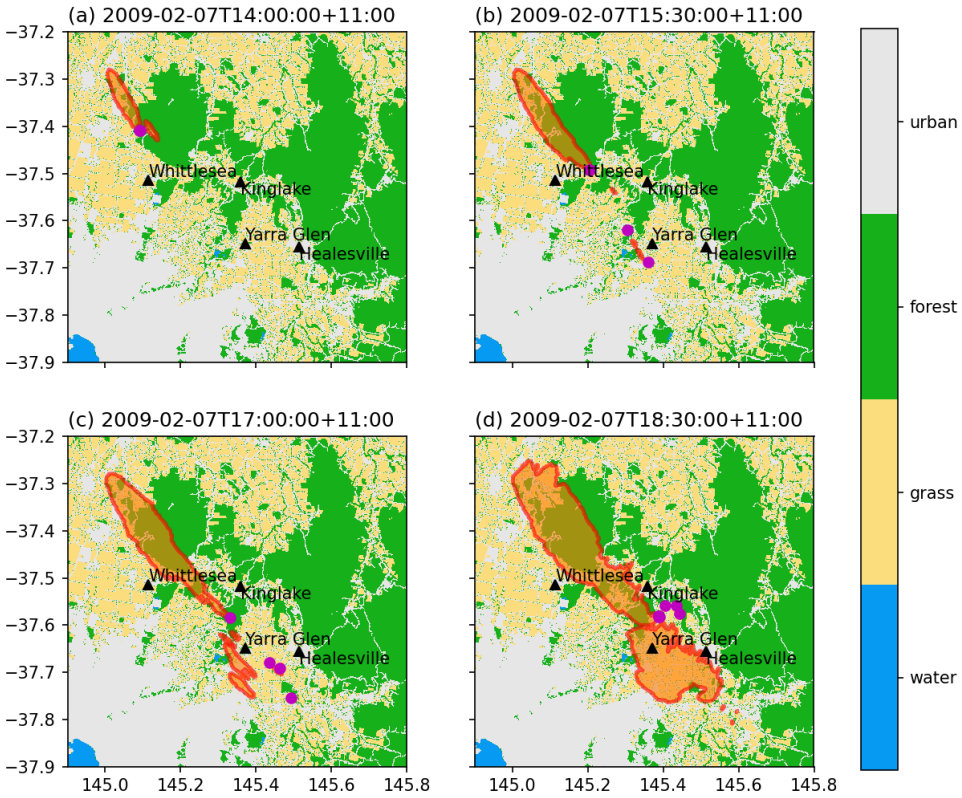

Figure 15: Simulated fire perimeters, with ember transport, at 90-minute intervals through the afternoon. The fire perimeters are shown by the red contours and orange shading, and fuel types according to the colour bar. Magenta dots indicate the landing positions of embers launched at that time, but note that the landing time of those embers can be several tens of minutes later.

Some properties of the fire and ember transport are shown in Figure 16. The total area reaches over 1000 km^2 during the evening, in good agreement with Cruz et al.'s (2012) statement that the fire burnt "100,000 ha in less than 12 hours". Their estimates of rate of burning increased steadily during the afternoon to around 100 km^2/hr prior to the change, increasing to 350 km^2/hr with the change and then declining, in excellent agreement with our results. Prior to the wind change, the estimated total power (Figure 16b) is about 600 GW between 3 pm and 5 pm, before almost doubling to about 1.1 TW at 6 pm, whereas the estimates of Cruz et al. increase from about 150 GW to 2.1 TW during this period. The total fire power peaks at over 1.5 TW on the wind change substantially lower than Cruz et al estimated (8.8 TW), and likely pointing to a deficiency in our crude fuel model. Most of the time, a single fire patch containing the original ignition accounts for nearly all the power, but at times secondary patches are up to about 20% of the total prior to the change, and 50% afterwards, before being incorporated into the main perimeter. Figure 16c plots the 50th and 90th percentile of spotting distance. The latter value persists at about 30 km for several hours in the afternoon. Individual spotting distances, as chosen by the random number generator, similarly peak near this value, in excellent agreement with Cruz et al's (2012) analysed maximum spotting distance of 34 km. The maximum spotting distance of a little over 30 km appears to be set by the meteorological

environment, particularly by the increase in static stability above about 4.5 km altitude, and the fire power. Notably, very few of our numerous simulations of this event produced spotting distances exceeding 35 km. The omission of a burn-out time in our simulation, in conjunction with the good agreement, suggests that the maximum spotting distance analysed by Cruz et al. (2012) was not limited by the combustion characteristics of the embers in this case.

The fire spread following the wind change is less than observed, even with the contribution of spotting. The coupling of the ember transport parameterisation is not considered satisfactory in this case, because the assumption of a single plume emanating from the head-fire is untenable when the over 50 km long flank has become the head. A more sophisticated coupling to the fire spread simulator, and improved knowledge of plume behaviour in fires of this nature, will be necessary to resolve this issue.

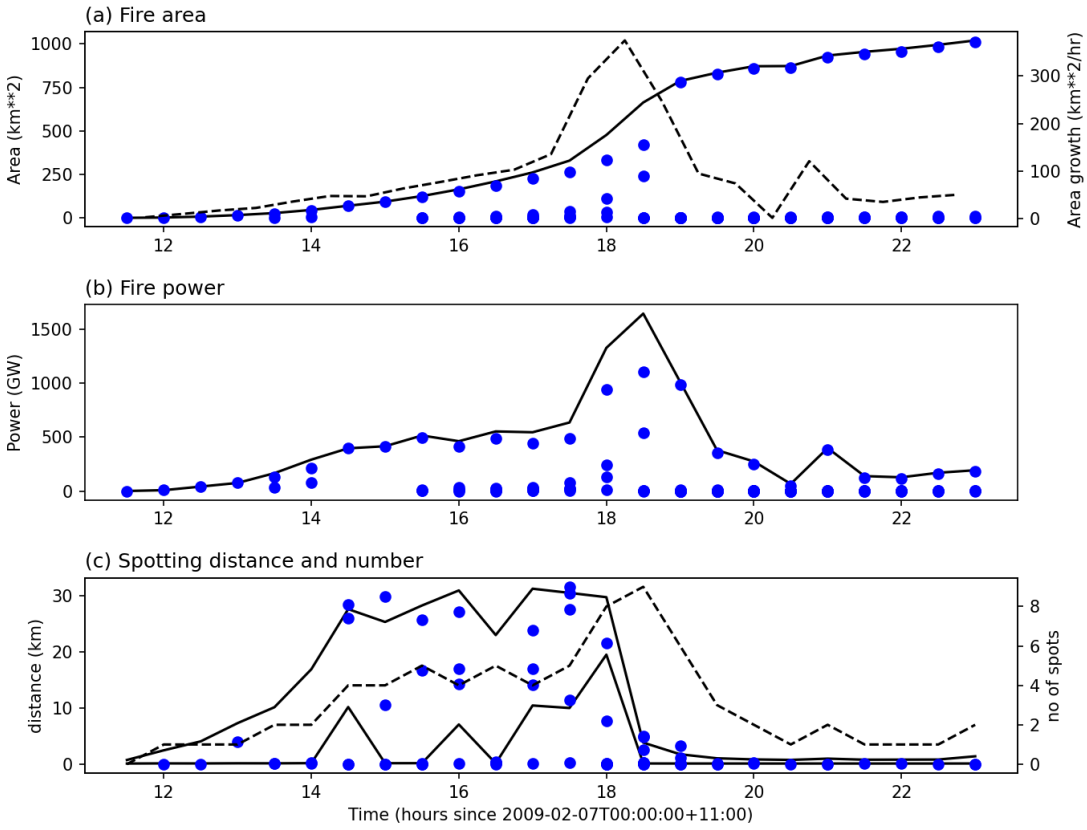

Figure 16: Statistics of the simulated fire spread. (a) Area of individual fire patches (blue dots) and total (black curve). (b)) Power of individual fire patches (blue dots) and total (black curve). (c) 90th and 50th percentiles of spotting (black curve), actual spotting distances from that patch (blue dots), and number of spots (dashed black curve, right-hand axis). All for just the most powerful fire patch at each time

Ensemble

Since the parameterisation explicitly incorporates the stochastic nature of the spotting process, it is natural to consider an ensemble simulation. We generated another 29 simulations with the above settings, differing only in the state of the random number generator that generates landing positions. A complete

ensemble method would also represent uncertainty in other aspects of the simulation, including fuel characteristics and the meteorology, but for simplicity and ease of interpretation these were omitted.

Maps of the proportion of ensemble members that burnt a given point, prior to either 18:00 (about the time of the wind change) or midnight, is shown in Figure 17. For convenience, we shall call these probabilities, but recall that our ensemble methodology does not account for several important sources of uncertainty. The northwestern part of the fire run has probabilities of almost one, these areas were burnt in the no-ember simulation. A narrow fringe of low probabilities extends from the southwest flank of this region to the town of Whittlesea. The ensemble members that contributed to this band experienced one or two relatively long spots early in the simulation, when the winds were more northerly, extending them to the south. The southeastern portion of the fire-affected area has much lower probabilities, consistent with our examination of the deterministic run which showed that the spread into the Yarra Valley was caused by spotting. Prior to the change, these probabilities are quite low, indicating that they consist of individual spots with little overlap between members. After the change, these ignitions extend to the northeast then north, with more overlap between ensemble members and consequently higher probabilities. Although the fire stopped a few kilometres to the west and south of Healesville in reality, we see a significant probability of spread further to the southeast after 6 pm. This additional area of fire potential originated as spots immediately prior to the change with flight times of half an hour or more. We presently calculate ember trajectories using the atmospheric flow at the time of launch, a source of unrealism, but relaxing this assumption might not produce large changes in the results, noting that these embers were simulated to reach heights of several kilometres and the change was less than 1 km deep (Engel et al. 2012). The towns of Yarra Glen and Healesville show significant probability of fire but were not burnt on the day. Reasons for this may include the success of suppression (several spotfires near Yarra Glen were suppressed by fire crews, M. Cruz, personal communication 2020), and that at this late stage of the summer and after a prolonged drought, paddocks and cropland likely had lower fuel loads than is typical for the grassland spread model used here.

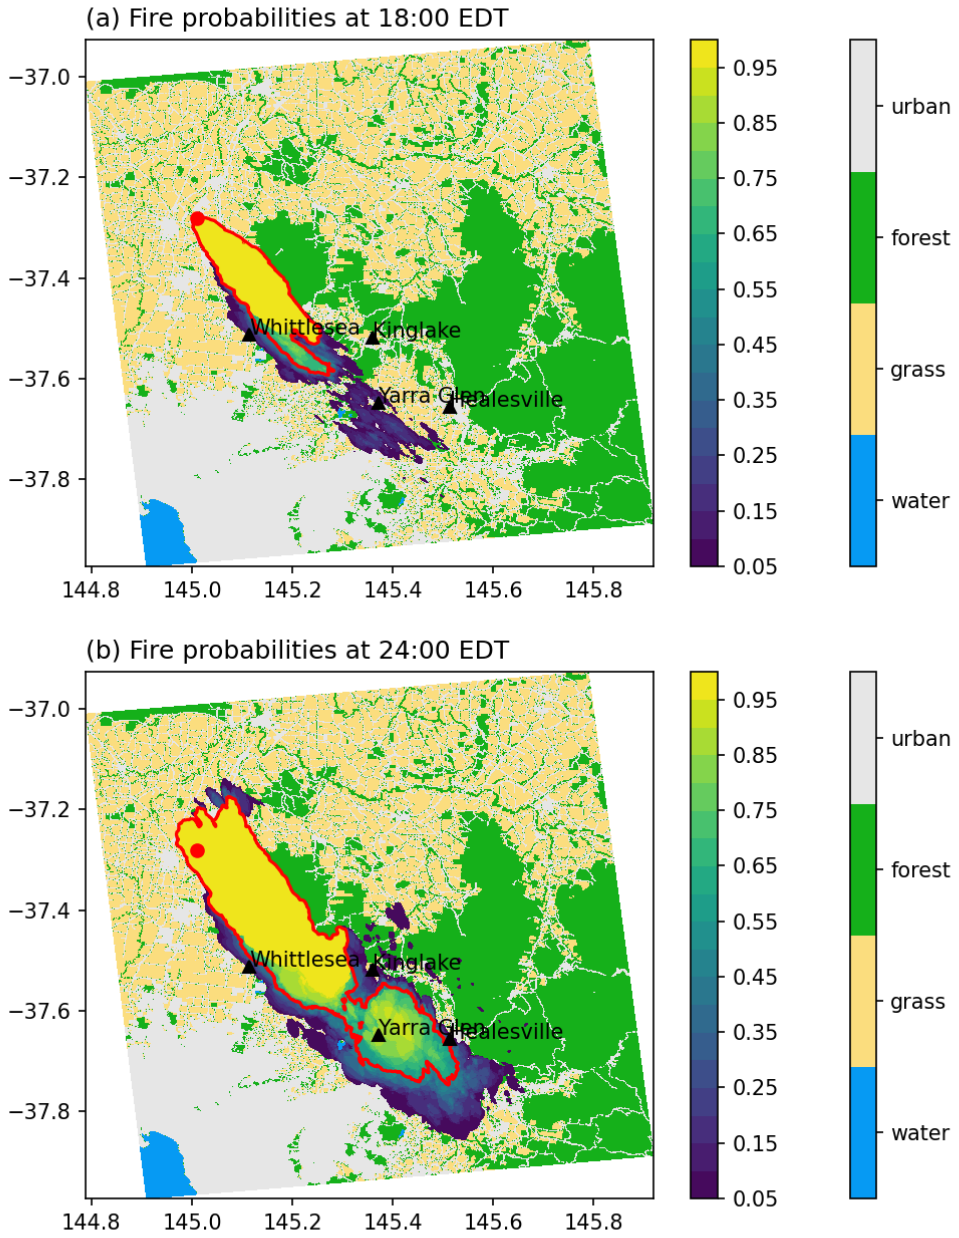

Figure 17: Fire probabilities at (a) 18:00 EDT, about the time of the wind change and (b) 24:00 EDT for the ensemble simulation described in the text. Grey shading shows the fuel type, and colour shading the probabilities, with the 50% probability highlighted by the red contour. Key locations are indicated by black triangles, and the ignition point by the red dot.

Conclusions

We have developed a physically based parameterisation of ember transport in bushfire plumes, which sheds light on the long-range spotting process and is sufficiently computationally fast to be useable in fire management operations.

The parameterisation consists of four components: a bulk model of the time-mean plume, a representation of turbulence within the plume, a probabilistic representation of ember transport within the plume, and a model of transport of falling embers beneath the plume

The parameterisation produces results that are reasonably consistent with the explicit calculations of Thurston et al (2017) and Thomas et al (2019).

Ember landing distributions appear to fall into two regimes, with a transition between. These regimes are a somewhat flat distribution, which occurs for intense fires and low fall velocities, or an exponential decay with distance for less intense fires and/or higher fall velocities. These regimes are consistent with the analysis of Storey et al (2020), although they note three regimes. We consider that a simple tool for quantifying the risk of spotting in the long-range regime, based on the work herein, would be possible and would be of significant benefit to bushfire management.

We have examined the variation of the predicted maximum transport distance with meteorological parameters, fire characteristics and ember fall velocity. More powerful, or smaller, fires create plumes with stronger updrafts that transport embers further. Embers with a higher terminal fall velocity travel less distance. The effect of wind speed is complex. Plumes from very hot fires send embers further as the wind speed increases, but for weaker fires, the increased entrainment into the plume with stronger winds reduces the ability of the plume to carry embers and hence the maximum spotting distance. Inversion height can have a marked effect in parts of the parameter space. Provided that the plume is strong enough to carry a significant number of embers to the inversion, the marked drop in plume updraft there leads to a large ember fallout rate and often defines the maximum transport distance. In these cases, varying the inversion height can substantially change the maximum spotting distance.

We have coupled the ember transport parameterisation to the Spark surface fire spread model and presented deterministic and ensemble simulations of the Kilmore East fire of 7 February 2009. Without embers, the simulated spread is much less than observed. With embers, there is substantial variation in the area burnt due to the stochastic nature of the parameterisation, but the contour of 50% probability of area burnt is in good agreement with the observed propagation of the fire prior to the wind change.

These results are promising at the current stage of development but point to several areas where improved knowledge is required and which we envisage will be refined in the future. Knowledge of embers is incomplete, particularly the rate at which they are produced, and their distribution of initial mass. The extraordinarily low ember production rate we used requires further study and clarification. Our parameterisation relies heavily on a series of LEM simulations, but the hottest fire these simulations considered was H_{plume} = 39 GW[1], much less than the H_{fire} of several hundred GW pre-change and over 1 TW post-change in our coupled simulation, and all were for a fire consisting of a 250-m radius circle. The representation of plume structure (both mean and turbulent) needs to be validated against hotter and more realistically shaped fires, with the problem of very long fire-fronts following wind changes being a particular issue. We have

[1] The LEM used for these experiments became unstable with hotter fires.

also considered only relatively simple boundary layer structures so far, with little known about plume rise through directional wind shear, a potentially important issue after the wind change. Plume rise in the presence of significant vertical motion, e.g. due to mountain waves, also seems an important topic. Nor have we yet considered the possible impact of condensation within the plume.

The analysis herein also points towards future operational needs for more quantitative prediction of fire spread and behaviour. Plume dynamics is an essential ingredient of long-range ember transport, with fire power being a poorly known but key factor. In contrast, the meteorological factors are much better known, both during real time operations and for post-event analyses. Plume dynamics, and hence fire power, also form an important input to use of the recently developed pyrocumulonimbus forecasting technique described by (Tory et al. 2018, Tory and Kepert 2020), and can, in principle, be estimated with sufficient accuracy from the rate of fuel consumption. We recommend that techniques be developed and implemented to estimate this quantity to aid in real-time fire management.

Practical Implications

We have demonstrated a way of calculating ember transport which requires little computer time. There are several possibilities for making this into a practically useful product:

- As part of a fire-spread simulator, as demonstrated here with Spark.
- As part of a coupled fire-atmosphere model such as ACCESS-Fire.
- A stand-alone system for a given fire, that gave a predicted ember transport range. This prediction could be based on either observed or predicted meteorological data and given for a range of fire power values.
- A threshold product for spotting exceeding some range. For example, we could calculate and display maps of the minimum fire power that would give 90th-percentile spotting distance over 5 km, based on NWP data. These would indicate, for example, if conditions were expected to become much more favourable for long-range spotting later in the day. Interpretation would be similar to the PFT (described below): low values of the threshold would indicate atmospheric conditions that were favourable for long-range transport.

A limitation of the present product is that it has no information on the characteristics of the embers. For forest types which produce only short-burning embers, the transport time will greatly exceed the burnout time and the maximum spotting distance will be set by the burnout time. However, even in these cases spotting still requires sufficient plume development, which the model will diagnose. The utility of an operational product will likely be improved by incorporating information on the expected ember characteristics.

LARGE-EDDY SIMULATION OF PYROCUMULUS

Summary

Pyro-convective clouds have the potential to increase surface rates of wildfire spread, enhance ember transport, generate lightning and tornadic storms, and alter Earth's radiative balance through the stratospheric injection of combustion products. It is well known that sufficient moisture is required for a buoyant plume above a wildfire to generate pyro-convective cloud, and indeed this happens with some degree of regularity. However, the typical source of this moisture is less clear, with conflicting evidence available regarding the significance of pyrogenous moisture (that is, moisture released by combustion) in the development of moist pyro-convection. Here we attempt to fill this knowledge gap by quantifying the relative importance of environmental and pyrogenous moisture in pyrocumulus development. We expand upon previous case-study-based modelling approaches by performing large-eddy simulations of buoyant wildfire plumes of different intensities and over the parameter space of environmental moisture and pyrogenous moisture. Intense fires in sufficiently moist atmospheres can produce pyro-convective cloud in the absence of any moisture from the combustion process and the cloud is deeper for the more-intense fires and moister atmospheres. The addition of pyrogenous moisture only has a small effect on the development of pyro-convective cloud. Further analysis reveals that this is due to substantial entrainment of relatively dry environmental air into the sub-cloud buoyant plume. Entrainment dilutes pyrogenous moisture to the extent that very little remains once the plume has ascended to the level of the cloud base, suggesting that relative to environmental moisture, pyrogenous moisture plays only a modest role in pyrocumulus development. This has implications for forecasting pyro- convective cloud and its impacts.

Introduction

Intense heating of the air in the vicinity of a wildfire causes deep ascent. If deep enough to raise air above the lifting condensation level, this ascent leads to the formation of pyrocumulus (pyroCu) or pyrocumulonimbus (pyroCb). The generation of such pyro-convective cloud by wildfire plumes is significant because of the many ways in which it can potentially influence fire behavior. For example, moist pyro-convection has been associated with (i) amplified burn and spread rates (Peterson et al., 2015); (ii) tornadogenesis (McRae et al., 2013); (iii) ignition of new fires by pyroCb lightning, noting that pyroCb lightning conditions favor hotter and longer-lived lightning strikes (Rudlosky and Fuelberg, 2011); and (iv) enhanced spotting potential due to deeper and stronger updrafts (e.g., Koo et al., 2010) as a result of latent heat release in the convective column. The invigorated updrafts in moist pyro-convection are also capable of penetrating the tropopause, leading to stratospheric injection and hemispheric-scale transport of aerosol and other combustion products (Fromm et al., 2006, 2010). Therefore, knowledge of the processes that contribute to the development of moist pyro-convection in wildfire plumes is an important component in understanding the behaviour of fires, their impacts and potential climatic influences.

Here we use large-eddy simulations to investigate the generation of pyro-convective clouds by buoyant wildfire plumes. Firstly, we perform simulations over a range of fire intensities and environmental moisture levels. We then repeat a subset of these simulations with differing amounts of pyrogenous moisture, allowing us to assess the relative importance of environmental moisture and fire-derived moisture in the development of pyro-convective clouds.

Methodology

The Large-Eddy Model

The simulations presented here are performed with Version 2.4 of the UK Met Office Large-Eddy Model (LEM). The LEM is a high-resolution cloud-resolving model which solves a quasi-Boussinesq, anelastic, equation set on a three-dimensional Cartesian grid. Sub-grid stresses are parametrised using a stability-dependent version of the Smagorinsky-Lilly scheme, described by Brown et al. (1994). Moist processes are represented by the three-phase microphysics scheme of Swann (1998), which accounts for phase changes between vapour, liquid (cloud and rain) and frozen (snow, ice and graupel) water species. Our simulations are performed without any radiation parametrisation and the Coriolis force is also neglected, a reasonable assumption given the scale of the plumes simulated here. For further details about the LEM, the reader is referred to Gray et al. (2001) and references therein.

The LEM has traditionally been used to simulate a wide range of common turbulent and cloud-scale atmospheric phenomena, from the dry, neutral boundary layer (e.g. Mason and Thomson, 1987) to tropical deep convection (e.g. Petch and Gray, 2001), although more recently it has been used to study the behaviour of buoyant plumes from fires. Devenish and Edwards (2009) used the LEM to perform simulations of the Buncefield oil depot fire in the United Kingdom, the largest fire in Europe since the Second World War. The model was able to successfully reproduce the observed height and multi-directional spread of the plume. Devenish et al. (2010) performed further validation by comparing the entrainment and radial spreading rates of buoyancy and vertical velocity in LEM-simulated plumes to theoretical plume models.

All simulations presented here are performed on a model domain of 256 by 256 grid points in the horizontal, with a uniform horizontal grid length of 50 m, resulting in a domain size of 12.8 km by 12.8 km. The domain has a stretched vertical grid of 256 points in the vertical, with a vertical grid length of 20 m near the surface, increasing smoothly to 50 m at a height of 1.5 km. The vertical grid length is constant at 50 m between heights of 1.5 and 10 km and then smoothly increases to 60 m between 10 km and the model top, which is situated at 12.7 km. The lateral boundary conditions of the model are doubly periodic and a no-slip condition is applied at the lower boundary, with a specified roughness length of 0.05 m. A free-slip condition is applied at the upper boundary of the model domain in conjunction with a Newtonian damping layer in the top 2.5 km of the domain, to prevent the reflection of vertically propagating gravity waves. The horizontal extent of the domain is sufficient that over the duration of the simulations presented here the pyro-convective clouds are not affected by the periodic lateral boundary conditions. None of the pyro-convective clouds

deepen enough to encounter the Newtonian damping layer near the model top.

Experimental design

The wildfire plumes presented here are simulated using a two-stage process. The first stage is a "spin-up" phase, in which a realistic atmospheric boundary layer is generated. The model output at the end of the spin-up phase is then used as a set of initial conditions for the second stage, in which the buoyant plume from a wildfire is simulated.

At the beginning of stage one the model is initialised with horizontally homogeneous conditions, designed to be representative of a high fire danger day. An initial potential temperature profile is specified, consisting of a 4.0-km deep well-mixed layer with a constant value of 310 K. Above this, a troposphere with a constant lapse rate of 3 K km^{-1} is specified from the top of the mixed layer to the model top. The initial water vapour mixing ratio profile is also specified as a constant value throughout the 4.0-km deep well-mixed layer, and then reduced above the mixed-layer top at a rate such that the relative humidity remains constant throughout the troposphere. The model is initialised with zero environmental wind, allowing us to concentrate on the effects of the thermodynamics on the formation of pyro-convective cloud.

In order to explore the sensitivity of the development of moist pyro-convection to the environmental moisture, a range of simulations is performed in stage one. These simulations differ only in their initial value of the water vapour mixing ratio. Five simulations are performed, with constant mixed-layer water vapour mixing ratios, q_{vML}, of 2.0, 2.5, 3.0, 3.5 and 4.0 g kg−1. Random perturbations, with maximum amplitude of ±0.1 K, are added to the initial potential temperature field within the atmospheric boundary layer, in order to break the model symmetry and promote the development of turbulence, and a sensible heat flux of 50 W m^{-2} is applied uniformly at the surface. The model is then integrated forward in time with a variable time-step, determined by the Courant–Friedrichs–Lewy condition, until the boundary layer has spun up to a quasi-steady state, as determined by the time-series of the domain-average turbulent kinetic energy (TKE). At the end of the spin-up period we have five sets of initial conditions for use in stage two of the process, the simulation of wildfire buoyant plumes. These initial conditions are shown in Figure 18, in which the temperature and dew point traces together form the classic inverted-V profile, which is characteristic of conditions under which pyro-convective cloud forms (e.g. Peterson et al., 2017).

Then, in stage two of the simulation process we take the model state at the end of the first stage of the simulation and generate a buoyant wildfire plume, using the method employed by Devenish and Edwards (2009). A sensible heat flux anomaly is applied at the model lower surface boundary over a circular area of radius 250 m, located in the centre of the model domain, equidistant from both lateral boundaries. It should be noted that there is no direct feedback of the atmosphere back onto the fire in these simulations and therefore no capacity to simulate fire spread. This means that in this configuration it is possible to isolate and consider only the interaction between the buoyant plume and the thermodynamic environment in which it exists, without the extra complication of surface-based fire behaviour feedbacks.

Fires of four different intensities, Q, of 5, 10, 20 and 30 kW m^{-2} are simulated and in each case the heat flux is linearly ramped up from zero to its peak value over five minutes, held constant for one hour, and then linearly ramped down back to zero over five minutes. The combination of five different environmental moisture profiles and four different fire intensities gives us twenty simulations in total. We denote these simulations "dry" fires, as there is no representation of the pyrogenous moisture flux, either from moisture within the fuels or from the combustion process itself. These dry fires permit the assessment of the sensitivity of pyro-convective cloud to environmental moisture and fire intensity.

Additionally, a set of "wet"-fire simulations is performed, in which all five Q = 30 kW m^{-2} dry-fire simulations are repeated twice, once with a latent heat flux, L_E, of 2.1 kW m^{-2} and once with L_E = 11.4 kW m_{-2}. These values are chosen to represent the range of realistic ratios of heat to moisture fluxes in wildfires, based on the calculations of Luderer et al. (2009), and are denoted the "least-moist" and "most-moist" fires, respectively. These wet fires permit the assessment of the sensitivity of pyro-convective cloud to pyrogenous moisture, and in combination with the dry fires an assessment of the relative roles of environmental and pyrogenous moisture.

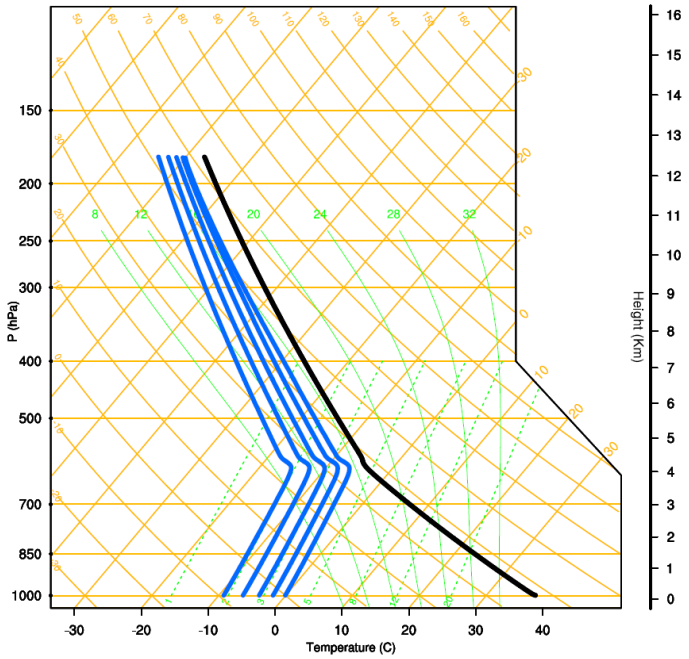

Figure 18. Skew-T log-p diagrams of the environmental atmospheric conditions spun up at the end of stage one. All simulations have the same temperature profile, shown in black, and the five different dew-point profiles are shown in blue.

Results

"Dry" fires

Figure 19 shows a snapshot from the Q = 30 kW m^{-2} fire intensity and q_{vML} = 4.0 g kg^{-1} mixed-layer water vapour mixing ratio simulation, 22 minutes after the fire has reached maximum intensity. This is the simulation with the moistest background conditions and the highest fire intensity. The wildfire buoyant plume has a strong updraft throughout most of the extent of the boundary layer, with a maximum vertical velocity in excess of 25 m s^{-1} (Figure 19 (a)). The plume decelerates as it penetrates into the stably stratified troposphere, but the ascent is deep enough to trigger the formation of a pyro-convective cloud. The cloud base is located at 4.5 km above ground level (AGL) and the cloud extends to an altitude of 7.5 km AGL (Figure 19 (c)). Although the updraft core of the buoyant plume is reasonably slender within the boundary layer, the horizontal extent of the pyro-convective cloud is much greater, having a width in excess of 3 km (Figure 19 (d)).

Extensive condensation within the cloud, particularly notable in the liquid water mixing ratio field around (x, z) = (0.5, 6.5) km, leads to the release of latent heat which is evident in the potential temperature perturbation field at the same location (Figure 19 (b)). There is a co-located updraft greater than 10 m s^{-1} within the pyrocumulus cloud, which is separated from the main plume updraft within the boundary layer. This resurgence of the updraft is due to latent heat release increasing the local plume buoyancy.

Figure 20 shows a snapshot of the liquid water mixing ratio vertical cross-sections for all twenty combinations of Q and q_{vML}, at 22 minutes after the fires have reached full intensity. The potential formation of a pyro-convective cloud and the properties of that cloud are dependent on both the intensity of the fire and the environmental moisture. Pyrocumulus is more likely to form if the environment is moist and the fire intensity is high. The size of a pyro-convective cloud and its cloud-top height both increase with environmental moisture and fire intensity. In the driest environment, with a mixed-layer water vapour mixing ratio of 2.0 g kg^{-1}, pyro-convective cloud only forms for the most intense fire and the cloud that does form is very small. As the environmental moisture increases, pyro-convective cloud forms for weaker and weaker fire intensity, and by the time the mixed-layer water vapour mixing ratio reaches the maximum value of 4.0 g kg^{-1}, pyro-convective cloud forms for all values of fire intensity. For fixed fire intensity, increasing the mixed-layer water vapour mixing ratio lowers the cloud-base height, from 5.7 km AGL for the driest to 4.5 km AGL for the moistest. Conversely for a fixed mixed-layer water vapour mixing ratio, increasing the fire intensity does not substantially affect cloud-base height.

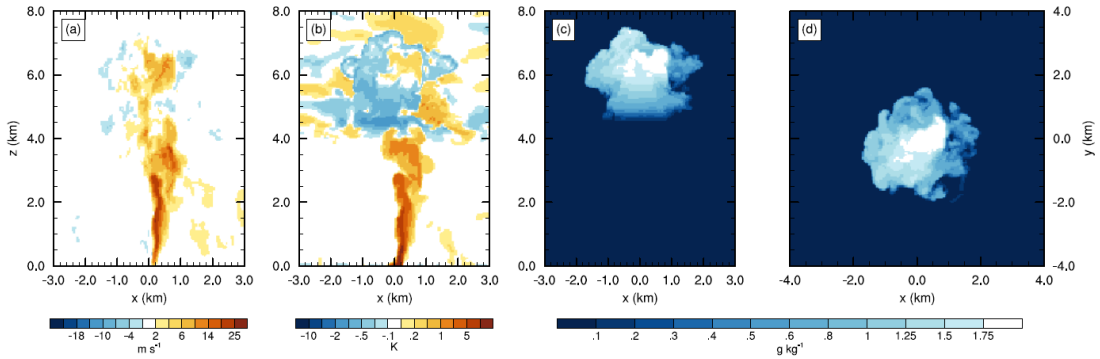

Figure 19. Vertical cross-sections of the instantaneous (a) plume vertical velocity, m s^{-1}, and (b) plume potential temperature perturbation, K, both in the y = 0 plane. Instantaneous liquid water mixing ratio, g kg^{-1}, maximum values in the (c) y– and (d) z–directions. All plots are shown at 22 minutes after the fire has reached full intensity and are for the Q = 30 kW m^{-2} fire intensity and qv = 4.0 g kg^{-1} mixed-layer water vapour mixing ratio case.

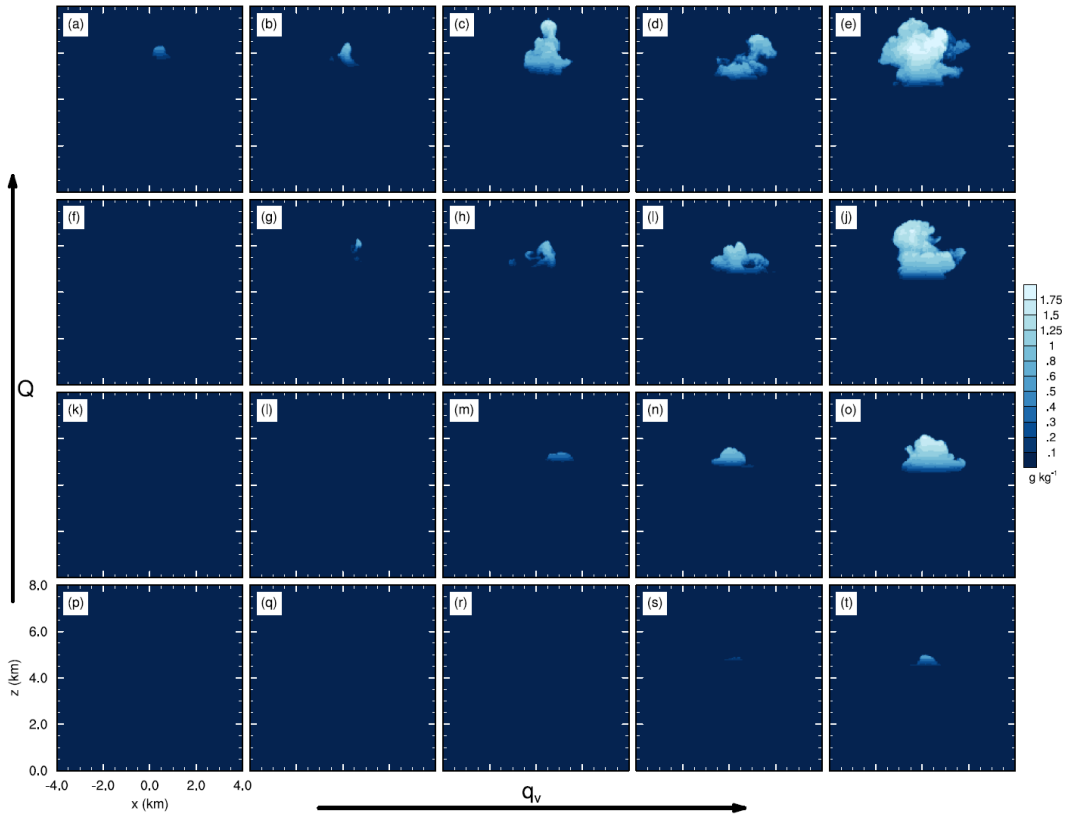

Figure 20. Vertical cross-sections of the maximum instantaneous liquid water mixing ratio in the y–direction, g kg^{-1}, at 22 minutes after the fires have reached full intensity. Panels show all twenty combinations of fire intensity, Q = (5, 10, 20, 30) kW m^{-2}, arranged vertically and mixed-layer water vapour mixing ratio, qvML = (2.0, 2.5, 3.0, 3.5, 4.0) g kg^{-1}, arranged horizontally. Fire intensity increases from the bottom row of panels to the top row of panels and boundary-layer water vapour mixing ratio increases from the left-hand column of panels to the right-hand column of panels, as indicated by the arrows.

"Wet" fires

Snapshots of the liquid water mixing ratio vertical cross-sections for all five Q = 30 kW m^{-2} fire intensity simulations, with latent heat fluxes LE of 0.0, 2.1 and 11.4 kW m^{-2} at 22 minutes after the fires have reached full intensity are shown in Figure 21. The effect of adding a latent heat flux to the fire is very small, slightly increasing cloud amount and cloud top height, for all of the mixed-layer water vapour mixing ratios shown.

To investigate why the addition of a moisture source, in the form of a latent heat flux, to the fire has only a small effect on pyrocumulus formation, we construct time-height diagrams of the maximum water vapour mixing ratio at each model level (Figure 22). In the zero latent heat flux case the boundary-layer water vapour mixing ratio is constant throughout the simulation and the lifting of the boundary layer air by the plume into the relatively dry free troposphere is evident after about 10 minutes. In the L_E = 2.1 kW m^{-2} case, the water vapour mixing ratio is greater near the surface, above which the moisture becomes more diffused with height and is similar to the zero latent heat flux case. Even in the L_E = 11.4 kW m^{-2} case, the moisture released by the fire is only just reaching the top of the boundary layer, where it is then able to have a minimal effect on pyrocumulus formation. The amount of moisture that gets transported upwards in the L_E = 11.4 kW m^{-2} case is made clear by the difference plot of time-height water vapour mixing ratio in Figure 22 (d).

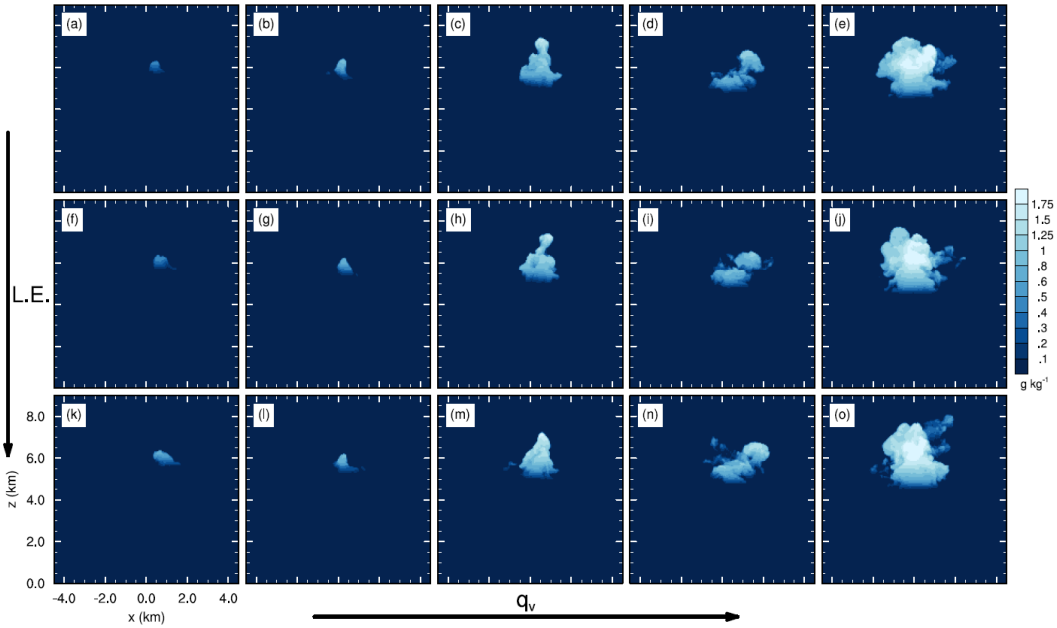

Figure 21. Vertical cross-sections of the maximum instantaneous liquid water mixing ratio in the y−direction, g kg^{-1}, at 22 minutes after the fires have reached full intensity. Panels show all fifteen combinations of fire latent heat flux, LE = (0, 2.1, 11.4) kW m^{-2}, arranged vertically and mixed-layer water vapour mixing ratio, q_{vML} = (2.0, 2.5, 3.0, 3.5, 4.0) g kg^{-1}, arranged horizontally. Fire latent heat flux increases from the bottom row of panels to the top row of panels and boundary-layer water vapour mixing ratio increases from the left-hand column of panels to the right-hand column of panels, as indicated by the arrows. All simulations are performed with a fire intensity Q = 30 kW m^{-2}.

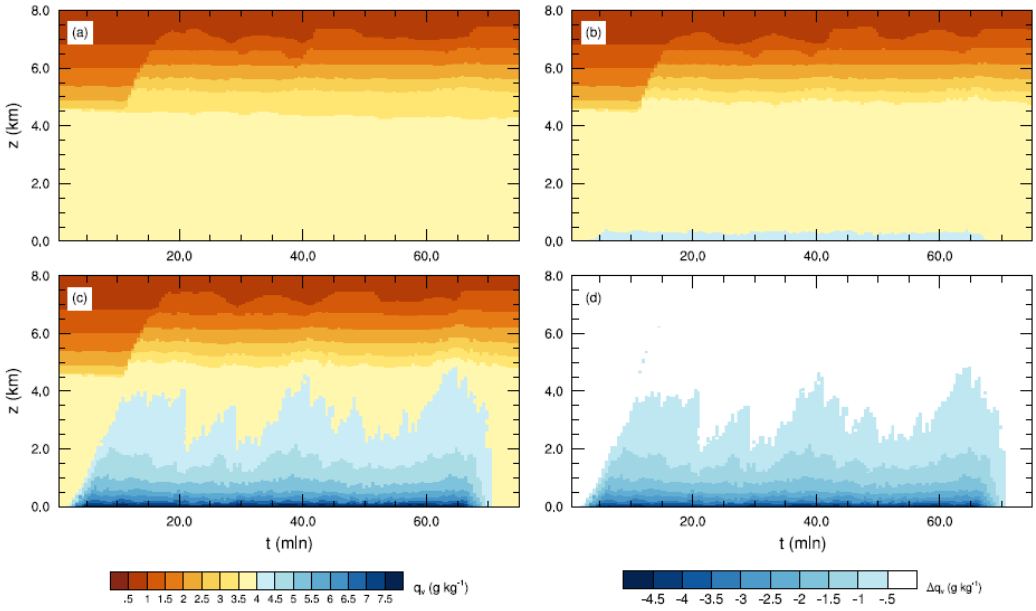

Figure 22. Time-height diagrams of the maximum instantaneous water vapour mixing ratio, g kg^{-1}, for simulations with fire latent heat fluxes, LE, of (a) 0.0 kW m^{-2}, (b) 2.1 kW m^{-2} and (c) 11.4 kW m^{-2}. Panel (d) shows the difference between (a) and (c). All simulations shown are for the Q = 30 kW m^{-2} fire intensity and q_{vML} = 3.5 g kg^{-1} mixed-layer water vapour mixing ratio case.

Discussion and conclusion

Large-eddy simulations of buoyant wildfire plumes have been performed in background atmospheric conditions representative of high fire-danger days. The potential for the formation of pyro-convective cloud has been investigated by varying the environmental moisture and the fire intensity. Moister back-ground atmospheres produce larger, more intense pyro-convective clouds and if the background atmosphere is moist enough, intense fires can produce pyro-convection even in the absence of any moisture from the combustion. Increasing the environmental moisture reduced the cloud-base height, whereas increasing the fire intensity had no discernible effect on cloud-base height.

The relative importance of environmental and fire-derived moisture has been assessed by repeating a set of fixed-intensity runs twice, with latent heat fluxes representative of dry and moist fuels respectively. Pyrogenous moisture was found to only have a small effect on pyrocumulus formation. This is because even for high rates of moisture release, this moisture becomes diluted as it is transported by the plume to the top of the boundary layer.

The resurgence of the plume updraft due to the latent heat released within a pyro-convective cloud has the potential to increase the height to which firebrands are lofted, particularly as the pyrocumulus updrafts are in excess of 10 m s^{-1}, considerably greater than the fall velocity of typical firebrands of 4–6 m s^{-1} (e.g. Ellis, 2010). Lofting the firebrands higher into the atmosphere has the potential to increase the distance they travel before landing, although detailed calculations similar to those described in the previous chapter (see also Thurston et al. 2017) would need to be performed to accurately assess the potential impact of pyro-convection on spotting distance.

Although we are currently unable to forecast the precise location and timing of an individual pyro-convective updraft or downburst, by performing studies such as this we are able to learn about the conditions under which they form, their dynamics and potential impact on fire behaviour.

Practical implications

The ingredients for cloud formation are well known – moist air, and a cooling mechanism, which is usually ascent. In the case of pyrocumulus, there are two possible moisture sources, the atmosphere and the fire, and their relative role has been controversial. The moisture from the fire is difficult to characterise, for it depends on the total heat output from the fire, and the fuel characteristics. Hence, finding that forecasts are insensitive to a poorly known parameter is good news, for it means that the forecast can probably make do with quite rough estimates of fire moisture output. Further good news is that the atmospheric contribution to moisture, particularly on likely pyroCu days, can be reasonably well estimated from surface observations near the fire.

The following sections describe the development of a pyrocumulus forecasting tool. This knowledge became a crucial underpinning of that tool.

THERMODYNAMIC CONSTRAINTS ON PYROCUMULUS FORMATION

Summary

In favourable atmospheric conditions, large hot fires can produce pyrocumulonimbus (pyroCb) cloud in the form of deep convective columns resembling conventional thunderstorms, which may be accompanied by strong inflow, dangerous downbursts and lightning strikes. These in turn may enhance fire spread rates and fire intensity, cause sudden changes in fire spread direction, and the lightning may ignite additional fires. Dangerous pyroCb conditions are not well understood and are very difficult to forecast.

Here, a conceptual study of the thermodynamics of fire plumes is presented to better understand the influence of a range of factors on plume condensation. Recognising that plume gases are undilute at the fire source and approach 100% dilution at the plume top (neutral buoyancy), we consider how the plume condensation height changes for this full range of dilution and for a given set of factors that include: environmental temperature and humidity, fire temperature, and fire moisture to heat ratios. The condensation heights are calculated and plotted as saturation point (SP) curves on thermodynamic diagrams for a broad range of each factor. The distribution of SP curves on thermodynamic diagrams provides useful insight into pyroCb behaviour. Adding plume temperature traces from Large-Eddy Model simulations to the thermodynamic diagrams provides additional insight into plume buoyancy, how it varies with height, and the potential for dangerous pyroCb development.

Introduction

Pyrocumulus (pyroCu) clouds are produced by intense heating of air from fire or volcanic activity that leads to ascent and subsequent condensation when the rising air becomes saturated due to cooling from adiabatic expansion. The process is similar to conventional convective cloud formation, when a lifting mechanism (e.g., orographic lifting, intersection of two air masses) raises air above the level at which cloud forms (the lifting condensation level). Additional lifting and condensational heating may raise the air to the level of free convection, above which it is positively buoyant. Turbulent entrainment of cooler and drier air from outside the rising airmass dilutes the cloud buoyancy, which can limit the size and growth of the cloud (e.g., fair weather cumulus). Larger and more intense lifted regions can accelerate to the tropopause (e.g., cumulonimbus thunderstorms). The main difference between conventional cumulus and cumulonimbus and fire-sourced pyroCu and pyroCb (hereafter referred collectively as pyroCu/Cb) clouds is that the initial lifting in the latter cloud types is provided by the buoyancy from the heat and perhaps moisture released by the fire. In large fires with an intense convection column the cloud may resemble towering cumulonimbus with updrafts that penetrate into the stratosphere. We refer to these plumes as pyroCb. (See Tory et al 2015, 2016 for a review of pyroCu/Cb literature and forecast techniques respectively).

There is abundant anecdotal evidence to suggest that the presence of pyroCb activity can have a significant impact on fire behaviour, including: (i) the amplification of burn- and spread-rates (Fromm et al. 2006, Trentmann et al. 2006, Rosenfeld et al. 2007, Fromm et al. 2012), (ii) enhanced spotting due to larger,

taller and more intense plumes (e.g., Koo et al. 2010), and (iii) ignition of new fires by pyroCb lightning strikes due to pyroCb conditions favouring hotter and longer-lived lightning strikes (e.g., Rudlosky and Fuelberg 2011, Peace et al. 2017).

Given the potential threat posed by pyroCb there is great interest in being able to predict its development. Unfortunately, pyroCb are very difficult to forecast. Current forecast techniques draw on similarities between pyroCb and conventional thunderstorms, and the recognition that conditions that favour thunderstorm development will also favour pyroCb development. Ideal pyroCb conditions are thus similar to ideal thunderstorm conditions but with a dry rather than moist lower troposphere. These conditions appear on a thermodynamic diagram as the classic inverted-V profile, in which a dry adiabatic temperature profile of constant potential temperature (θ_{env}) forms the right side of the inverted-V, while the constant specific humidity (q_{env}) moisture profile makes up the left side.

In this paper we construct an idealised theoretical plume model in an inverted-V environment to aid our understanding of how the environment and fire properties influence pyroCu/Cb formation.

Methods: Plume model

The fire plume is a mixture of hot combustion gases and entrained air from the immediate environment. For simplicity, it is assumed that the plumes are well-mixed and develop in a well-mixed atmospheric boundary layer of constant potential temperature (θ) and constant specific humidity (q), and that the plume condensation level (CL, which is the saturation point on a thermodynamic diagram) occurs within this well-mixed layer. While this assumption is likely to be unrealistic for deeper CLs, it has no impact on the conclusions. For simplicity the thermodynamics of the plume above the CL are not considered. However, useful information on plume thermodynamics can be determined from the plume θ and q at the CL, and diagnostic quantities derived from these variables.

The plume θ and q are expressed as functions of the environment θ and q and the fire θ and q, and the plume dilution fraction α,

$$\theta_{pl} = \alpha\theta_{env} + (1-\alpha)\theta_{fire} \quad (1)$$

$$q_{pl} = \alpha q_{env} + (1-\alpha)q_{fire}. \quad (2)$$

θ_{env} and q_{env} are specified for each scenario, and α is varied from 0 to 1, representing a range of plume dilution rates from pure combustion gases to pure environmental air (i.e., infinitesimal quantities of combustion gases). These parameters and other variables introduced below are listed and described in Table 2.

Table 2: Plume model variables and constants

θ	Potential temperature (units K)
θ_{env}	Constant environment potential temperature (up to the condensation level)
θ_{fire}	Potential temperature of the fire (flames)
θ_{pl}	Plume potential temperature
$\Delta\theta_f$	Plume potential temperature contribution from combustion
q	Specific humidity (units $kg\ kg^{-1}$, mass of water vapour to total mass of air)
q_{env}	Constant environment specific humidity (up to the condensation level)
q_{fire}	Fire specific humidity (includes moisture from the air consumed in combustion)
q_{pl}	Plume specific humidity
Δq_f	Plume moisture contribution from combustion (evaporation of fuel moisture and moisture as a product of combustion)
α	Plume dilution factor. Ranges from 1 (100% dilute = environment value) to 0 (pure combustion gas)
β	Plume buoyancy factor. Ranges from 0 (plume 100% dilute) to $\gamma - 1$. (Useful range $0 \rightarrow \sim 10^{-1}$.)
γ	Fire heat multiplication factor to express θ_{fire} as a multiple of θ_{env}.
φ	Fire moisture to heat ratio (units $kg\ kg^{-1}\ K^{-1}$).

The potential temperature of combustion, θ_{fire}, can be expressed as a multiplier (γ) of the environment potential temperature, $\theta_{fire} = \gamma\theta_{env}$. Assuming $\theta_{env} \sim 300\ K$ γ ranges from $2 \rightarrow 5$ representing flame temperature estimates from forest fires of 600 K (flame tips) to 1500 K (flame base). The heat produced by the fire is, $\Delta\theta_f = (\gamma - 1)\theta_{env}$. The fire moisture to heat ratio, φ, is specified, which means fire moisture, $\Delta q_f = \varphi\Delta\theta_f$. Taking into account a 6 to 1 air to fuel ratio in combustion, the combustion gas includes an additional moisture contribution from the environment, $q_{fire} = \Delta q_f + 0.86 q_{env}$.

An iterative process is used to calculate the CL based on estimates of the CL pressure (P_{CL}). The CL temperature is calculated from θ_{pl} and P_{CL}, which are used to calculate the saturation vapour pressure at P_{CL}. If this saturation vapour pressure is less than (greater than) the actual vapour pressure the plume must be saturated (unsaturated) at P_{CL} and the process is repeated at a lower (higher) level until the plume CL is approached to the nearest 1 hPa.

A number of diagnostic equations have been developed to illustrate plume characteristics. Each can be expressed as a function of a buoyancy parameter,

$$\beta = (\theta_{pl} - \theta_{env})/\theta_{env} = (1-\alpha)(\gamma - 1), \qquad 3.$$

which reduces the experimental parameter space by replacing α and γ with β.

Results

Two well-mixed boundary layer profiles are considered, one warm (Fig. 1a) and the other cold (Fig. 1b,). The first has $\theta_{env} = 303\ K = 30\ C$, $q_{env} = 5 \times 10^{-3} kg\ kg^{-1}$, 19% relative humidity, and an elevated lifting condensation level (LCL) about 3 km above the surface. The second has $\theta_{env} = 271\ K = -2\ C$, $q_{env} = 2 \times 10^{-3} kg\ kg^{-1}$, 61% relative humidity, and a relatively low LCL. This is representative of the Flatanger fire in Norway, January 2014, which destroyed 140 houses. The LCL is located at the apex of the θ_{env} and q_{env} curves. For clarity we refer to this LCL as the environment LCL (ELCL) to distinguish it from any in-plume LCLs.

Saturation point curves

Figure 23 includes SP curves for the hottest fire ($\gamma = 5$) in the warm environment (Figure 23a) and the coolest fire ($\gamma = 2$) in the cold environment (Figure 23b), and each with the two values of fire moisture to heat ratios that represent Luderer et al.'s (2009) driest ($\varphi = 3 \times 10^{-5} kg\ kg^{-1} K^{-1}$, red) and wettest ($\varphi = 15 \times 10^{-5} kg\ kg^{-1} K^{-1}$, blue) realistic fires. Each SP curve represents the position of the plume condensation level corresponding to α varying from 1 (100% dilution, lower left) to 0 (pure combustion gas, upper right) for the specified environment conditions, fire moisture to heat ratio, and fire heat. For example, the dots on the SP curves represent the apex of plume temperature and moisture traces for a plume consisting of 95% environment air and 5% combustion gas.

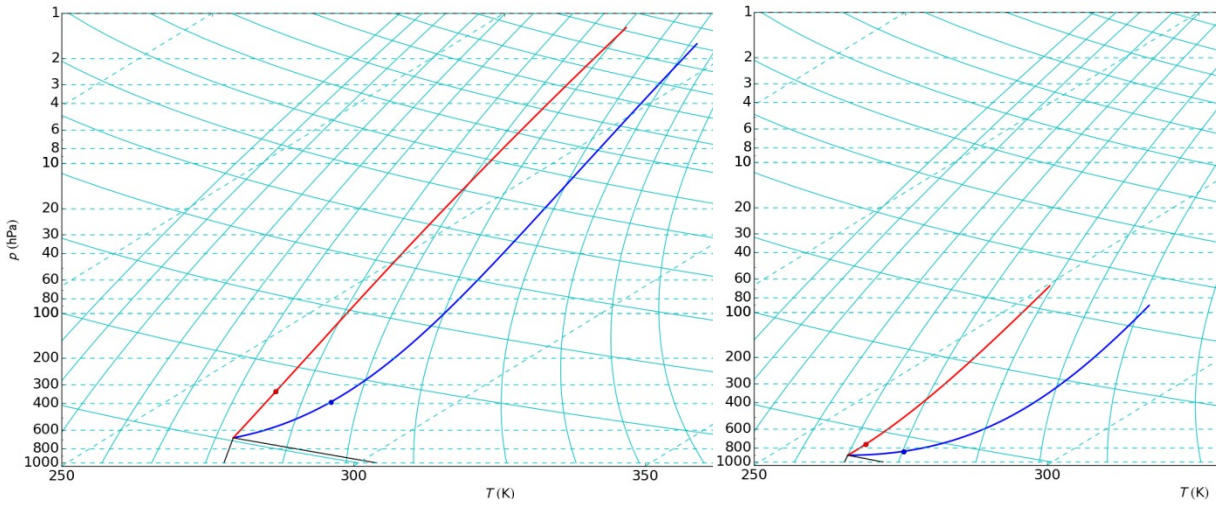

Figure 23: Saturation point curves for the dry (red) and moist (blue) fires, for the two cases (a) hot fire in a warm environment, and (b) cool fire in a cold environment, on a skewT-logp diagram. The 95% dilution points are indicated by dots. The environment LCL is located at the apex of the grey lines of constant θ_{env} and q_{env}. The pale blue lines are lines of constant pressure (dashed, horizontal), temperature (dashed, diagonal), potential temperature (solid, shallow gradient) and specific humidity (solid, steep gradient).

Of the parameter space investigated, the two most extreme cases are included in Fig. 1: hottest and driest fire in the warm environment (red curve in Figure 23a), and coolest and moistest fire in the cold environment (blue curve in Figure 23b).

100% dilution coincides with the ELCL, and zero dilution (at the upper right end of the coloured curves) shows exceptionally high CLs. These are ~1.5 hPa (> 40 km, Fig. 1a) and ~90 hPa (> 20 km, Fig. 1b).

Figure 23 shows that for pyroCu/Cb to form in the lower troposphere (e.g., below 500 hPa) significant dilution is necessary (e.g., > 95 % dilution for the hot fire, Figure 23a). Furthermore, Large-Eddy Model simulations (LEM, Thurston et al. 2016 and the previous chapter) suggest typical dilution amounts in the condensing plume are likely to be closer to 99%. It follows that substantial amounts of dilution must occur in typical pyroCu/Cb plumes, which corrects a common misconception that relatively undiluted plumes are necessary for pyroCu/Cb to form.

From Eq. 3 it can be seen that a plume from a hot fire (γ) with moderate dilution (α) could have the same buoyancy (β) as a plume from cooler fire with less dilution. This overlapping parameter space produces overlapping SP curves (e.g., the $\gamma = 2$ fire SP curve is almost exactly overlaid by the $\alpha = 1 \rightarrow 0.75$ part of the $\gamma = 5$ fire, not shown). Hereafter, we discuss plume buoyancy represented by the parameter β instead of fire heat and plume dilution. We conclude from Fig. 1 that for pyroCu/Cb to form in the lower troposphere, $\beta \leq O(10^{-1})$. Typical condensation level values of β in the LEM simulations are one to two orders of magnitude smaller than this.

Because β is small at condensation, the origin and gradient of the SP curves provide useful information about the height at which a plume will condense, and thus the potential for pyroCu/Cb formation. The origin ($\beta = 0$) coincides with the ELCL, which provides a first order estimate of the plume condensation level. The difference between the actual plume condensation height and the ELCL is of second order importance (for small β) which can be estimated from the product of β and the SP curve gradient. Thus for a given β, steep SP curves (e.g., dry fires) correspond to greater condensation heights than flatter (e.g., moister fires) SP curves.

Fortunately, the most important factor for estimating plume condensation heights (the ELCL) does not require any information about the fire. The secondary factor is dependent on the environment plus the fire moisture to heat ratio (φ) and the plume buoyancy (β). To determine β at the plume condensation height, we expect detailed knowledge would be required of how plume buoyancy is affected by fire size, distribution and intensity, and how the atmosphere (e.g., wind and thermodynamic stability) affect the entrainment rate (plume dilution), and thus the distribution of β with height. One might also expect detailed knowledge of the fire and fuels would be required to determine φ. However, within the assumptions of the simple theoretical plume model (i.e., plume moisture and heat are diluted at the same rate), φ remains constant throughout the plume (and is independent of β), which means it can be estimated from a single plume measurement anywhere in the plume and at any time,

$$\varphi \cong \frac{q_{pl} - q_{env}}{\theta_{pl} - \theta_{env}}. \qquad 4.$$

In reality φ is likely to vary with time (and perhaps spatially), as the fire burns through a variety of fuels, but as long as the measurement is taken above the flaming zone (so that additional radiative heat losses are relatively small) the measured plume element should maintain a constant φ throughout its ascent

through a well-mixed boundary layer (constant θ_{env} and q_{env}). Thus, in practice it should be possible to produce thermodynamic diagrams with the ELCL and SP curves plotted, similar to Figure 23, from observations at the fire ground, provided a representative sample of q_{pl} and θ_{pl} measurements can be made.

Plume temperature traces

The SP curves provide insight into the height at which a plume or plume element might condense for a given environment (θ_{env}, q_{env}), fire moisture to heat ratio (φ)and plume buoyancy (β), but they do not tell us anything about specific plumes. In Figure 24 temperature and moisture traces from two LEM plume simulations (taken from Thurston et al. 2016) have been plotted on thermodynamic diagrams with SP curves included. An extra SP curve has been added representing the fire moisture to heat ratio from one of Potter's (2005) fireCAPE thought experiments (green curve), which Luderer et al. (2009) argued was unrealistically moist.

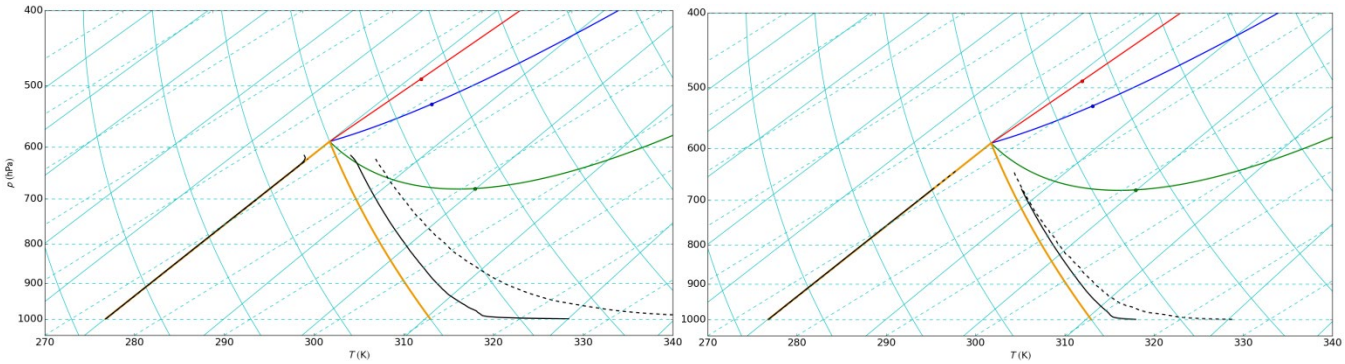

Figure 24. As in Fig. 1 but with mean (solid) and maximum (dashed) plume temperature and moisture traces in a hot and dry ($\theta_{env} = 310$ and $q_{env} = 4 \times 10^{-3}$), zero wind environment from LEM simulations with a constant circular surface heat flux (Q) of 250 m radius. Saturation point curves for the dry (red, $\varphi = 3 \times 10^{-5}$) moist (blue, $\varphi = 15 \times 10^{-5}$) fires and an extremely moist fire (green, $\varphi = 100 \times 10^{-5}$) are included. (a) hot fire ($Q = 30\ kWm^{-2}$) (b) cool fire ($Q = 5\ kWm^{-2}$).

As the LEM plume air ascends and approaches the ELCL it begins to entrain warm and dry environment air from above the boundary layer, which is a process that cannot be incorporated in our theoretical model. Thus, we discount the plume traces higher than about 620 hPa, and instead mentally extrapolate them to the SP curves. Additionally, in order to make a clear distinction between plume and environment air, only plume elements that are at least 1 K warmer than the environment were included in the plume-average temperature trace. In reality the plumes are expected to contain a mix of air parcels of temperature varying from the environment temperature (recently entrained parcels) up to the maximum temperature indicated by the dotted lines (least diluted air parcels).

For the hot fire (Figure 24a) the mix of plume element temperatures would be expected to have a range of condensation heights extending from the ELCL to where the extrapolated dashed line meets the SP curve corresponding to the fire's moisture to heat ratio. The corresponding buoyancies range from $\beta = 0 \rightarrow 0.14$, with the plume average, $\beta = 0.04$. This simulation produced deep pyroCb with rain and evaporatively cooled downdrafts. Whereas the plume dilution was generally too great for condensation to occur in the cool fire simulation (Fig. 2b)

as it mostly lost buoyancy near 650 hPa before intersecting any of the SP curves. A few parcels of buoyant air did occasionally reach the condensation level, producing short-lived puffs of shallow cloud.

Discussion: What can we learn from these diagrams?

There is a surprising amount of information about conditions that support pyroCu/Cb development and plume behavior that can be gleaned from Figure 23 and Figure 24.

- The positive gradients of the SP curves corresponding to the realistic range of fire moisture to heat ratios (red and blue) demonstrate that buoyant plume elements condense at levels higher than the ELCL (consistent with Luderer et al. 2009).
- An exception is for very moist fires (e.g., green SP curves in Figure 24) and/or very dry environments (e.g., Figure 23b) where the SP curve may have a negative gradient, in which case some buoyant elements might condense at levels lower than the ELCL (e.g., the Potter 2005 fireCAPE thought experiment). This is more likely to occur in cold and dry (small q) environments.
- Buoyant elements from moister fires will condense at lower levels than for drier fires.
- There is a broad range of temperatures and hence buoyancy within plumes, that decrease with height (due to dilution from entrainment, Figure 24).
- Plumes with non-trivial buoyancy near their condensation level (e.g., Figure 24a), contain plume elements with a range of buoyancy from zero to a maximum value corresponding to the least dilute plume element, with a corresponding range of condensation heights. These condensation heights are determined by the intersection of the plume element temperature trace and the relevant SP curve.
- Plumes that produce pyroCb (e.g., Figure 24a) have non-trivial buoyancy near the condensation level, suggesting the fireCAPE concept may be useful for pyroCb forecasting.
- The same LEM heat sources in environments with lower ELCLs (e.g., that might occur with the passage of a cold front or sea breeze) might produce very much more energetic pyroCb. At 900 hPa the hot fire mean and maximum plume values of β are about three times greater than at 600 hPa, and the cool fire mean and maximum plume values of β at 900 hPa are very similar to the hot fire at 600 hPa.

More insight will be described in a journal article (in preparation), based on a mathematical exploration of the model parameters. This paper describes among other things how the environment affects the SP curves, what values of plume buoyancy are important for pyroCu/Cb activity, and the sensitivity of fireCAPE to plume buoyancy.

Conclusions

PyroCb can produce dangerous fire behaviour, through changes in fire rate of spread and direction, increased spotting, and additional ignitions from lightning strikes. Unfortunately, pyroCb are difficult to predict and not well understood.

In this paper we have introduced a simple theoretical model that provides useful insight into the conditions that favour pyroCu/Cb formation, the thermodynamic composition of fire plumes, and the relative sensitivity of environment conditions to fire properties that determine pyroCu/Cb formation and behaviour. Some of the more general results are summarised here:

- A substantial amount of dilution is required for pyroCu/Cb cloud elements to condense. However, too much dilution and the plume may lose buoyancy before ascending high enough for condensation to occur.
- The environment lifting condensation level (ELCL) provides a good first order estimate of the plume condensation height.
- Typical forest fires that produce pyroCu/Cb will have buoyant plume elements that condense at elevations higher than the ELCL, because the additional heat provided by the fire contributes to raising the condensation level more than the additional moisture contributes to lowering the condensation level.
- PyroCu/Cb formation and behaviour is relatively insensitive to fire heat and moisture, but could be very sensitive to environment changes, such as might be experienced with the arrival of a cold front or sea-breeze that lowers substantially the ELCL due to the arrival of cooler and moister air.

Practical implications

This work confirms the potentially controversial finding of the previous section, of the small role that fire moisture plays in pyrocumulus formation for typical forest fire conditions. It thus supports the simplifying assumptions that we use in the next section to develop a pyrocumulus forecasting tool. The elegant simplicity of the modelling in this component will also help inform our development of that tool and ensure that it has a rigorous scientific foundation.

PYROCUMULUS PREDICTION: METHOD

This section describes the development of the first version of a pyroCb prediction tool. It was tested on more than 30 events, including both pyroCb and non-pyroCb producing fires. These tests identified substantial variation in pyroCb threat from day to day. They also identified that conditions highly favourable for pyroCb formation do not favour wildfires and vice-versa. Consequently, a method to combine the potential for intense fire development with the potential for pyroCb formation was identified to be necessary for pyroCb prediction. This report reviews the development of the original pyroCb prediction tool (the Pyrocumulonimbus Fire Power Threshold, PFT) and documents the development of the combined PFT and fire potential diagnostic, termed the PFT-flag. The next section describes testing of these tools in a real time trial over the summer of 2019-20.

Introduction

In favourable atmospheric conditions, suitably large and hot fires can produce pyroCb cloud in the form of deep convective columns with many similarities to conventional thunderstorms. They may be accompanied by strong inflow, dangerous downbursts and lightning strikes, which may enhance fire spread rates and fire intensity, cause sudden changes in fire spread direction, and the lightning may ignite additional fires. Dangerous pyroCb conditions are not well understood and can be very difficult to forecast.

In recent CRC research, a method for determining how favourable the atmospheric environment is for pyroCb development was developed. This method is combined with a plume-rise model (originally developed for pollutant dispersion prediction) to determine how much heat a fire must produce for pyroCb to develop in a given atmospheric environment. More specifically, this fire heat is the rate at which heat enters the fire plume (which has units of power), often termed the "power of the fire" or "firepower". A theoretical minimum firepower required for pyroCb to develop in a given atmospheric environment is calculated, termed the Pyrocumulonimbus Firepower Threshold (PFT).

Forecast spatial plots of PFT are being trialled that provide an indication of how the favourability of the atmosphere for pyroCb development varies in space and time over typical weather forecast periods. It is anticipated that such plots will provide useful guidance for fire weather forecasters and fire agencies. Preliminary studies have shown that the PFT can vary substantially from day to day, and that days that favour pyroCb formation do not necessarily favour large-hot fires. A PFT-flag is also under development that identifies when both pyroCb and large-hot fires are favourable.

Background – Previous CRC research

In previous CRC research an idealised theoretical plume model was introduced (Tory et al. 2018, Tory and Kepert 2018 – see also the previous section titled "Thermodynamic constraints on pyrocumulus formation") that can identify at what temperature and pressure (or height) condensation will begin to form in a fire plume. These condensation properties vary according to how much warmer

the plume is than the environment (i.e., how buoyant the plume is). Plotted on a thermodynamic diagram, the condensation properties are represented by a single point, termed the Saturation Point (SP), and the SPs for a range of plume buoyancies form a SP curve (Figure 25, solid blue curve). Each SP curve is unique to an assumed mixed-layer environment comprising constant mixed-layer potential temperature (θ_{ML}, Figure 25, thick red line) and specific humidity (q_{ML}, Figure 25, thick pale-blue line), and an assumed fire moisture to heat production ratio (φ). Fortunately, the SP curves are not sensitive to a range of realistic values of φ, and neither is the PFT.

The SP curve defines where a hypothetical ascending plume parcel begins to follow a moist adiabat on a thermodynamic diagram. For pyroCb formation any moist ascending parcel needs to remain buoyant (warmer than the environment, rightmost thin red line in Figure 25) until it reaches some designated height at which pyroCb is deemed to have been achieved. Here, that height is the so-called electrification level, -20 ℃ (Figure 25, pale-blue dashed line). The coolest moist adiabat that satisfies these criteria represents the coolest possible pyroCb plume-element pathway, and thus defines the pyroCb moist-adiabat limit (Figure 25, yellow curve). Where this moist adiabat and the SP curve meet is the free-convection height limit. Any buoyant plume element that reaches or exceeds this height will, in theory, freely convect to the electrification level. The intersection of these two curves defines the free-convection height (z_{fc}, Figure 25, fine blue dashed line), which is one of the key inputs to the PFT (see below). Another key PFT input is a measure of the plume-element buoyancy at this height. Specifically, it is the potential temperature difference ($\Delta\theta$) between the plume element at z_{fc} ($\theta_{pl,fc}$) and θ_{ML}, which can easily be read off the thermodynamic diagram (Figure 25). $\Delta\theta$ is a proxy for the plume buoyancy, and represents the buoyancy required for the plume to overcome any stable layers (inversions) that might inhibit the plume from reaching the electrification level. The third key PFT input is the average mixed-layer windspeed, which can be determined directly from the wind data available or estimated from the winds barbs on the edge of a thermodynamic diagram (Figure 25, highlighted in the green ellipse).

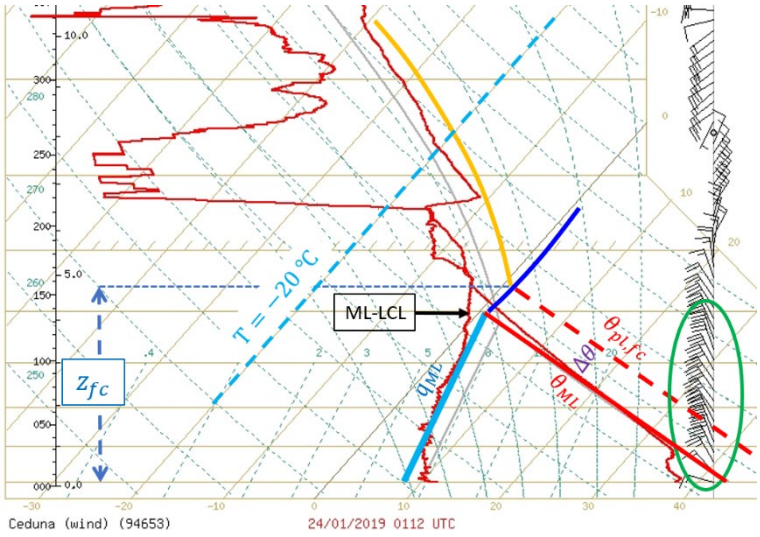

Figure 25: A sample atmospheric sounding (thin red lines and wind barbs) applied to an F-160 thermodynamic diagram, with quantities required for the PFT calculation overlaid: Mixed Layer potential temperature (θ_{ML}, thick red line); Mixed-Layer specific humidity (q_{ML}, thick pale-blue line); Mixed-layer lifting condensation level (ML-LCL, apex of the θ_{ML} and q_{ML} lines); Saturation Point curve (SP curve, blue curve emanating from the ML-LCL); Free-convection moist adiabat (yellow curve); Electrification level ($T = -20$°C, pale-blue dashed line); Free-convection height (z_{fc}, blue dotted line corresponding to the intersection of the SP curve and free-convection moist adiabat); Free-convection plume potential temperature ($\theta_{pl,fc}$, red dashed line); Plume excess potential temperature ($\Delta\theta$, difference between $\theta_{pl,fc}$ and θ_{ML}); and the winds used to calculate the mixed-layer wind speed (U, green ellipse).

PFT equations

The PFT is derived from an equation that describes the buoyancy flux distribution along the plume centerline for a Briggs plume (Briggs 1975, 1984) in a constant horizontal wind crossflow (U) and neutrally stable environment (Tory and Kepert 2018, Tory 2018). A schematic representation of a Briggs plume is shown in Figure 26. The plume geometry is described by two equations, an equation for the plume centreline height (Figure 26, yellow line) with downwind distance (x),

$$z_c = \left[\left(\frac{3}{2\beta^2}\right)\frac{B_{flux}}{\pi}\right]^{\frac{1}{3}}\frac{x^{\frac{2}{3}}}{U}, \qquad 1.$$

and an equation that describes an upright circular plume cross-section,

$$R = \beta z_c. \qquad 2.$$

Here B_{flux} is the buoyancy flux at the plume source, which is proportional to the heatflux or firepower entering the plume. β is a constant entrainment coefficient and π is the circle constant. Eq. 2 describes the radius of the dynamic plume (pale blue lines, Figure 26), which includes the plume gases (internal plume) plus the surrounding environment lifted by the rising plume. The internal plume (dark blue lines, Figure 26) radius is smaller and is given by,

$$R' = \beta' z_c. \qquad 3.$$

Both entrainment parameters (β and β') have been measured in numerous observational and laboratory studies yielding $\beta = 0.6$ and $\beta' = 0.4$ (e.g., Briggs

1975, 1984). The PFT equation is based on buoyancy flux conservation within the internal plume, and thus β' is the appropriate entrainment parameter in the following equations.

Assuming only some fraction of the plume area (α) needs to reach the free-convection height, represented by a fraction of the plume radius, α' (Figure 26, red arrow), then the plume centreline expressed as a function of the free-convection height becomes,

$$z_c = \frac{z_{fc}}{1+\alpha'\beta'}, \qquad 4.$$

and the PyroCb Firepower Threshold becomes,

$$PFT = \left\{\frac{\pi C_p}{R_d}\left[\frac{\beta'}{1+\alpha'\beta'}\right]^2\right\}\left(\frac{P_c}{\theta_{pl,fc}}\right) z_{fc}{}^2 U \Delta\theta_{fc}. \qquad 5.$$

Here C_p and R_d are the Specific heat at constant pressure and gas constant for dry air. P_c is the pressure at the plume centreline corresponding to the free-convection height, and $\theta_{pl,fc}$ is the plume potential temperature at the free-convection height (by definition $\theta_{pl,fc} = \theta_{ML} + \Delta\theta$). Equation 5 is a slightly more accurate PFT formulation than that provided in Tory and Kepert (2018). (Note, this work is still under development, and further tuning and changes are likely in the future.)

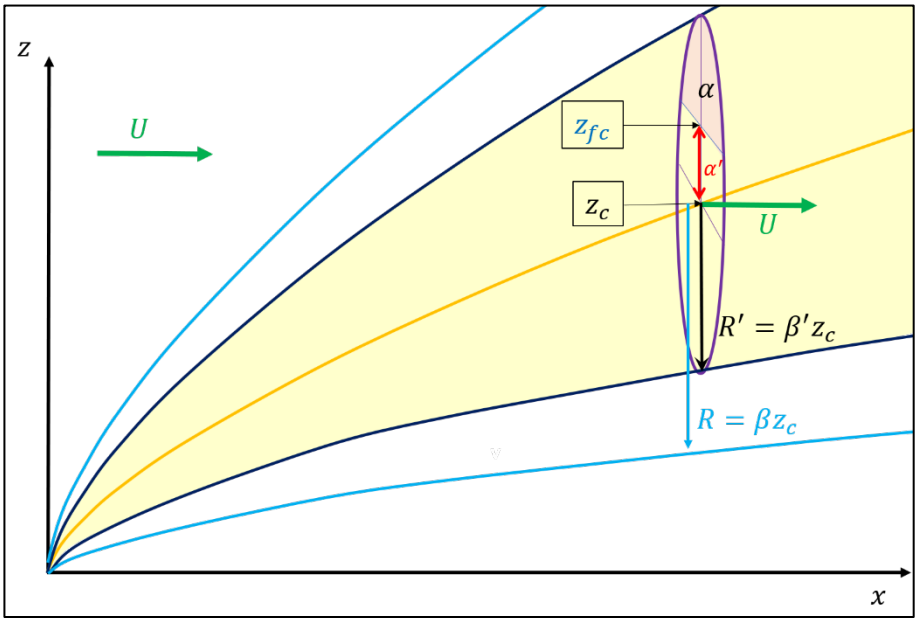

Figure 26: Schematic representation of a Briggs plume, bent-over in the downwind direction (x) by a constant crossflow (U). The internal plume (yellow shading) has a vertical circular cross-section about the plume centreline height (z_c, yellow line) of radius R', which is linearly proportional to the plume centreline height via the constant internal plume entrainment parameter (β'). Similarly, the dynamic plume (inside the pale blue lines) is defined by the dynamic plume radius and entrainment rate, R and β. Only a fraction of the plume area (α) is required to reach the free-convection height (z_{fc}), which can be expressed as a fraction of R' (α').

PFT equation insights

Since the entrainment coefficients (β and β') are constant, the plume centreline height (Eq. 1) at any given downstream distance x is a function of only two

variables, the buoyancy flux (B_{flux}), and the background cross-flow wind (U). Considering these variables independently, it is clear that the plume centreline height is much more sensitive to changes in the wind than the firepower. For example, to double the plume height at a given x would require the windspeed to be halved, whereas it would require the firepower to be increased eight times. This result suggests that observed temporal changes in plume height and slope are more likely to be associated with variations in windspeed than fuel type and/or fuel loads.

The curly bracket term in Eq. 5 is constant, and the curved bracket term typically varies by 10% or less, which means the majority of PFT variation comes from the remaining terms: the free convection height, the background wind speed and the plume excess potential temperature (the buoyancy proxy). The relationship between the PFT and each of these three terms makes sense intuitively. The larger the free convection height, the more firepower required for the plume to rise to that height. The greater the windspeed, the more firepower is required to counter the tendency for the plume to be tilted over by the wind. The greater the buoyancy required for the plume to penetrate stable layers or inversions above z_{fc}, the more firepower that is required. The real insight provided by Eq. 5, however, is the relative power of each term and how they combine to determine an overall pyroCb formation threat. For example, if z_{fc} decreases while $\Delta\theta_{fc}$ increases, perhaps with the passage of a cold front or sea breeze, Eq. 5 will determine if the net effect is more or less favourable for pyroCb formation.

Results

PFT spatial distributions

Spatial plots of PFT generated from computer forecast models can provide valuable insight into how the pyroCb threat varies in space and time (Figure 27). The PFT colour scale in Figure 27 is logarithmic in order to capture the substantial PFT variability at any given time across the landscape. Both cases presented show increased threat near a wind change. This is a very common result. Ahead of the wind change the winds are often very hot and dry (contributing to large z_{fc}) and the windspeeds are often very high (large U), which together contribute to large PFT (Eq. 5) corresponding to highly unfavourable conditions for pyroCb formation. If the wind change brings cooler and moister air to the fire ground (reduced z_{fc}), and a short-term (or longer) lull in the winds (reduced U), conditions become much more favourable for pyroCb formation. Furthermore, a shift in wind direction can lead to increased firepower as long flank fires become head fires.

The two PFT plots in Figure 27 also demonstrate very large differences in PFT in the vicinity of the fires at the time the pyroCb were observed (about 300 GW for Sir Ivan and 10 GW approaching Licola). This suggests pyroCb formation conditions are much less favourable for the Sir Ivan fire than the Licola fire, since 30 times the firepower would be required in the former than the latter. Indeed, conditions that highly favour pyroCb plume development (high humidity and light winds) do not favour large, hot fires, and vice-versa.

Fire-weather forecasters and fire-behaviour analysts using the PFT diagnostic would require some knowledge of the size and intensity of any going fires to be able to assess pyroCb formation potential. This knowledge is rarely available and will never exist for fires that have yet to be ignited. Without such knowledge they can only identify a relative threat, and the PFT forecast performance cannot be verified.

It may be many years before sufficiently reliable observations of firepower become available to enable a rigorous verification program to be undertaken. In the short term, a dataset of past events could be constructed, to identify fire-types that will produce pyroCb for a specific PFT forecast (e.g., a Sir Ivan-scale fire is required for pyroCb to form when the PFT = 300 GW). However, this approach is not ideal because it assumes the yet to be verified PFT is stable and performs consistently across the full breadth of fire weather conditions.

An unverifiable forecast tool such as the PFT has limited prediction value. It must be combined with other information to have value. Ideally, the pyroCb prediction tool would allow forecasters to know that a fire burning in a specific location at a specific time will or will not produce pyroCb.

Returning to Figure 27 one clear difference between the Sir Ivan and Licola fires is the fire-danger conditions. Sir Ivan had catastrophic fire conditions, which would support a much hotter fire than the very-high fire danger conditions present in Licola. Combining some measure of the fire danger conditions with the PFT could produce a verifiable prediction tool.

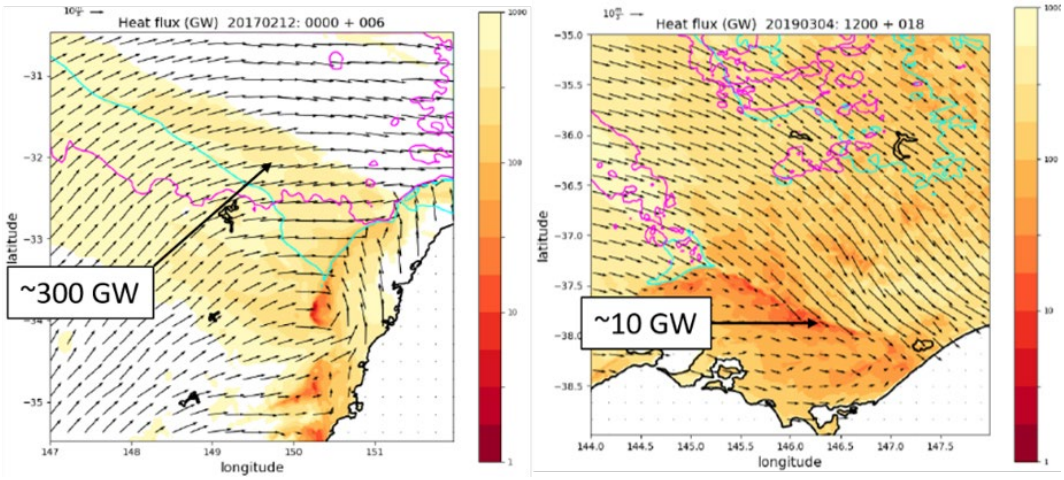

Figure 27: PFT forecasts for two pyroCb events, Sir Ivan (5 PM, 12 February 2017, left) and Licola (5 PM, 5 March 2019, right). The PFT scale is logarithmic (units GW). The wind barbs represent the mixed-layer wind velocity. The PFT label for Sir Ivan points to the fire site, whereas for Licola it points to a minimum value on the wind change that is about to impact Licola.

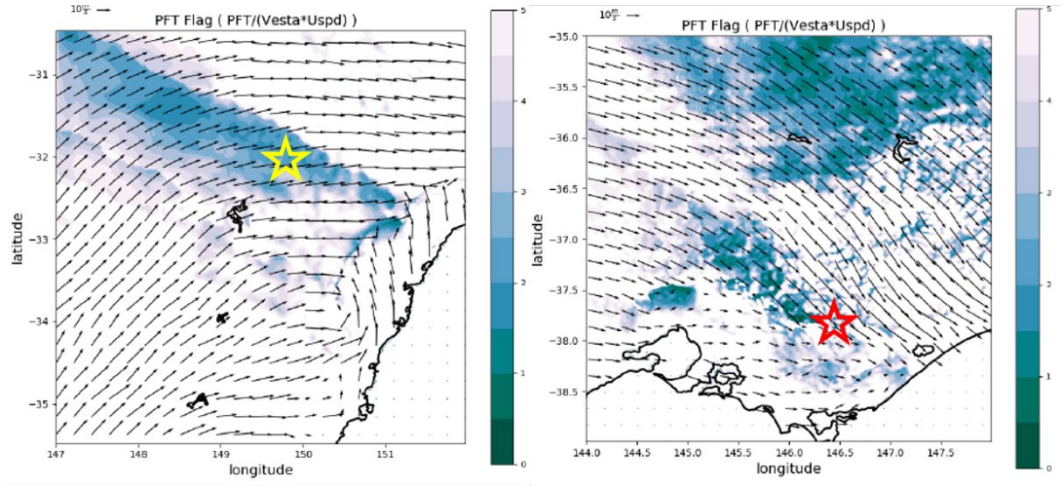

Figure 28: Same as Figure 27 but the PFT-flag forecasts. The stars indicate the fire locations.

PFT-flag formulation

The PFT-flag is designed to represent a ratio of pyroCb plume-development potential (PFT) to fire-intensity potential, to identify when a favourable combination of plume and fire potential is present. For simplicity, the PFT-flag uses only atmospheric variables readily available in Numerical Weather Prediction (NWP) models. The atmospheric components of a variety of fire-weather indices were tested, with the best performing being the product of the Project Vesta (Cheney et al. 2012) and the near-surface windspeed. The main test was for the PFT-flag to produce a similar value for the two extreme cases introduced above, Sir Ivan and Licola. Despite this very rudimentary testing regime the only changes required, after application to more than twenty cases, were the addition of low-wind and low fire-danger limits.

In order to separate the atmospheric component from the fuel components in the Vesta function, two assumptions need to be made. The near-surface windspeed is greater than 5 $km\ hr^{-1}$ (1.39 $m\ s^{-1}$), and a constant term in the rate-of-spread equation is small compared to the windspeed/fuel term and can be ignored. The latter assumption is good for moderate to high fuel loads. Such fuel loads may be necessary to support deep flaming, observed in pyroCb producing fires (McRae and Sharples 2014). The resulting Vesta atmosphere-only equation can be expressed as,

$$V = 18.35{V_M}^{-1.495}.V_U, \qquad 6.$$

where,

$$V_M = \begin{cases} 2.76 + 0.124RH - 0.0187T, & \text{Period 1} \\ 3.60 + 0.169RH - 0.0450T, & \text{Period 2} \\ 3.08 + 0.198RH - 0.0483T, & \text{Period 3} \end{cases} \qquad 7.$$

$$V_U = 1.531(\max(U_{10}, 3.0) - 1.39)^{0.858}. \qquad 8.$$

Here V_M and V_U are the fuel moisture and wind speed contributions to the Vesta function (V). The wind speed term uses the near surface or 10 m wind (U_{10}, units $m\ s^{-1}$). The fuel moisture term is a function of relative humidity (RH, units %) and air temperature (T, units °C), which varies with time of day and time of year, expressed as three distinct periods (Eq. 7). Period 1 extends from midday to 5 PM from October to March. Period 2 is used otherwise for daylight hours, and Period 3 for night hours. Note, these periods are valid for southern hemisphere low- to mid-latitude regions. Application to the northern hemisphere, high latitudes requires additional consideration about when best to apply Period 1. (Experiments in the mid-summer Arctic, suggest Period 1 should apply continuously.)

For comparison with the better-known McArthur Forest Fire Danger Index (FFDI, McArthur 1967, Noble et al. 1980), $V = 15$ is roughly equivalent to an FFDI of 50, and $V = 30$ is similar to an FFDI of 100.

The PFT-flag is given by,

$$PFTflag = \frac{PFT}{V.U_{10}}. \qquad 9.$$

To avoid divide by zero errors and to eliminate the PFT-flag triggering in light-wind conditions, or low fire danger conditions the PFT-flag is set to a very large number (to indicate pyroCb is impossible) whenever $V \leq 2.0$ or $U_{10} \leq 2.0$. The smaller the PFT-flag value, the more favourable the combined plume and fire conditions are for pyroCb formation. When calculating PFT-flag, PFT in units of GW is used in Eq. 9 to generate manageable PFT-flag units for plotting (e.g., Figure 28).

Like the PFT, the PFT-flag is still in development and will undergo further tuning and editing in the future.

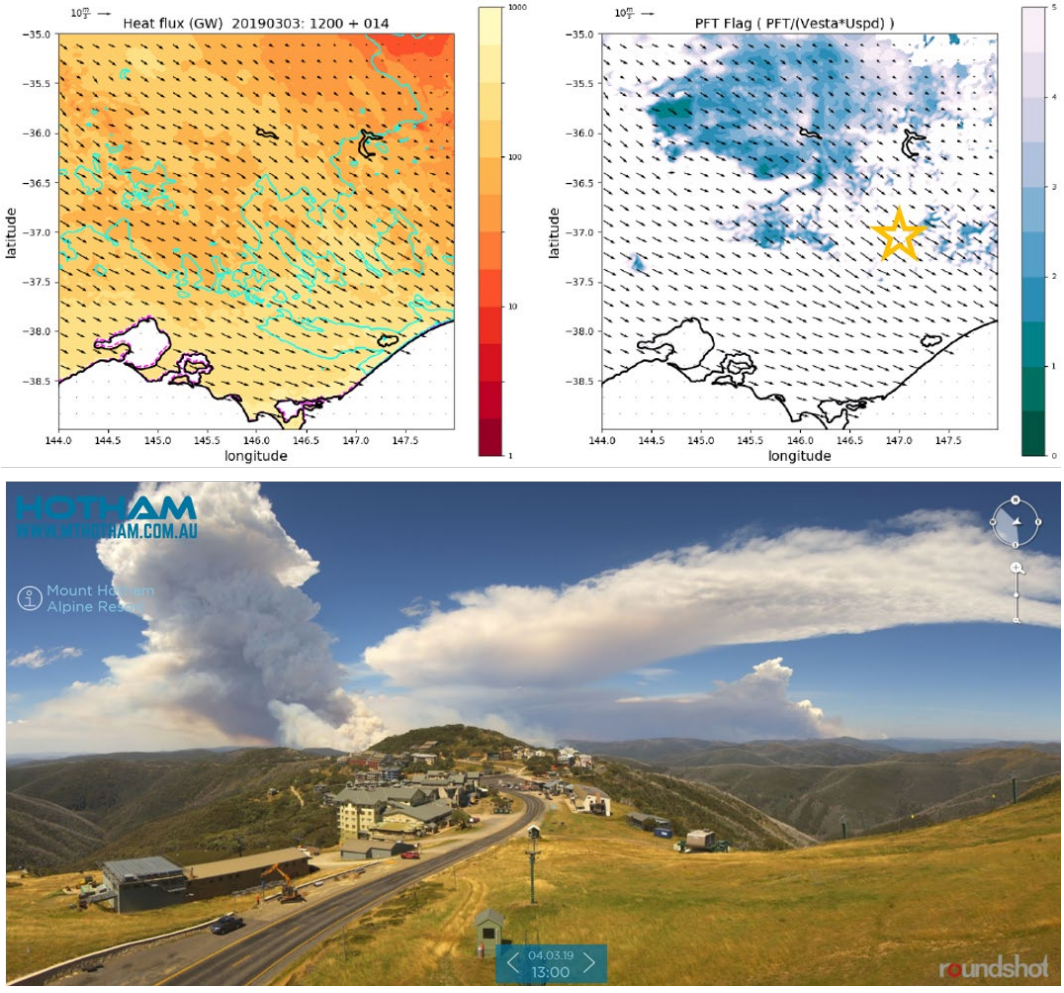

PFT-flag results

A rigorous assessment of the PFT-flag is yet to be undertaken. At present a prediction of favourable pyroCb formation is based on a somewhat arbitrary choice of colour-scale in the PFT-flag plots. Both Sir Ivan and Licola yielded PFT-flag values between about 1 and 3 (Figure 28), which led to the choice of a colour-scale between 0 and 5. Initial interpretation of these values is as follows: less than 1 is highly favourable for pyroCb formation; between 1 and 3 is favourable; between 3 and 5 is possible; and greater than 5 is unlikely. This is very preliminary and is likely to change in the future.

As mentioned above many cases have since been examined, with very promising results. A good example is the dual fires that produced pyroCb in view of the Mt. Hotham webcam on 4 March 2019: The Mayford-Tuckalong Track fire and the Mount Darling-Cynthia Range Track fire. Early in the day, while impressive pyroCu developed, there was no indication of pyroCb, and the PFT-flag indicated pyroCb was unlikely in the immediate vicinity of the fires (although favourable conditions were indicated nearby, Figure 29). Throughout the afternoon, the PFT flag showed a steady southward progression of favourability, with the shading becoming increasingly darker (c.f., Figs 5 and 6). Both fires produced pyroCb, with the northern fire (Mayford-Tuckalong Track closest to the camera) triggering about an hour earlier than the southern fire.

The sequence of 10-minute images (not shown) showed the two plumes produced multiple bursts of deep convection throughout the afternoon, progressively becoming larger and more energetic, with clear evidence of rain falling from the nearby fire plume, and a mature anvil in the more remote fire plume. This behaviour was well-matched by the ever darkening and southward progression of the PFT-flag shading. As a forecast tool the PFT-flag would have provided excellent guidance for pyroCb prediction on these two fires.

Summary

A series of CRC studies beginning with a little "blue-sky" research into the thermodynamics of smoke plumes led to the ability to identify potential condensation heights in plumes, and the minimum plume buoyancy required for plumes to freely convect to the electrification level. With this knowledge equations for a theoretical minimum firepower required for pyroCb formation were derived using the Briggs plume model (PyroCb Firepower Threshold, PFT).

The work has culminated in the development of a diagnostic that seeks to determine when the atmosphere is conducive to both deep plume development and large, hot fires (PFT-flag). Originally designed as a flag to alert users when to examine the PFT, the PFT-flag may prove to be a more valuable prediction tool than the PFT itself. It was developed and tuned to identify the atmospheric conditions corresponding to two pyroCb events at opposite ends of the pyroCb spectrum. The first (Sir Ivan) occurred in catastrophic fire weather conditions, when pyroCb formation conditions were not especially favourable. The second (Licola) occurred in much milder fire weather conditions, when plume formation conditions were considerably more favourable. The PFT-flag has now been applied to more than 20 cases covering multiple days and time periods. While no quantitative performance assessment has yet been made the

tool appears to be working surprisingly well. It not only identifies days of pyroCb occurrence, but also reproduces the diurnal variation in pyroCb threat, plus variations in threat associated with atmospheric features such as troughs, fronts and sea-breezes.

The next section describes real-time testing of the PFT and PFT-flag during the southern Australian fire season of 2019-20.

Practical implications

This section describes a scientifically rigorous, objective tool for helping predict the formation of pyroCb. There are two key advances:

- The problem has been formulated as a firepower threshold and depends only on meteorological quantities.
- We have developed a simple way of calculating an approximate PFT, suitable for use on observed or modelled data.

Meteorological quantities are more accurately predictable than fire (for reasons including that the difficult problems of ignition and suppression are avoided) and available days in advance. Hence, a quantity that depends only on the meteorology will be available on the same forecast horizons as other weather data, providing information up to a week in advance of the highest-risk periods and locations for pyroCb development. We expect that in the future we may refine the calculation of the PFT, either by refining the existing method or introducing a substantially new one. However, the basic paradigm of predicting a threshold will, we believe, remain a useful conceptual framework.

A partial barrier to operational use is the lack of power data from fires. This barrier does not prevent qualitative use but is a limitation for more quantitative interpretation. Many fire agencies use fire simulators to assist in fire management. Information sufficient to calculate the power output is available within these simulators but is not always presented to the user. We encourage that these data be made available.

PYROCUMULUS PREDICTION: RESULTS OF THE REAL-TIME TRIAL

The previous section describes a method for predicting the thermodynamic favourability of the atmosphere for pyrocumulonimbus cloud formation was developed (Tory et al. 2019). This was extended to produce a tool that quantifies this favourability, while also incorporating the effects of wind on plume rise. The quantifiable quantity is a theoretical minimum firepower required for pyroCb to form in a given environment, termed the pyrocumulonimbus firepower threshold (PFT). This initial version (hereafter PFT1) was reported on in the previous annual report. It was found that conditions that favour large and intense fires with the capacity for initiating pyroCb (hot, dry and windy) were opposite to the ideal conditions for deep, moist plume growth (cool, moist and calm), suggesting that perhaps only a somewhat limited range of 'compromise' conditions might be capable of supporting pyroCb formation. A PFT flag was developed (essentially the ratio of the PFT to a modified fire danger index) to identify these ideal conditions (see previous section and Tory 2019). Version 1 of the PFT flag is hereafter referred to as 'PFT1 flag'.

A suite of PFT1 products were tested in a real-time forecast trial during the southern Australian 2019/2020 fire season. These included PFT1 and its three main ingredients, plus the PFT1 flag and some of its ingredients.

The trial prompted extensive discussions with operational personnel, from both the Bureau of Meteorology and fire agencies, and it became apparent that the specific focus on pyroCb was too narrow, because any towering moist convection can be hazardous. Here we introduce the term Moist Pyro-Convection (MPC) to explain any cloud that forms in smoke plumes, and hereafter refer to 'deep' MPC to represent towering pyrocumulus or pyrocumulonimbus events, in recognition that both can be hazardous to people working on the fire ground. We also introduce the term fire-generated thunderstorms (FGT), which is a more descriptive name than pyrocumulonimbus, used by forecasters to improve public messaging.

Prior to the trial efforts were made to raise awareness of the products to Fire-Weather Forecasters (FWF) and Fire-Behaviour Analysts (FBAn) in the form of presentations at conferences and invited lectures. Training material was also prepared and delivered to end-users who received the PFT forecast products. The trial was a success, with some indication of deep MPC threat identified for almost all cases. The PFT and products were strongly endorsed by users. The trial raised a number of issues regarding MPC prediction in general and identified opportunities for further development of the PFT product suite. The first journal paper on the PFT was submitted to the American Meteorological Society journal, Weather and Forecasting, and is now in revision.

Pyrocumulonimbus Firepower Threshold: PFT

The PFT used in the trial (PFT1) takes into consideration the magnitude of any inversion or stable layer the smoke plume must penetrate, the height the smoke plume must rise before sufficiently buoyant cumulus clouds form in the smoke plume, and incorporates the impact of background wind on plume-rise via the Briggs (e.g., Briggs 1984) plume-rise model. The highly complex process of moist plume development is reduced to a single analytic equation at the core of PFT1,

which necessarily requires several assumptions. However, the success of PFT1 during this trial, together with the extensive prior history of the Briggs model, suggests that the assumptions are reasonable. Below the condensation height it is assumed that the smoke plume can be described by a Briggs plume rising through a neutrally stable layer of constant background wind. Above the condensation height, it is assumed the rising moist plume can be described by simple parcel theory (as applied to common thunderstorm forecast products). One of the closure requirements for PFT1 is the specification of a moist, plume-top height. Pressure levels corresponding to typical tropopause heights were initially considered, but a desire to include potentially intense, shallower convection (including shallow Fire-Generated Thunderstorms, FGT), led to the choice of a temperature dependent height corresponding to a level where cloud electrification is likely (-20 ℃). It was recognised that this limit would over-predict lightning producing storms, but have the benefit of identifying the potential for development of other types of potentially dangerous moist pyro-convection.

Early testing of PFT1 highlighted the fact that atmospheric conditions that favour the development of large and intense fires with deep flaming (hot, dry and windy) do not favour moist plume development (cool, moist and calm), and vice versa. Deep, moist pyro-convection only develops when there is some overlap in these opposing sets of conditions. It became clear that a useful forecast tool based on PFT1 would need to incorporate conditions that relate to potential fire size and intensity, either built into the diagnostic or separate to it. A PFT1 flag was trialled designed to flag forecaster attention whenever a favourable combination of moist plume formation and large intense fire formation conditions were present (Tory 2019). The PFT1 flag is essentially a ratio of PFT1 to a modified Vesta fire-danger index (Cheney et al. 2012). Forecast plots of the PFT1 flag and the modified fire-danger index were added to the suite of PFT1 forecast products (hereafter PFT1 forecasts).

An example of the PFT1 products is shown in Figure 30 corresponding to the most active day of the season, 30 December 2019 with seven confirmed fire-generated thunderstorms (FGT) and several other deep MPC bursts.

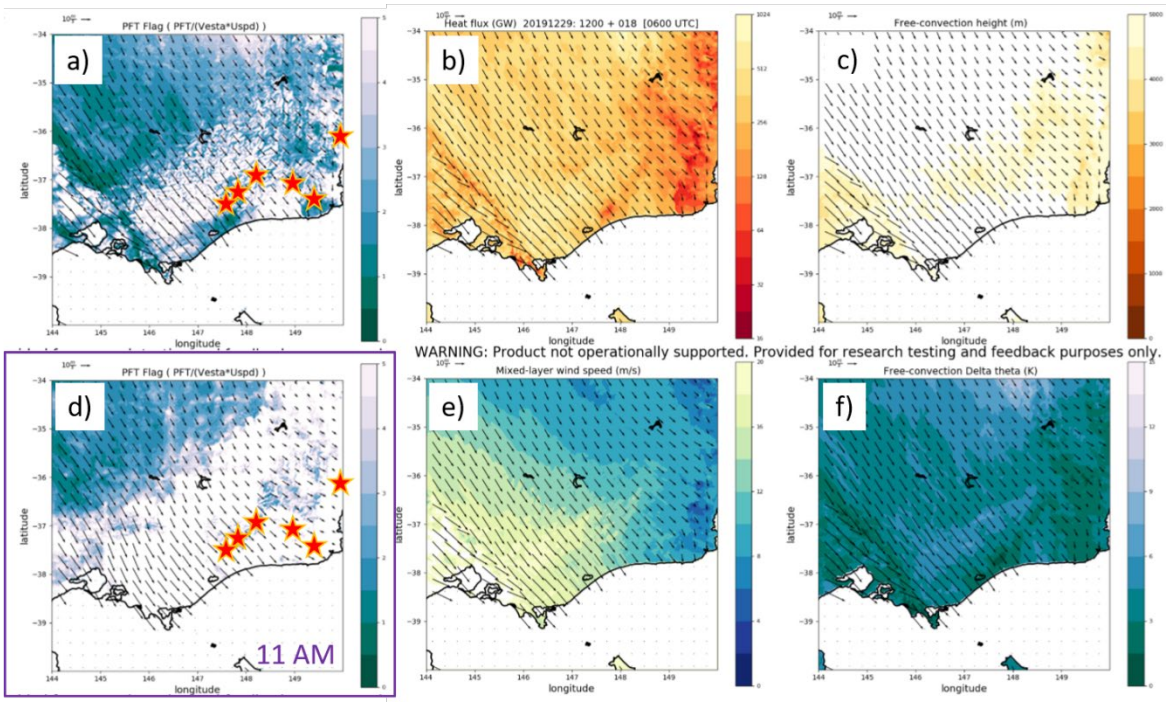

Figure 30: PFT1 forecast products for 30 December 2019 at 5 PM local time (0600 UTC) except panel (d) which is 11 AM (0000 UTC). The stars in panels (a) and (d) indicate the location of notable deep MPC associated with fires (from left to right) Bullumwaal, Ensay, Diggers Hole, Wulgulmerang, Cann River (the Cann River fire also produced multiple FGT later that night) and a fire east of Cooma. The panels include: (a),(d) PFT flag, arbitrary units; (b) PFT, units GW (log base 2 colour scale from 16 to 1024); (c) free-convection height, z_{fc}, units m (colour scale 0, 5000); (e) mixed-layer wind speed, U, units ms^{-1} (colour scale 0 to 20); (f) free-convection Delta theta, $\Delta\theta_{fc}$, units K or ℃ (colour scale 0 to 15). The quantities in panels (c), (e), (f) represent the terms on the right-hand-side of Eq. 1. In each panel darker shading indicates higher favourability for deep MPC.

Preparation for the real-time trial

In August 2019, one-hour presentations on the PFT and PFT sub-products were delivered to Fire-Behaviour Analysts (FBAn) at the NSW Regional Fire Service headquarters, and to Fire-Weather Forecasters (FWF) in the Bureau of Meteorology office in Sydney. There was considerable interest in the products in both organisations.

Scripts were completed that generated twice-daily PFT-product forecasts for five sub-domains around Australia plus an Australia-wide domain. The forecast period was initially 48 hours, but this had to be reduced to 36 hours for technical reasons (since overcome). Requests from users during the real-time trial for longer forecast periods, led to the addition of 36—72-hour forecasts for the NSW, Qld and Australia-wide domains.

Training material was prepared to educate FWF and FBAn users on the theory behind the products, how to understand them, and how they might be used to help in the forecast process of deep MPC. The training included a lecture, and a practical exercise designed to build a deeper understanding of the theory and an appreciation of how it relates to an existing fire-thunderstorm diagnostic concept.

The practical exercise made use of an approximation to PFT1 (Tory and Kepert 2020) designed to facilitate a simple process for estimating and visualising PFT1 on a thermodynamic diagram,

$$PFT1 \sim 0.3 \times (z_{fc})^2 \times U \times \Delta\theta_{fc}. \qquad 1.$$

z_{fc} is the free-convection height in units of km (the minimum height the plume must rise to initiate deep moist pyro-convection), U is the average velocity magnitude of the mixed layer horizontal wind, units ms^{-1}, and $\Delta\theta_{fc}$ represents how much warmer the plume needs to be than the mixed-layer potential temperature, units ℃ (or K). These units yield PFT1 values with units of GW. An example is shown in Figure 31 during the evening immediately following the disastrous Black Saturday fires (e.g., Cruz et al. 2012).

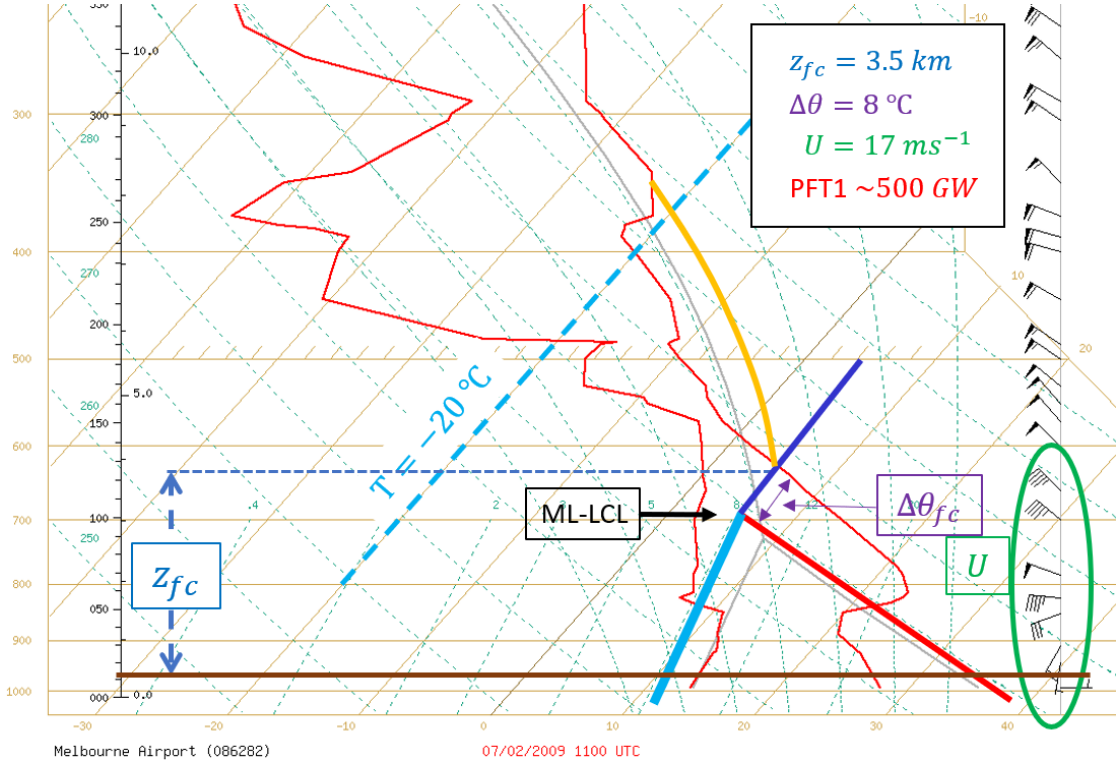

Figure 31: Melbourne Airport sounding 1100 UTC, 7 February 2009 (10 PM local time), with markers used in the manual PFT calculation (Eq. 1). Thick red and cyan lines represent the mixed-layer potential temperature and specific humidity respectively, with the mixed-layer lifting condensation level (LCL) marked at the apex of these two lines. The saturation point (SP) curve (approximated by the blue line) identifies the potential plume condensation positions. The minimum-buoyancy plume-path is indicated by the gold curve. The free-convection height (z_{fc}) is the height at the base of the gold curve. The plume must be at least $\Delta\theta_{fc}$ warmer than the mixed-layer potential temperature (thick red line), and the average mixed-layer wind term, U, can be estimated from the wind barbs highlighted in green.

The mixed-layer potential temperature (thick red line) and specific humidity (thick cyan line) represent the average thermodynamic properties of air entrained into the plume. The saturation point curve is approximated by a straight line representing a 15:1 ratio of plume temperature increment to plume moisture increment (units $K/g\ kg^{-1}$, blue line). It represents a range of possible positions where the plume can reach saturation (Luderer et al. 2009, Tory et al.

2018). The minimum-buoyancy plume-path that freely convects to a level cooler than -20 ℃ is indicated by the gold curve. The free-convection height (z_{fc}) is the height at the base of the gold curve. Here the plume must be at least $\Delta\theta_{fc}$ warmer than the mixed-layer potential temperature (thick red line), and the average mixed-layer wind term, U, can be estimated from the wind barbs highlighted in green.

Unfortunately, the early start to the fire season resulted in the cancellation of in-person training, requiring a shift to online training in the form of a webinar (https://www.bnhcrc.com.au/resources/presentation-audio-video/7033). The PFT product forecasts were made available to FWFs and FBAns who had completed the training.

General results of the real-time PFT forecast trial

A real-time trial of PFT1 forecasts took place during the 2019/2020 southern Australian fire season. Forecast maps at 6-hourly intervals were automatically generated from ACCESS-R data (the Bureau of Meteorology's numerical weather prediction model) for five regions around Australia, plus an Australia-wide region, and sent to end-users including FWFs within the BoM and FBAns in the state fire agencies. Additional high-resolution forecasts at hourly intervals were prepared on demand for high threat days when possible (this process was manually initiated). The trial was very successful, exceeding our expectations. It provided very useful guidance to forecasters briefing emergency services and identified the potential for development of nearly all intense MPC events. This success provides strong evidence that MPC are inherently predictable, and that given the substantial simplifications and assumptions incorporated into PFT1, there may be a fairly distinct separation between typical favourable and unfavourable conditions from day-to-day. The success also means that improvements could focus more closely on timing and location (i.e., reduce the false-alarms without degrading the forecast by increasing the number of misses) and perhaps focus on refining the products to target specific types of MPC. To this extent the trial highlighted several promising development options.

The trial helped identify important issues regarding MPC and FGT prediction and verification, more generally. The first issue stems from a general vagueness surrounding MPC with respect to:

- Threat variability (i.e., can we distinguish dangerous from benign);
- Understanding plume structure and behaviour and the relationship to specific threats (i.e., can we distinguish specific threats from plume observations);
- Limited methods for verification and observational confirmation; and a lack of clarity on what is being verified or confirmed.

Each point makes the forecasters task more difficult.

The second follows from the first but incorporates the broader class of forecasting challenges. The general vagueness surrounding MPC mean there is an inevitable lack of clarity in messaging and the end-user perceived messages on the predicted MPC and associated threats. The forecast challenges have been exacerbated by the sudden increase in the observed frequency of MPC events,

from a few events per decade to tens of events per season. Compared with many other severe weather events forecasters have limited experience with MPC, and associated prediction tools and procedures are less well-developed. The great unknown surrounding FGT and MPC in general and the unprecedented threat that fire agencies had to deal with, created a high-pressure forecast environment, in which decisions had to be made on minimal information. This resulted in PFT1 forecasts being used in the decision-making process more than we anticipated (rather than primarily being monitored to assess performance). Furthermore, individual usage and interpretation varied among FWFs and FBAns contributing to subjectivity and variability in messaging. The relative frequency of events in each jurisdiction also had an impact on how PFT1 products were used and the value they provided to the forecast process. In NSW, the progression of intense fires from the north to the south over a six-month period, provided a good opportunity for FWFs in particular, to test and learn how to use the PFT1 to full advantage, whereas FWFs in less-active states, such as SA and Tasmania, did not get this opportunity.

A third issue arose from questions regarding whether PFT1 products could be used to investigate the sudden increase in MPC events. Given that only atmospheric ingredients are included in these products, any climate application can only identify possible changes in atmospheric conditions, and only those conditions explicitly included in the PFT1 products. It became apparent during the trial, in consideration of the first and second issues, that any climate study should ideally be delayed until additional atmospheric phenomena important for MPC development be incorporated into the two diagnostics, and until we are clear about what we wish to identify (e.g., just FGT or the broader class of deep MPC).

Specific results of the real-time trial

FWFs had limited observations available to identify when deep MPC become dangerous. Lightning observations (generally available in almost real-time) are good indicators that potentially dangerous indrafts and other hazards have developed, although not all intense, deep MPC produce lightning. Lightning being one of the only concrete indicators, led to a focus on predicting lightning producing events, since they were "verifiable". However, this inadvertent FGT lightning focus may have proved a distraction as considerable effort was employed to try and predict lightning with few tools at hand, and an uncertainty whether conventional thunderstorm tools and knowledge was valid for FGT (e.g., if the cloud electrification process is greatly affected by smoke). This was confounded by enigmatic behaviour when very deep MPC on one day failed to produce lightning and very shallow MPC the next day produced abundant lightning. In examining this behaviour, it became apparent that parcel theory, used to represent moist plume-rise, could not explain observed cloud-top temperatures, and using parcel theory in PFT1 would bias PFT1 towards over-prediction of deep MPC when the middle troposphere was dry and/or relatively stable. It also became clear that a general lightning-prediction rule-of-thumb, cloud-top temperatures < -20 ℃, was not a good predictor of lightning in MPC (although it was a good indicator of deep MPC). A cloud-top temperature threshold of -40 ℃ was a more reliable indicator of FGT.

Case studies discussed in an upcoming report illustrate the limitations of parcel theory, and the apparent inconsistencies in lightning formation in deep MPC. They showed that reliance on cloud-top temperature alone was problematic, and that it may be necessary to consider minimum cloud-base temperatures (to ensure the presence of super-cooled water), and minimum plume buoyancy (to accelerate the plume updraft to charge-separation speeds). Other cases showed that even with all these ingredients, lightning did not form, suggesting we still have much to learn about cloud electrification in MPC.

Two "missed" events were noted. In the first, an intense nocturnal FGT on the Marthavale fire, PFT1 was successful but the PFT1 flag was not triggered. During the hours leading up to and following the FGT, firepower estimates in real time suggested a large increase in activity from 100 to 1000 GW, during which time PFT1 dropped from about 500 to 250 GW, indicating a cross-over from very much less to very much more firepower required for deep MPC formation around the time the FGT developed. Model-predicted near-surface conditions were too cool, moist and calm for the PFT1 flag to be triggered. Observed near-surface conditions, however, were much warmer, drier and windier.

The second missed event also occurred overnight when model-predicted near-surface conditions were probably too cool, moist and calm. However, in this case it seems possible that either the plume dynamics caused a local reversal of the flow from offshore to onshore, or the model failed to predict onshore flow, which in either case would have entrained cooler and moister maritime air into the plume. Also, if these flow differences were combined with a semi-stationary convergence line between the terrestrial and maritime airmasses, the plume lift may have amplified. Together, a lowered plume condensation level and increased plume-rise height may have been expected, suggesting the actual PFT1 may have been an order of magnitude smaller than values obtained from the model fields (i.e., conditions much more favourable than the model predicted).

The dramatic increase in observed FGT and deep MPC in general in southern Australia and globally, particularly in the last two years is at least partly due to warming and drying of forests contributing to very high fuel loads and more frequent extreme fire conditions. It is also possible that atmospheric conditions have become more favourable. If PFT tools that are purely functions of atmospheric variables can be demonstrated to be good predictors of deep MPC events, then they can be applied to reanalysis data dating back a few decades to see if a significant atmospheric change can be identified, and can also be applied to climate models to assess possible future changes. For climate applications PFT products need to be more objective than for forecast applications (e.g., forecasters can easily dismiss obvious cases where deep MPC cannot develop), and it is also desirable to reduce false alarms to best represent the true climatology of deep MPC. To this end, support from a National Environmental Science Programme project, has contributed to the development of methodologies to largely eliminate false alarms from cold-air outbreaks, and to incorporate the negative impacts of entrainment on moist-plume buoyancy.

A number of desired outcomes have been identified in preparing this report.

- An improved understanding of MPC generally, plus their hazards (how, where and when they form) and concrete understanding of what it is we wish to predict.
- Further development of MPC forecast procedures and tools.
- Further development of the PFT for climate applications.

Specific lessons learned when using PFT1 include:

- The -20 ℃ plume-top assumption is a good indicator for deep MPC, but not FGT specifically.
- Parcel theory used for moist plume-rise will over-predict deep MPC conditions when the middle troposphere is dry and/or relatively stable.
- Incorporating moist-plume entrainment in a future PFT will address the previous point.
- To predict FGT lightning specifically would require a greater focus on cloud electrification processes and might require a better understanding of cloud electrification in smoky plumes.
- High spatial and temporal resolution forecasts of PFT products are often needed to identify deep MPC threat, especially when favourable conditions are tied to relatively small-scale transient features such as wind change lines.
- The PFT1 flag, which is a function of near-surface atmospheric variables, is susceptible to common nocturnal forecast model biases in these variables, leading to the under-prediction of fire-danger severity and the potential for deep MPC development.
- Offshore PFT values may need to be considered for fires burning very close to the coast, when a coastal PFT gradient is present separating warm and dry terrestrial air from cool and moist maritime air.

Summary

Preparation for the real-time PFT forecast trial, including awareness raising and training was successful, although we had to withstand the challenge of an early start to the season. The trial exceeded our expectations, with the PFT and PFT flag performing very well, and receiving frequent positive endorsements. The trial identified several broader issues surrounding PyroCb understanding and definitions, and forecasting challenges for a phenomenon that until recently occurred very rarely. The trial also identified a number of opportunities to further develop the PFT suite of products, many of which will be addressed in a CRC sponsored Utilisation project.

Practical Implications

The trial was very successful in raising awareness of the tool and for providing a strong impetus and demand for ongoing use. It raised important questions, discussed above, including what constitutes dangerous moist pyroconvection, and what operational actions should be taken when moist pyroconvection is expected. It also indicated some minor limitations of the current tool to address.

The success of the trial also strongly indicated the value of such exercises, and some valuable lessons in how to best run them. Partly because of the extraordinary 2019-20 season, the trial made great demands on the staff involved and the IT systems were barely fit for purpose. With the benefit of these lessons, we hope to be able to run similarly valuable trials in the future.

HIGH-RESOLUTION ENSEMBLE PREDICTION OF AN EAST COAST LOW

Summary

During autumn and winter months the eastern coast of Australia is periodically affected by rapidly developing and intense extratropical low-pressure systems that are known as East Coast Lows (ECLs). ECLs bring damaging winds and heavy rainfall with flooding that can last for several days. Due to their rapid development, many forecasting issues arise; such as: predicting the location of the system along the coast, and the intensity and location of maximum winds and rainfall. The use of ensembles can help overcome these forecast challenges and provide insight into how these systems form.

The event that is studied here occurred during 20-23 April 2015, with the worst impact on 21 April. This event was simulated using a 24-member ensemble of ACCESS nested models. Simulations show that forecast rainfall (averaged over the 24 ensemble members) is in good agreement with observed rainfall and the ensemble identifies Dungog as the area of substantial risk of extreme rainfall. A subset of ensemble members shows very little rain in the Dungog area; instead it is located further south or east, indicating that a large part of the coast is at risk of significant rain. Ensemble comparison shows that the strongest surface winds and highest rainfall occur to the south of the main low, together with the strengthening of the north-westerly jet and the upper cut-off low west of the surface low. However, localised differences between members produce large differences in location and strength of extreme winds and highest rainfall. A better understanding of the dynamics and predictability of ECLs has important implications not only for forecasters, but also for emergency services preplanning and preparedness.

Introduction

The focus of this study is a low-pressure system that affected the eastern coast of Australia between 20 and 23 April 2015 (known locally as an east coast low), with the worst impact on 21 April. This was a devastating event for the Dungog and Maitland area (map in Figure 32), with at least 4 deaths reported and widespread damage. Dozens of houses lost their roofs, over 200 000 houses were without power, and 57 schools closed. By studying this event, better understanding of the dynamics and processes that lead to the rapid intensification of these systems can be gained. Additionally, the predictability of the event can be studied as there are many issues associated with the forecasting of east coast lows (e.g., intensity and location of maximum winds and rainfall, location along the coast).

East Coast Lows (hereafter ECLs) are rapidly developing and intense extratropical low-pressure systems that periodically affect eastern coast of Australia during autumn and winter months. They form preferentially at night (Holland et al.1987) and although they are most common during winter months, they can occur at any time of the year (Speer et al. 2009). ECLs bring damaging winds and heavy rainfall with flooding that can last for several days. They are responsible for about 50-80 % of extreme rainfall events within the eastern seaboard of Australia (Dowdy et al. 2013), where a large majority of the Australian population live, and are also very important for catchment inflows and

water security east of the Great Dividing Range (Pepler and Rakich 2010). Recently, the "Pasha Bulker" storm of June 2007 caused the grounding of the Pasha Bulker carrier ship as well as major socioeconomic losses, making it one of the most significant meteorological events for Australia (e.g., Mills et al. 2010, Verdon-Kidd et al. 2016).

Holland et al.'s (1987) study showed that ECLs are mesoscale phenomena that form on the coast within a distinctive synoptic system that consists of a trough (or "dip") in the easterly wind regime over eastern Australia. This pattern is locally known as an "easterly dip" and its development is similar to the one associated with a "bomb" that develops off the eastern U.S. seaboard (e.g., Sanders 1986). ECL's formation is also associated with strong oceanic temperature gradients. Based on whether ECLs form east of, on, or west of the dip axis, Holland et al. classified them into three types.

Type 1 ECLs form when a developing easterly dip is present over the continent; they are small in size and form preferentially over the ocean. In contrast, development of ECLs of type 2 is rapid and occurs within an easterly dip on the coast. The majority of ECLs that were studied by (Holland et al. 1987) were of this type. Type 3 systems form on the western side of an easterly dip in the Tasman Sea; they are small (100 km across), develop rapidly and often disperse within 24 hours of their formation.

There are many forecasting challenges associated with ECLs due to their rapid development. These include predicting the location of the system along the coast, the intensity and location of maximum winds and rainfall, their relationship to the centre of the developing low as well as the likelihood of the associated storm surges (Leslie and Speer 1998, Mills et al. 2010). Another challenge associated with ECLs is their representation in numerical weather prediction (NWP) models, especially since the area of most intense winds and heavy rainfall is usually very narrow. For example, in the "Pasha Bulker" storm the strongest wind and rainfall bands were on the order of 100 km wide.

Large scale atmospheric processes associated with explosive cyclogenesis are reasonably well represented by today's NWP models and by coarse-resolution global climate models (Mills et al. 2010, Dowdy et al. 2013). However, smaller scale structures and processes that occur within the overall system and often contribute significantly to extreme weather are poorly represented by these models and require higher horizontal resolution. Fine details are also less predictable and one way of overcoming these challenges is by using an ensemble of high-resolution simulations. Ensemble forecasts provide opportunity to study predictability of ECLs by analysing differences between the simulations, as well as a more thorough analysis of processes that lead to rapid intensification of these systems.

Accurate forecasts of severe weather events such as ECLs are important not only for meteorologists but also for emergency services and communities as they enable better risk management and preparedness for the event. Therefore, the main aim of this study is to, by analysing a high-resolution ensemble of 24 simulations, elucidate processes associated with the development of the April 2015 east coast low and learn about the predictability of the event.

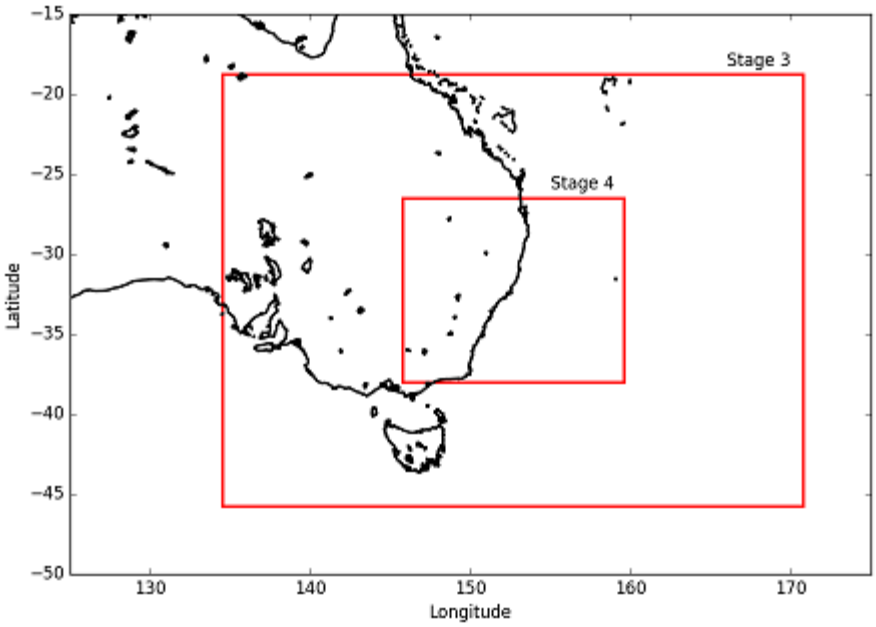

Figure 32. Outline of model domains, with larger domain having horizontal resolution of 4 km (stage 3) and smaller domain with horizontal resolution of 1.3 km (stage 4).

Set up and design of ensemble simulations

In this study, high-resolution ensemble forecasts are used to analyse the development and dynamics of the east coast low of April 2015. For all simulations, version 8.5 of the UK Met Office Unified Model (UM), the atmospheric component of Australian Community Climate and Earth-System Simulator (ACCESS) is used (Puri et al. 2013). The ensemble consists of 24 simulations (1 control + 23 perturbed) that are initialised from the 24 members of the Bureau of Meteorology prototype global ensemble prediction system (ACCESS-GE), which is based on the Met Office ensemble prediction system known as MOGREPS (Bowler et al. 2008). All simulations were initialised at 0000 UTC on 20 April 2015 (+ 10 hours for local time) and ran for 48 hours.

The model consists of a global model run that is nested down to 4 km (hereafter stage 3) and 1.3 km (hereafter stage 4) regional runs (Figure 32 shows the outline of stage 3 and stage 4 domains). The smallest grid spacing (1.3 km) sufficed to capture the dynamics of the event because of balance between the domain size and resolution. Each domain has 70 vertical levels and the model top for nested domains is 40 km. Results from the 1.3 km domain are shown in this report.

Results from ensemble simulations

Overview of the mean forecast and comparison with observed rainfall

In this section an overview of the ensemble mean forecast is presented. Figure 33 shows 1200 UTC 20 April 2015 near-surface (1 km) and upper-level (10 km) flow comparison for the ensemble mean forecast. Strong south-easterly near-surface winds are seen to the south of the main low in Figure 33a, and at upper levels there is a clockwise tilted (towards the northeast) trough that moves toward the coast, with a strong north-westerly jet to the northeast of the upper-level trough

(Figure 33b). This pattern seems to be common for most ECLs (Mills et al. 2010). However, there is large variability in the location and strength of maximum surface winds between the members, which is associated with differences in surface and upper-level flow characteristics and these are discussed in more detail in subsequent sections (note that this mean forecast has somewhat smooth features due to the ensemble averaging).

A comparison between the observed 48-hr rainfall total (up to 09 am 22 April 2015) and the forecast 48-hr rainfall total, which is averaged over the 24 ensemble members, is shown in Figure 34. It is evident that the forecast rainfall is in good agreement with the observed rainfall, even though the rainfall in the Dungog area is not as heavy as observed. Some individual members (Figure 35) had higher values than the ensemble mean and closer to the observed peak. The probability-matched ensemble mean (PMEM) of 48-hr rainfall total (Figure 34c) retains these peaks and shows even better agreement with the observed maximum rainfall. Forecast probability of 48-hr total rainfall exceeding 100 mm (Figure 34d), defined as the proportion of the ensemble exceeding this threshold, shows that the ensemble identifies the Dungog area as being of substantial risk of extreme rainfall.

Figure 35 shows that there is a large spread in 48-hr total rainfall between ensemble members. Some members predicted very little rain in the affected area (white circle in Figure 35); rather rain was located more to the south (e.g., member 11) or east-north-east of Dungog (e.g., members 17 and 20). This indicates that a substantial part of the coast is at risk of significant rain, and emphasises the importance of studying ensemble simulations in order to improve prediction of this event but also of other severe weather events.

To examine reasons for differences in the flow development and location of maximum winds and rainfall, ensemble members 11, 17 and 22 are compared and analysed in the next section. Ensemble member 11 is an example of a "dry" member; the majority of predicted rainfall was located to the south of the white circle in Figure 35. Ensemble member 22 is an example of a "wet" member; the highest predicted rainfall amounts (Figure 35) were located inside the white circle surrounding Dungog. In contrast to these, ensemble member 17 was different as it had 2 "centres" of highest rainfall; one to the south of Dungog and one to the north-east (Figure 35r). Analysis of surface and upper-level low development will show that this member differs largely from other members in the location and rotation of the main system, as well as secondary vortices that develop along the shear line.

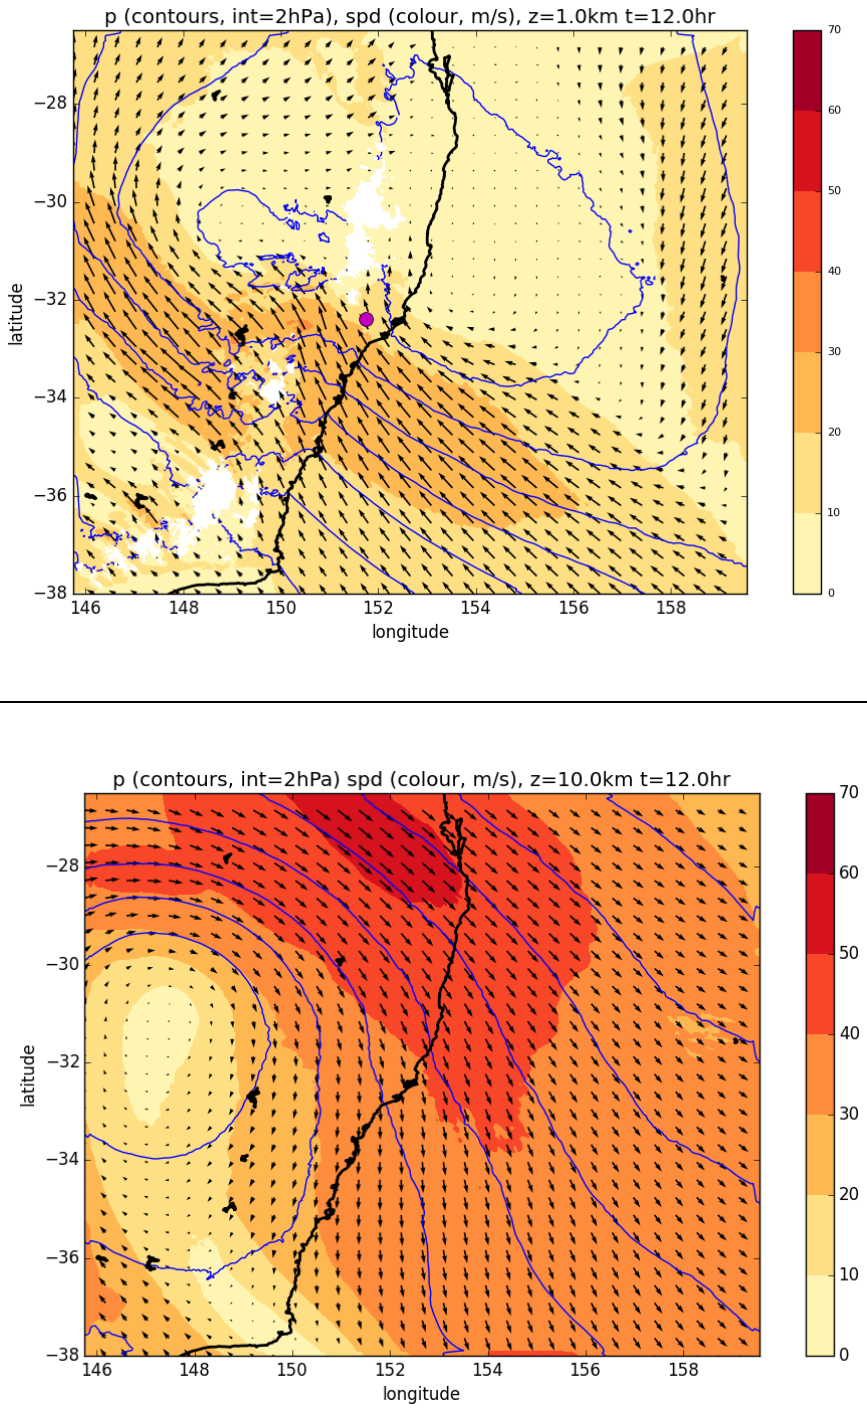

Figure 33: (top) Near-surface (1 km) and (bottom) upper-level (10 km) winds (m/s) and pressure (contours, hPa) for ensemble mean forecast at 1200 UTC 20 April 2015.

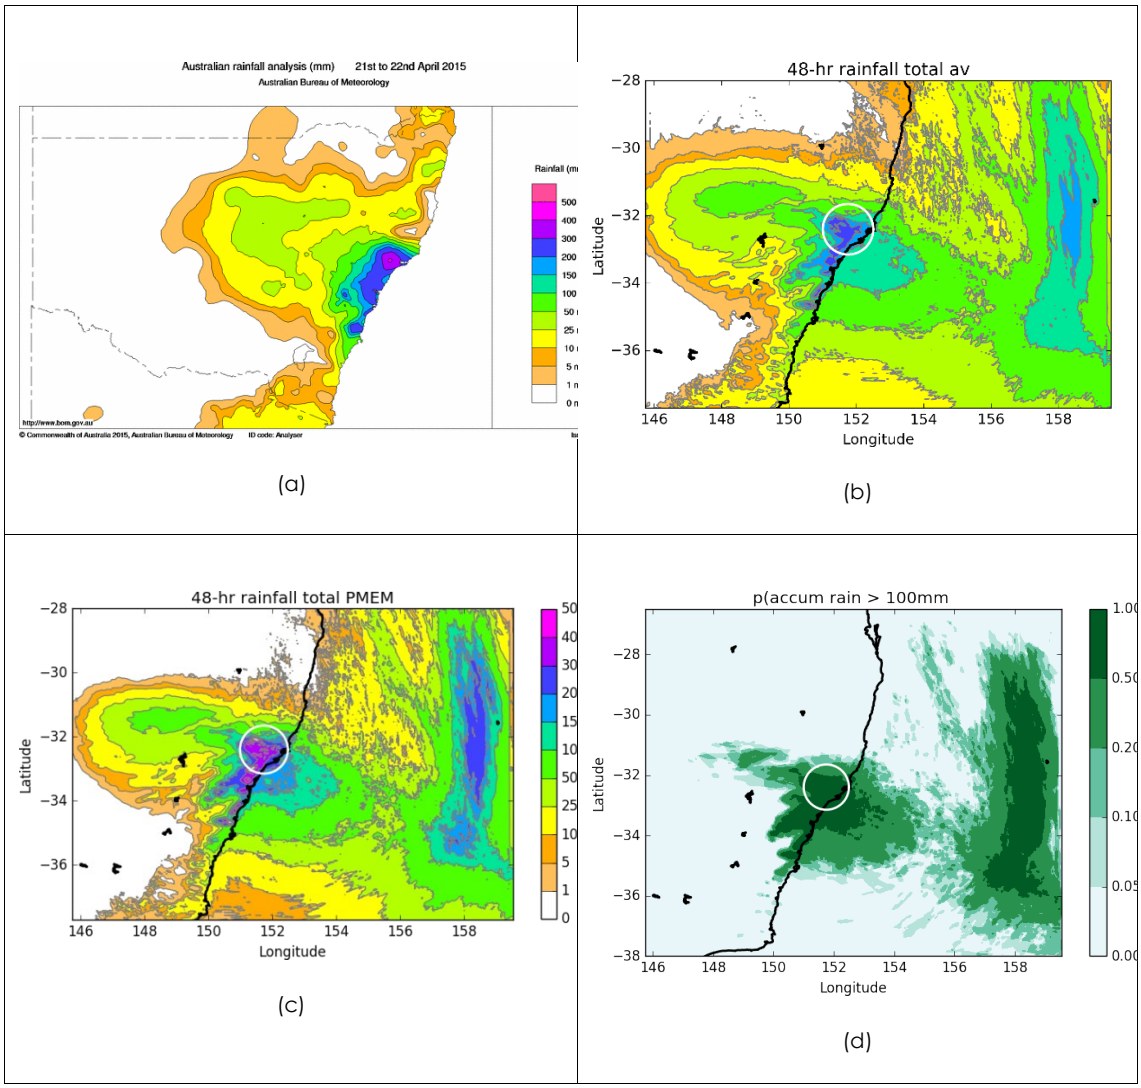

Figure 34: A comparison between (a) observed rainfall (up to 9 am 22 April 2015) and (b) forecast rainfall (averaged over the 24 ensemble members, for 20-22 April 2015). (c) Probability matched ensemble mean, and (d) rainfall probability exceeding 100 mm.

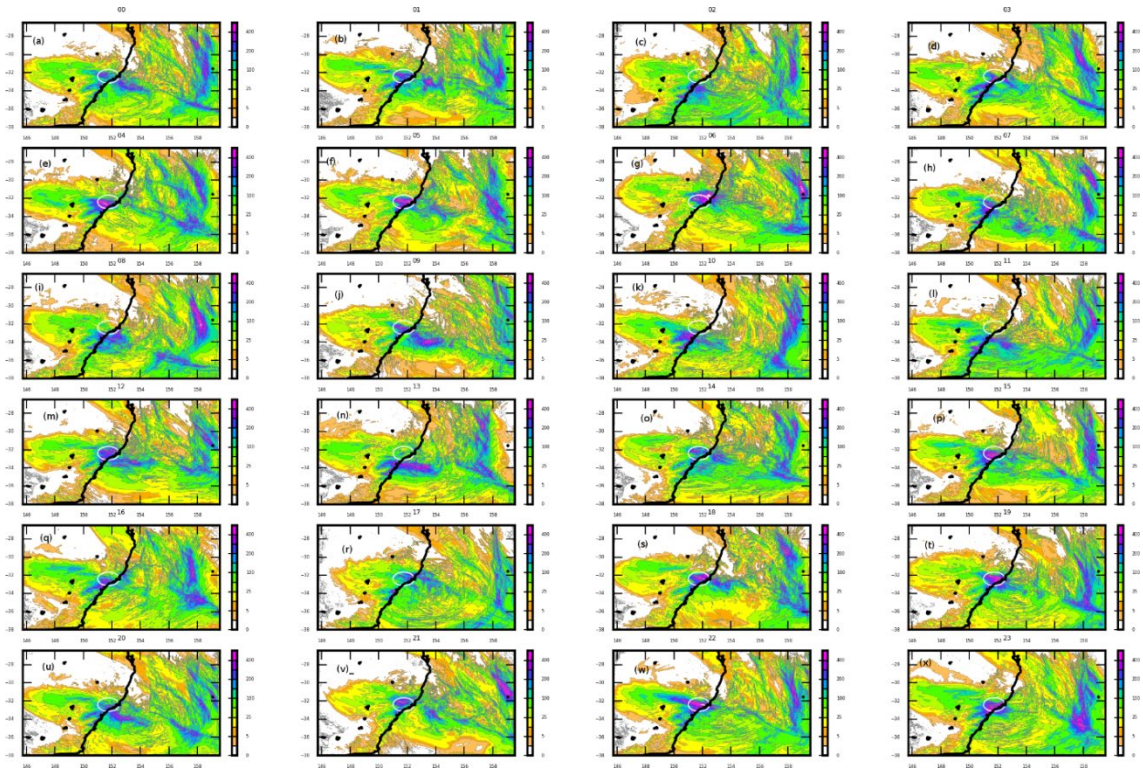

Figure 35. A postage stamp showing 48-hr total accumulated rainfall for all ensemble members.

Differences in surface- and upper-level flow development between the members

As mentioned previously, ECL development is characterised by the interaction of a pre-existing trough of low pressure and an upper-tropospheric cut-off low. Often, many small vortices form on the shear line at the southern edge of the main low, and these can bring the worst weather. Also, the strongest winds and heaviest rain usually occur to the south of the main low centre. A subsequent analysis investigates these features and their relationship in the April 2015 ECL in more detail. By studying the synoptic features as well as similarities and differences between ensemble members, a better understanding of the storm dynamics can be gained. Additionally, features that are more or less predictable can be identified in this way too.

Figure 36 and Figure 37 show near-surface (1 km) and upper-level (10 km) flow comparison for members 11, 17, and 22 at 1000 and 2000 UTC 20 April 2015, and at 0600 UTC 21 April 2015. All members in Figure 36 show strong south-easterly near-surface winds (> 20 m/s) to the south of the main low that are blowing across the isobars and toward lower pressure, with small low pressure centres along the shear line. This cross-isobar flow contributes to surface convergence which, as will be shown below, coincides with regions of highest rainfall. At the same time, the upper-level trough amplifies, cuts off and rotates clockwise, while moving eastwards towards the coast (Figure 37). The north-westerly jet on the northeast flank of the upper-level trough strengthens and splits into a cyclonic branch which curves to the south, and the anticyclonic branch which curves into the Tasman Sea (Mills et al. 2010). Although not evident in these figures due to the domain size, there is a strong ridge to the west and south of the upper-level low; this ridge extends all the way to the surface which contributes to the formation of a strong pressure gradient to the south of the low and subsequently to the strengthening of the low-level jet (LLJ).

Even though basic patterns of the evolution of the upper-level low and the strengthening of the surface low are similar among analysed ensemble members, there are localised differences which can produce large differences in location and strength of extreme wind and rain. It is evident from Figure 36 and Figure 37 that the south-easterly LLJ as well as the upper-level low have slightly different location and strength for the 3 members. These differences are also present amongst other ensemble members; however space does not permit us to show analyses for all ensemble members.

In member 11, the LLJ is located further south than in the other two members, with the rainfall maximum located also to the south rather than around Dungog (cf. Figure 35). The upper-level low is located to the west of the surface low. In member 22, the strong LLJ is directed towards Dungog and the upper-level low cuts off and is located to the west of the surface low. A secondary low forms offshore after 2000 UTC (Figure 36f), moves towards the coast, begins to counter rotate with the main low centre (Figure 36i) and aligns vertically towards the end of the simulation (not shown). Compared to members 11 and 17, member 22 has the highest rainfall accumulated around Dungog. Member 17 has very distinct vortices that develop on the shear line as evident from the pressure contours, and the strongest low developing offshore compared to the other two members (Figure 36b, e and h). Additionally, the LLJ has become southerly and the upper-

level low has moved further off-shore by 0600 UTC 21 April 2015 (Figure 36h and Figure 37h) compared to members 11 and 22.

A comparison of the low-level wind and hourly rainfall for the three members is shown in Figure 38. It is evident that the surface low is not symmetric, and that extreme winds are localised. Also, rain occurs in discrete regions within the low; along the low-level convergence lines to the south of the main low, and also on the eastern edge over the Tasman Sea. Strong near-surface temperature gradients are collocated with regions of extreme wind (not shown). From Figure 38 it is evident that for members 17 and 22 the rain band on the convergence line encompasses the Dungog area by 2000 UTC, however, in member 11 it is located further to the south.

In order to have a better understanding of the relationship between maximum winds and rainfall, mid-level winds and hourly rainfall are compared to a thermal advection rainfall diagnostic (Figure 38 and Figure 40). The diagnostic is based on the relationship between thermal advection and geostrophic winds turning with height (e.g. Holton 2004) which also applies for gradient winds that turn with height (Tory 2014). It has been used for decades by the Australian forecasting community in the prediction of subtropical and tropical heavy, and can be used successfully for most large scale and mesoscale systems, including ECLs. For both geostrophic and gradient wind formulations the relationship is the same, namely, anticyclonic (cyclonic) rotation with height is associated with warm (cold) air advection and isentropic ascent (descent). Here, the gradient wind formulation is used to calculate the thermal advection diagnostic (equation 15 in Tory 2014).

Comparison of mid-level winds with hourly rainfall in Figure 39 shows that areas of heavy rain coincide with winds that rotate anti-cyclonically with height, while the rain-free areas coincide with winds that rotate cyclonically with height. This agrees well with the thermal advection diagnostic shown in Figure 40, where positive warm air advection (WAA) or isentropic ascent (red in Figure 40) coincide with areas of heavy rain, while to the north negative WAA (blue in Figure 40) indicates cold air advection and rain-free areas. This comparison is valid for all three members analysed here, however as mentioned previously, they differ in location and strength of heavy rain and winds. This analysis also shows that the upper and lower-level circulations are not vertically aligned and that the tilt of the low is associated with ascent (corresponding to rain) and descent (no rain). Thus, localised differences in extreme wind and rain likely stem from different responses to upper-level and surface forcing and also to differences in tilting of the low core.

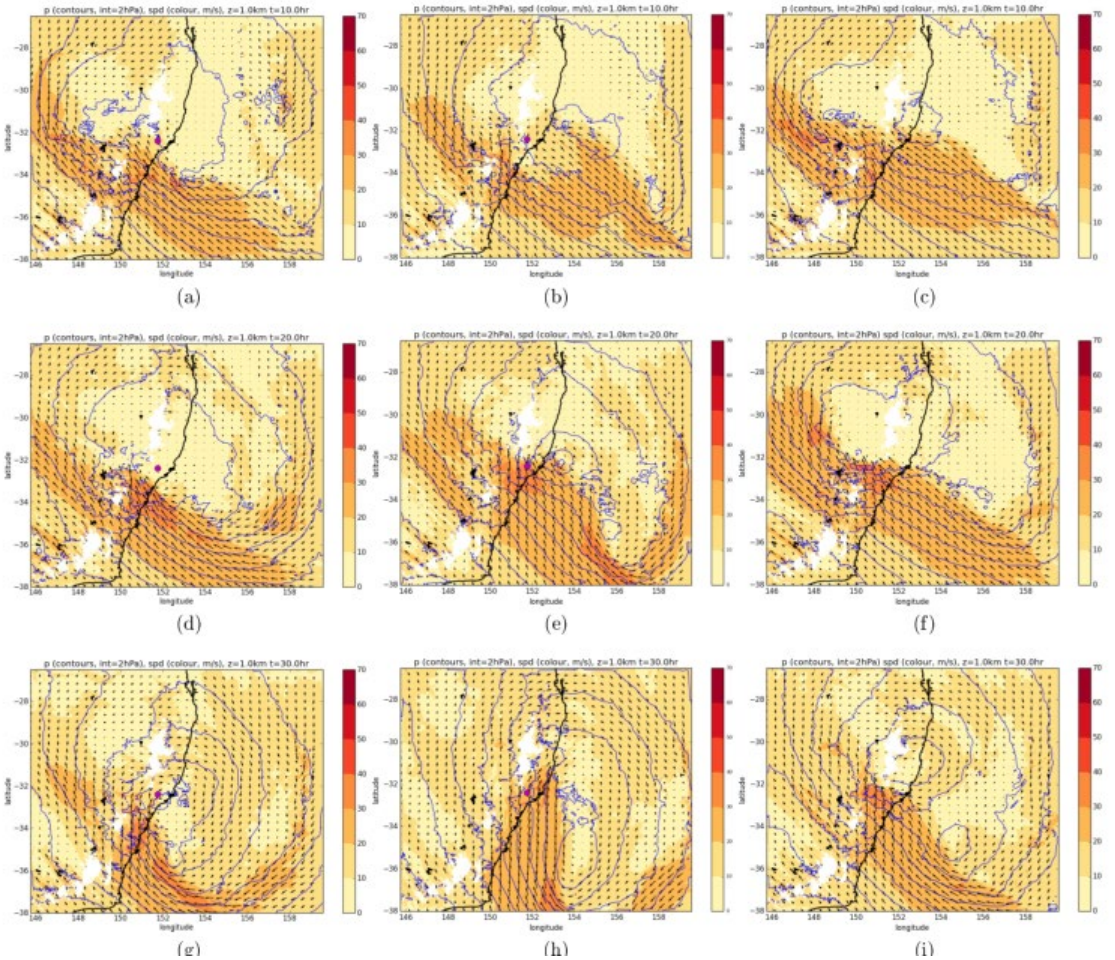

Figure 36: A sequence of 1 km winds (m/s) and pressure (contours, hPa) for (left) ensemble member 11, (middle) ensemble member 17 and (right) ensemble member 22 at (a,b,c) 1000 UTC 20 April 2015, (d,e,f) 2000 UTC 20 April 2015 and (g,h,i) 0600 UTC 21 April 2015.

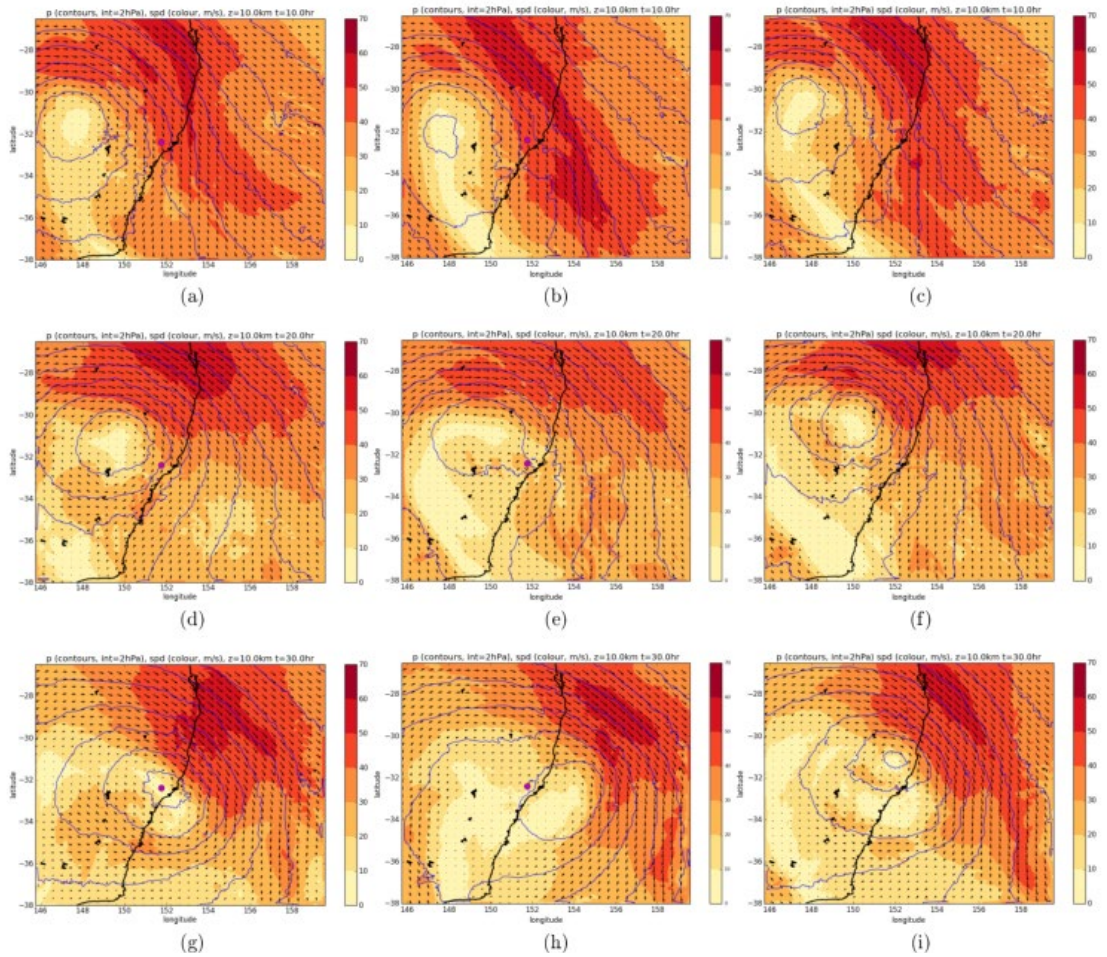

Figure 37: Same as Figure 36, except at 10 km.

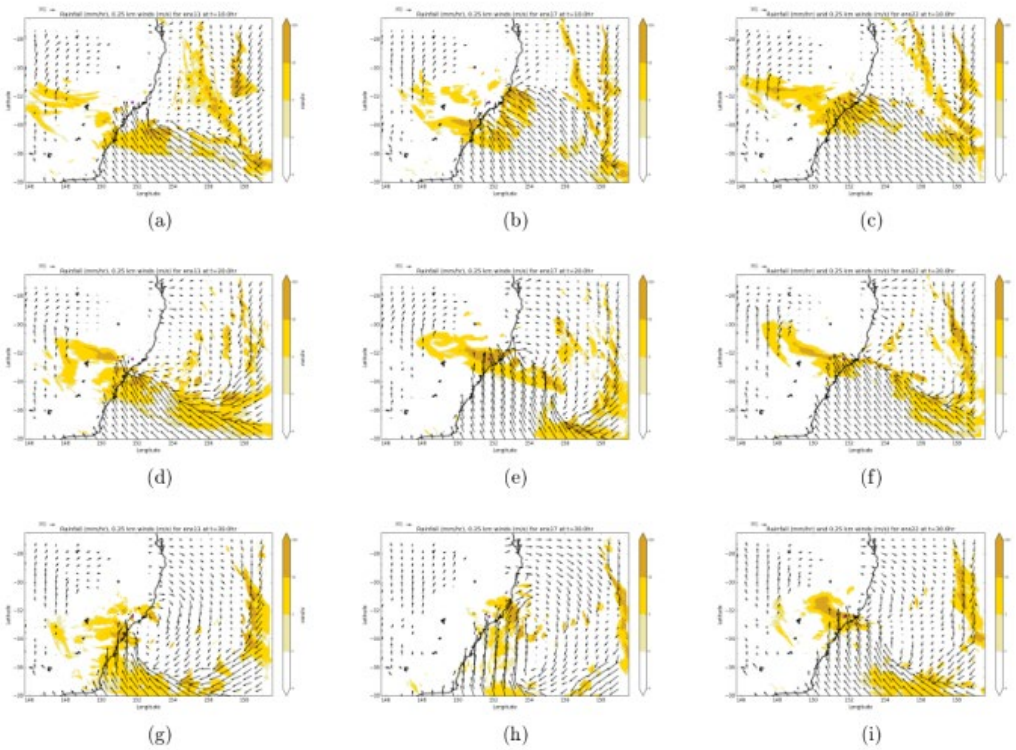

Figure 38. Same as Figure 36, except showing low-level wind (250-m wind vectors) and hourly rainfall (shading, mm).

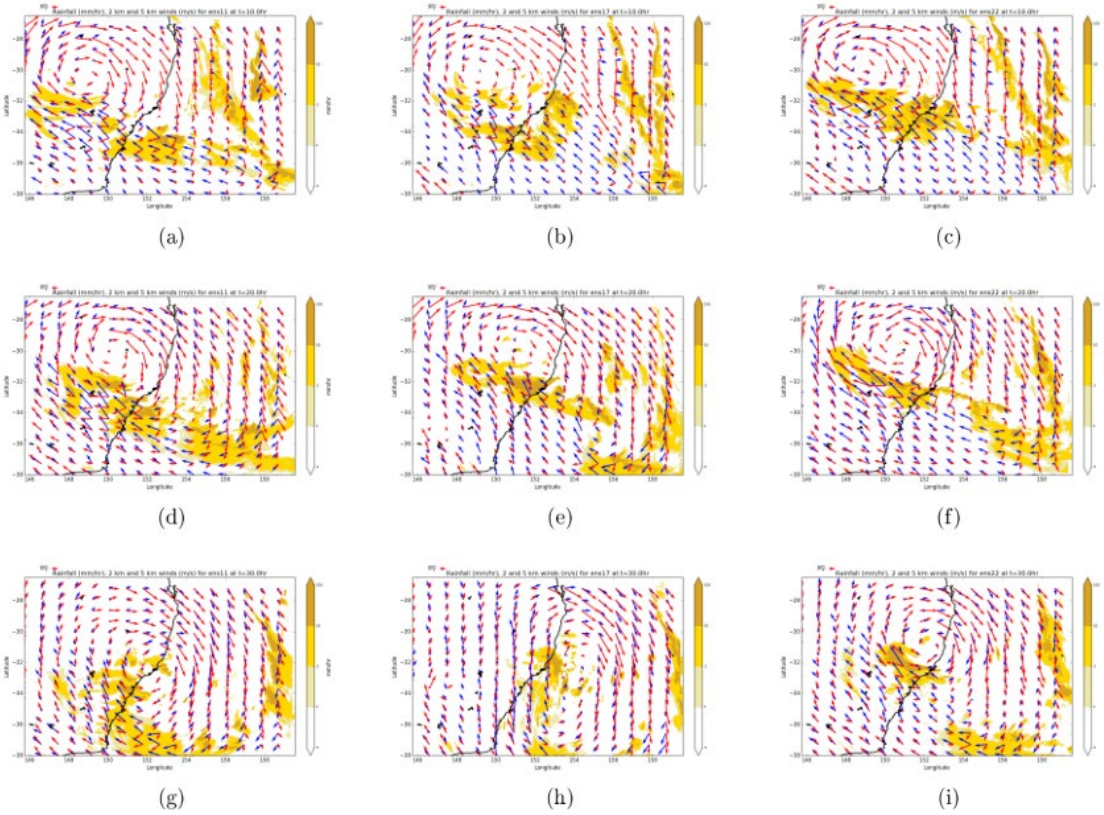

Figure 39. Same as Figure 36, except for 2 km (blue vectors) and 5 km (red vectors) winds, and hourly rainfall (shading, mm).

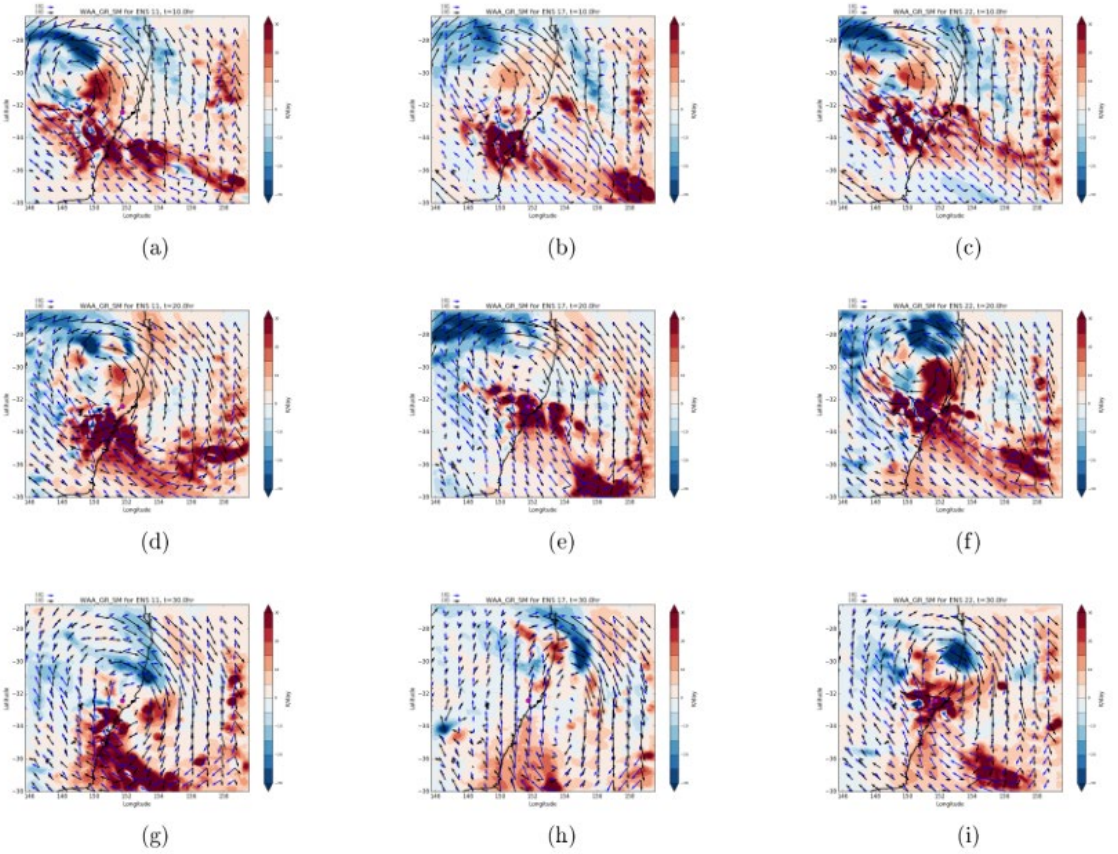

Figure 40. Same as Figure 36, except for thermal advection rainfall diagnostic (or WAA), with 1.5 km (blue vectors) and 5.5 km (black vectors) winds. Red indicates positive values, blue indicates negative values.

Summary and conclusions

An analysis of the East Coast Low (ECL) that affected the eastern coast of Australia between 20 and 23 April 2015 is presented, to get a better understanding of the dynamics of the storm as well as its predictability. For this, an ensemble consisting of 24 high-resolution simulations using the ACCESS model was prepared.

A comparison between the observed and forecast rainfall of the event (averaged over the 24 ensemble members) showed that the forecast rainfall is in good agreement with the observed rainfall. Even though peak rainfall in the Dungog area of the mean forecast was not as high as observed, the forecast probability of rainfall exceeding 100 mm showed that the ensemble identified Dungog as the area of substantial risk of extreme rainfall. Ensemble analysis also showed a large variation in total rainfall, with some members having more rain located to the south or east of the Dungog area. This indicated that a substantial part of the coast was at risk of significant rain.

Comparison of ensemble members 11, 17 and 22 showed that, common to most ECLs, the strongest surface winds occur to the south of the main low. Each also shows the strengthening of the north-westerly jet and the upper cut-off low west of the surface low. The intensification of the onshore low-level jet was associated with the cut-off low development and upper-level ridging that extended all the way to the surface, thus creating a strong near-surface pressure gradient and strong onshore winds. For all members, rain occurred along the low-level convergence lines to the south of the main low, and also on the eastern edge over the Tasman Sea. However, localised differences between members produced large differences in location and strength of extreme winds as well as in rainfall. While member 22 had an onshore jet directed towards Dungog and maximum rainfall accumulated around the Dungog area, the low level jet and maximum rainfall in member 11 were located further south. In contrast to these two members, member 17 had the strongest low developing offshore and maximum rainfall to the south and to the north-east of the Dungog area.

Comparison of the thermal advection rainfall diagnostic (calculated using gradient wind approximation) with mid-level winds and hourly rainfall showed the tilting of the low core is important and that areas of heavy rain coincided with anticyclonic rotation of winds with height and positive warm air advection (WAA), while rain free areas coincided with negative WAA and cyclonic rotation of winds with height. This analysis showed that this diagnostic approach works well with high-resolution data and could be used for analysing ECL development.

Practical implications

These results show how the use of ensembles in studying severe weather events such as ECLs provides a valuable insight into the processes that lead to the development and rapid intensification of these systems. It also provides the opportunity to study their predictability by analysing differences between the simulations. Ensemble prediction has advantages over deterministic methods, including that it is on average more accurate, and that it allows the calculation of probabilities and the identification of alternative scenarios. These advantages are evident in this study. In addition, a better understanding of the dynamics and

predictability of ECLs, as well as other high impact weather events, has important implications not only for forecasters, but also for emergency services preplanning and preparedness.

A CASE STUDY OF SOUTH AUSTRALIA'S SEVERE THUNDERSTORM AND TORNADO OUTBREAK (28 SEPTEMBER 2016)

Introduction

On 28 September 2016 one of the most significant thunderstorm outbreaks recorded in South Australia impacted central and eastern parts of the state. Multiple supercell thunderstorms were embedded in a Quasi-Linear Convective System (QLCS; Weisman and Trapp 2003) aligned with a strong cold front that was associated with an intense low-pressure system. The storms produced at least seven tornadoes, destructive wind gusts, large hail and intense rainfall. Transmission lines were brought down in four different locations, which contributed to a state-wide power outage.

Accurate prediction and understanding of tornadoes and other hazards associated with severe thunderstorms is very important, for timely preparation and announcement of warnings. By conducting high-resolution simulations, this study aims to offer a better understanding of the meteorology of the South Australia's thunderstorm and tornado outbreak. It also contributes to improving knowledge of how to best predict similar severe weather events, which in turn enables better risk management and preparedness for such events. Updraft helicity, a severe storm surrogate that indicates the potential for updraft rotation in simulated storms is used to investigate the ability of the model to predict supercell and tornado likelihood.

Modelling set-up

Here, a case study of the September 28 2016 event is conducted using high-resolution nested simulations of the Australian Community Climate and Earth-System Simulator (ACCESS) model. The model consists of a global model run (17 km that is nested down to 4.0 km, 1.5 km and 400 m (Figure 41), size chosen to capture the area where seven tornadoes were reported. Each domain has 80 vertical levels and the model top for nested domains is 38.5 km.

All simulations use Regional Atmosphere Mid-latitude first release (hereafter RA1M) science configuration (Bush et al. 2019), to assess model performance over the South Australia domain, and simulation results are compared to radar imagery. All simulations were initialised at 1500 UTC 27 September 2016 (+ 9.5 hours local time) and ran for 48 hours.

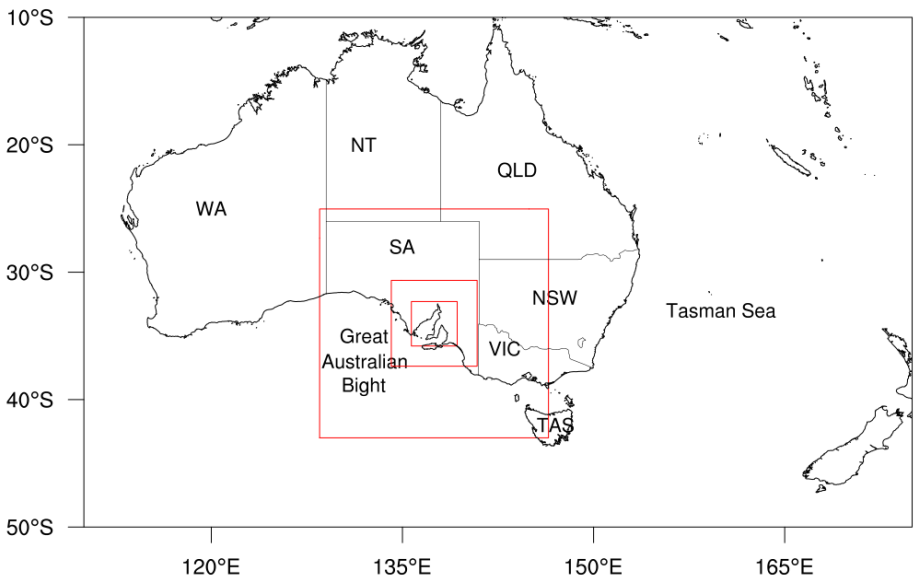

Figure 41: Outline of model domains, with larger domain having horizontal resolution of 4 km, and smaller domains with horizontal resolution of 1.5 km and 400 m.

Results

To investigate how well model represents severe thunderstorms, observed radar reflectivity (Figure 42, left) is compared to simulated reflectivity (Figure 42, right) from the 1.5 km model at 0600 UTC 28 September 2016. Figure 42a shows the line of thunderstorms that were associated with large hail and several reported tornadoes (Bureau of Meteorology 2016), and the overall timing and location of severe thunderstorms is captured well by the model (Figure 42b). Individual supercells are not depicted by the model as grid lengths on order of 1 km are not able to resolve supercell or tornado-like signatures.

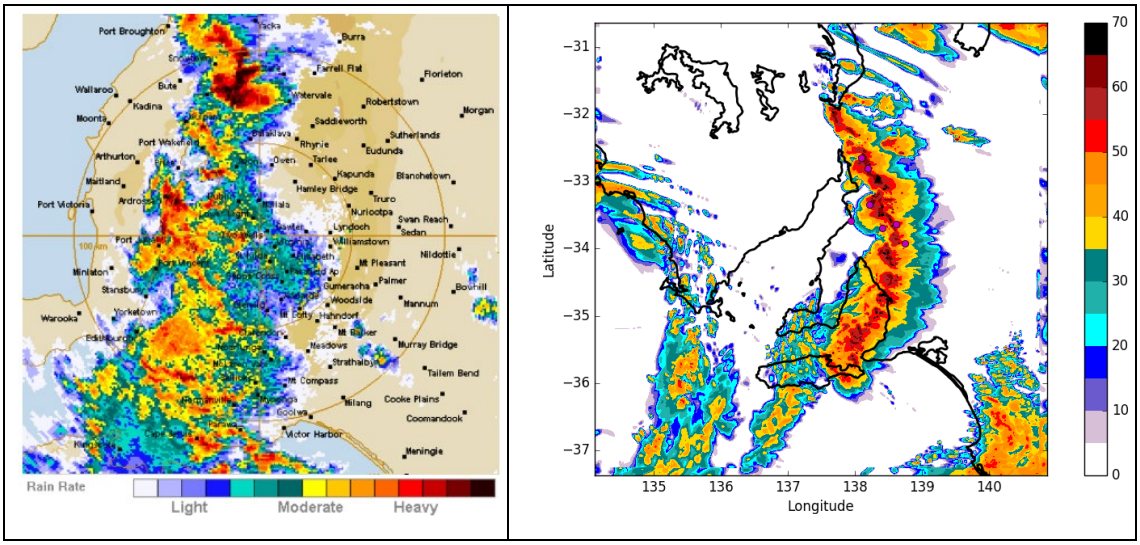

Figure 42: (Left) Radar reflectivity (dBZ) and (right) simulated radar reflectivity at 0600 UTC 28 September 2016. The magenta dots are the observed locations of the tornadoes.

Figure 43 shows simulated reflectivity from the 400-m simulation, where northern cells (black box in Figure 43) are stronger and better defined than the southern part of the convective system, indicated by differences in the simulated reflectivity. A close-up view of the simulated reflectivity and vertical velocity (Figure 44) in the northern part of the system (black box in Figure 43) reveals a hook-echo feature and a curved updraft that coincide with the location of one of the reported tornadoes (northernmost white dot in Figure 44). This indicates the presence of a mesocyclone (Davies-Jones 2015) at this location. The reported tornado is estimated to have started at approximately 0615 UTC 28 September 2016 (Bureau of Meteorology 2016, Sgarbossa et al. 2018), which is about 15 min later than in the simulation. This shows that the mesocyclone and features that indicate the possibility of a tornado are still reasonably well captured with the 400-m simulation.

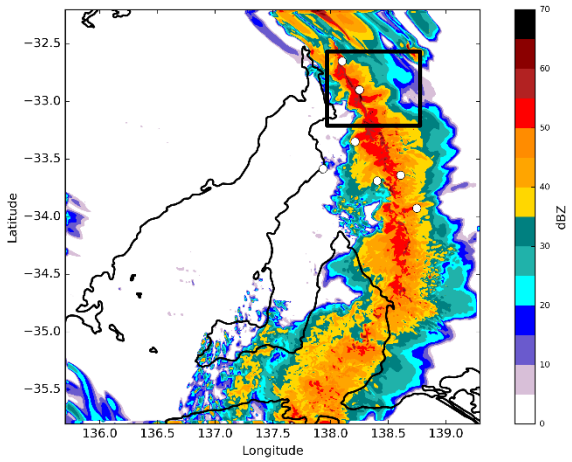

Figure 43: 0600 UTC 28 September 2016 simulated radar reflectivity (dBZ) at 2 km height for the 400-m simulation. The white dots are the observed locations of the tornadoes and the black box denotes the area shown in Figure 44.

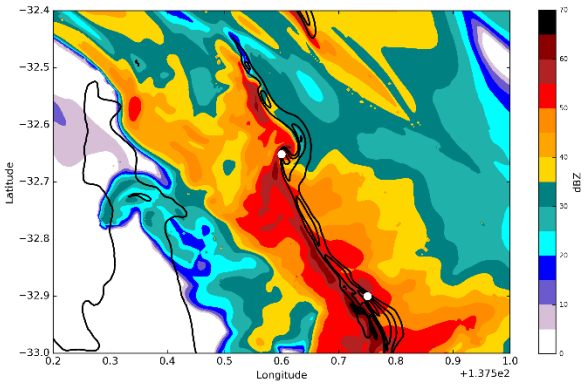

Figure 44: As per Figure 43, only for area denoted by the black box in Figure 43.

Based on Kain et al. (2008), a model diagnostic field (updraft helicity) is used to investigate the ability of the model to identify the potential for supercell thunderstorms, within which tornadoes may form. Updraft helicity (UH) is the product of vertical vorticity and vertical velocity, integrated vertically from 2 km to 5 km:

$$\boldsymbol{UH = \int_{2km}^{5km} w\left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}\right) dz},$$

where w is vertical velocity ($m\ s^{-1}$), $\left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}\right)$ is vertical vorticity (s^{-1}) and z is height (m). UH is negative for the Southern Hemisphere. While UH can be computed as an instantaneous value at a single model output time, Kain et al. (2010) developed hourly maximum UH that tracks the maximum value of the diagnostic at every grid point at any model time step within the previous hour (e.g., Clark et al. 2012, 2013; Sobash et al. 2016). Here, UH is computed as the output diagnostic in the Met Office UM model.

Figure 45 shows hourly minimum UH for the period between 0500-0600 UTC and 0600-0700 UTC 28 September 2016, which coincides with the time period of the reported tornadoes(Bureau of Meteorology 2016). It shows elongated swaths of hourly minimum UH exceeding -120 m^2s^{-2} and reaching values < -400 m^2s^{-2} in close proximity to the observed tornadoes (magenta dots in Figure 45). A long and coherent, vortex-like swath of hourly minimum UH < -500 m^2s^{-2} coincides with the location of the northernmost observed tornado, where mesocyclone was identified in simulated reflectivity and vertical velocity (cf. Figure 44). This indicates that the use of UH as a diagnostic field could provide useful guidance for identifying potential for tornado formation.

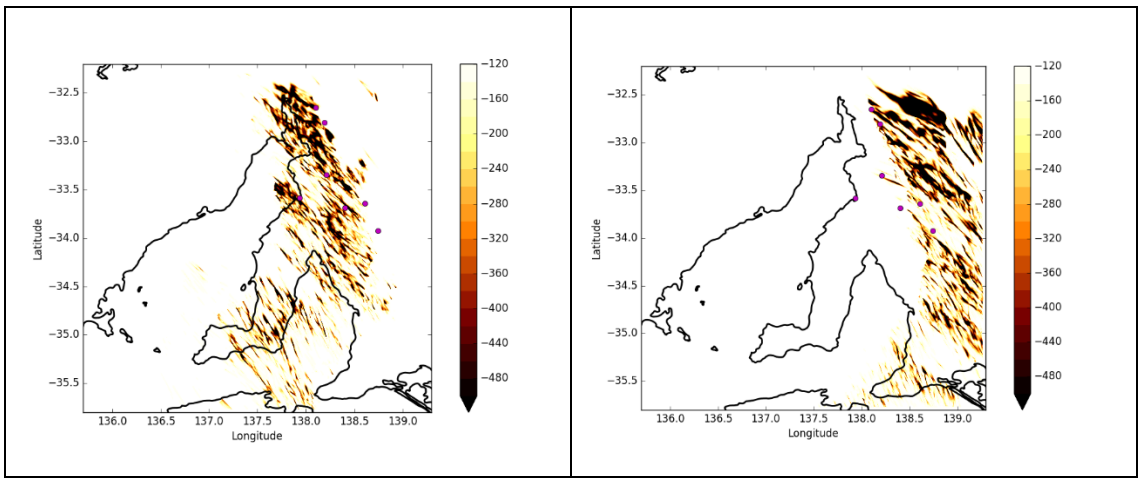

Figure 45: Hourly minimum updraft helicity (UH, $m^2\ s^{-2}$) for the period (left) 0500-0600 UTC 28 September 2016 and (right) 0600-0700 UTC 28 September 2016. The magenta dots denote the approximate location of observed tornadoes.

$$OW = \xi^2 - (E^2 + F^2) = \left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}\right)^2 - \left\{\left(\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y}\right)^2 + \left(\frac{\partial v}{\partial x} + \frac{\partial u}{\partial y}\right)^2\right\},$$

where ξ is the vertical vorticity, E is the stretching deformation, and F is the shearing deformation (Okubo 1970; Weiss 1991; Markowski et al. 2011). The OW parameter highlights flow regions where rotation dominates over strain, and therefore can be used to identify rotation (or vortices) associated with tornadic storms (e.g., Markowski et al. 2011; Coffer and Parker 2016).

Figure 46 shows the OW parameter calculated as per above equation, using layer averaging between 1 km and 4 km for the time interval between 0528 UTC and 0628 UTC 28 September 2016. It shows a coherent track of positive OW parameter that coincides with the location of the northernmost reported tornado, thus clearly identifying mesocyclone in this simulation. This analysis shows that OW parameter identifies mesocyclone and rotation in simulated storms and gives a clear indication of a tornado path. This suggests it can be used as an alternative diagnostic alongside UH to assess tornado potential, thus potentially reducing false alarms.

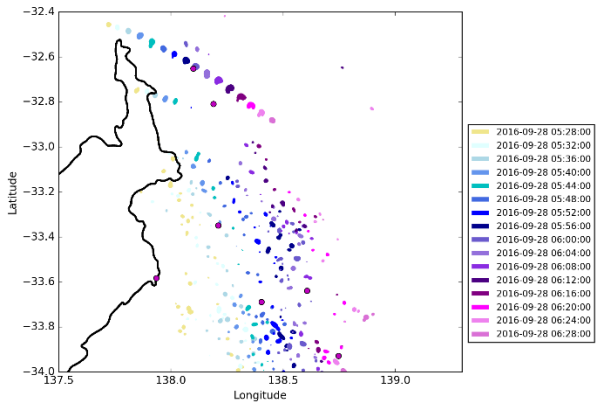

Figure 46: The Okubo-Weiss parameter (OW; positive values only), calculated as a 1 km-4 km layer average for the time interval between 0528 UTC and 0628 UTC 28 September 2016. The magenta dots denote the approximate location of observed tornadoes

Practical Implications

THE BLUE MOUNTAINS FIRES OF OCTOBER 2013

Summary

In October 2013, several fires broke out in the Blue Mountains of New South Wales, Australia, and were devastating to the region. Of particular interest is the State Mine Fire, which was lit on 16 October, then spread rapidly eastward during the afternoon of 17 October. To investigate the meteorological conditions during that day, the Bureau of Meteorology's numerical weather prediction (NWP) model was run at very high resolution, with a horizontal grid spacing of about 440 m, to simulate the weather conditions on the day. Several mesoscale features were identified which could account for the rapid fire spread. The first was a mountain wave structure with very strong downslope winds developing over the fire ground during the afternoon of 17 October, accompanied by hot and dry atmospheric conditions. The second was a slot of dry air with an approximate width of 50 km that passed over the fire ground around and shortly after midday which likely exacerbated the fires. The third was an undular bore that formed during the night of 16 October, which was related to the formation of the dry slot. The formation and evolution of these mesoscale features are described, and existing theories are considered to explain these phenomena.

Introduction

In October 2013, a number of significant fires broke out and created havoc in New South Wales, Australia. The most intensive fire activity occurred between 13 and 26 October, when there were 627 incidents and an estimated area of 164054 ha burnt. In particular, more than 200 houses in the Blue Mountains were damaged during the afternoon of 17 October, although, thankfully, no lives were lost (New South Wales Rural Fire Service 2014).

The fires in the Blue Mountains Region proper consist of the State Mine Fire and the Mount York Road Fire. The State Mine Fire was ignited on 16 October and spread to the south-east during the night and the morning of 17 October. Later, the fire spread very rapidly to the east from about noon. According to the analysis performed internally by New South Wales Rural Fire Services, the estimated burnt area increased from 1036 ha at 11:56 am to 12436 ha by 9:46 pm local time. The Mount York Road Fire was ignited on the afternoon of 17 October. While far less extensive than the State Mine Fire, it had an estimated burnt area of 492 ha by 7:46 am (AEDT) on 18 October (Figure 47).

The areal extension of the State Mine Fire during the afternoon of 17 October was phenomenal – an increase of 11400 ha (more than 10-fold) in less than 10 hours (Table 3). It is of both practical and academic interest to investigate whether the meteorological conditions had been conducive to the rapid spread of the fire during that day; if so, what were those meteorological conditions or features? We have tried to answer this question by performing a simulation with an operational NWP system but with much higher horizontal resolution, and have identified several mesoscale features of interest.

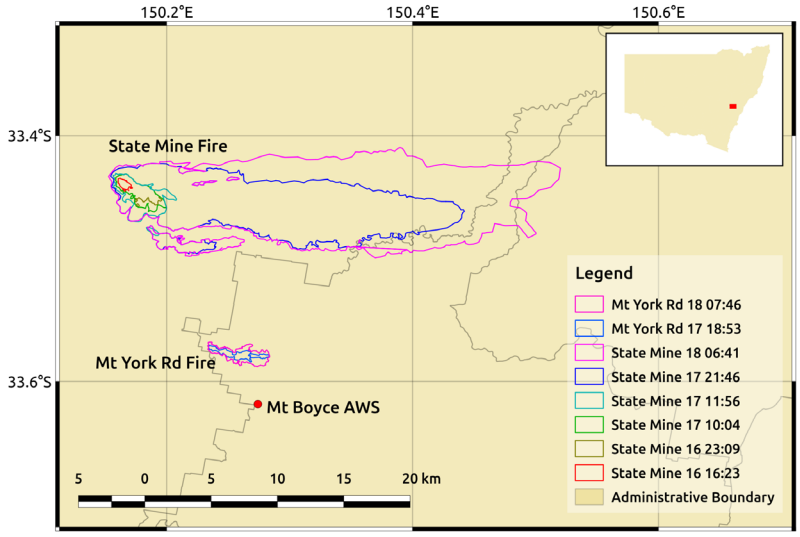

Figure 47. The progression of the perimeters of the State Mine Fire and the Mount York Road Fire. The times are in Australian Eastern Daylight Time (AEDT), which is 11 hours ahead of UTC. Location of Mount Boyce AWS is shown by the red dot. The inset shows the map of New South Wales, Australia, with the red box representing the areal extent of this map.

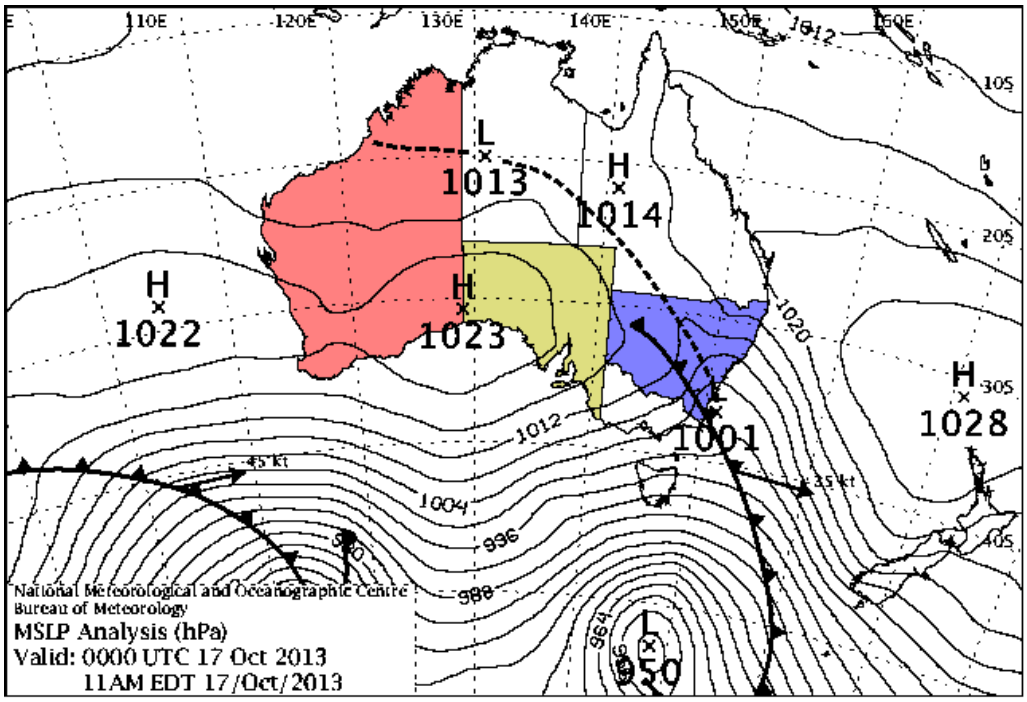

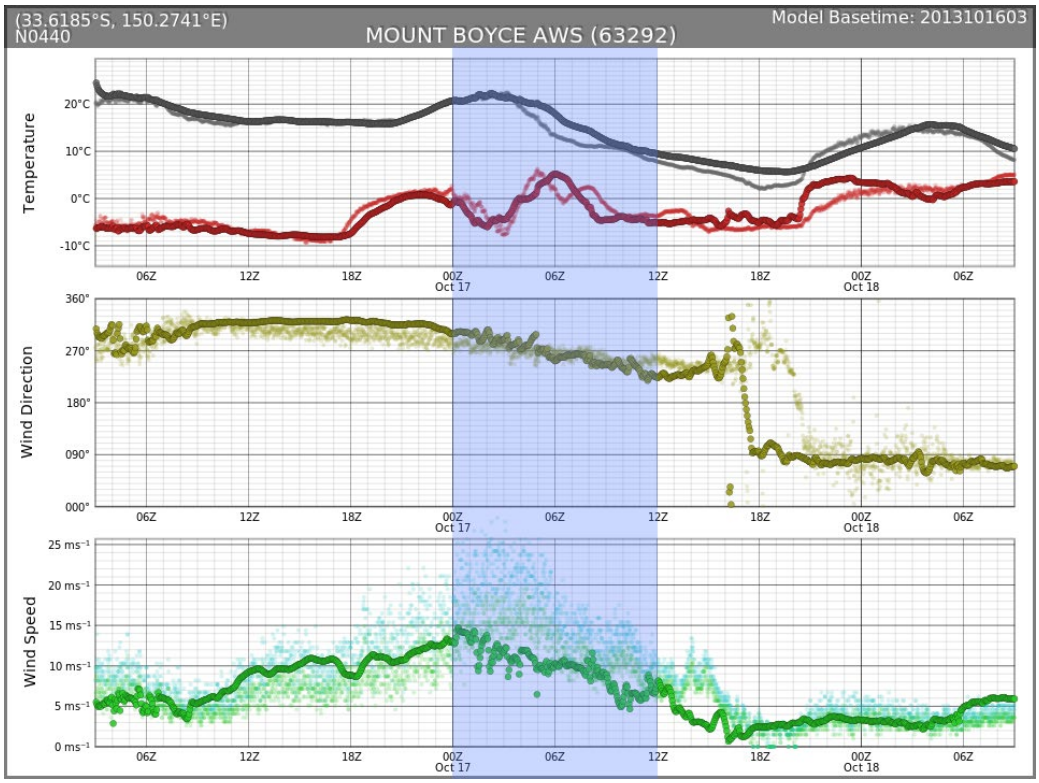

Figure 49. Meteogram of 2-metre temperature, 2-metre dew point, 10-metre wind direction, 10-metre wind speed and 10-metre wind gust at Mount Boyce with AWS observational data and N0440 model data. The solid dots in the foreground are NWP simulation data in 5-minute resolution. The fainter dots are AWS observation data in 1-minute resolution. Black dots are temperatures, red dots are dew point temperatures, yellow dots are wind directions, green dots are wind speeds and cyan dots are wind gusts. The translucent blue box highlights the time during the rapid spread of the State Mine Fire.

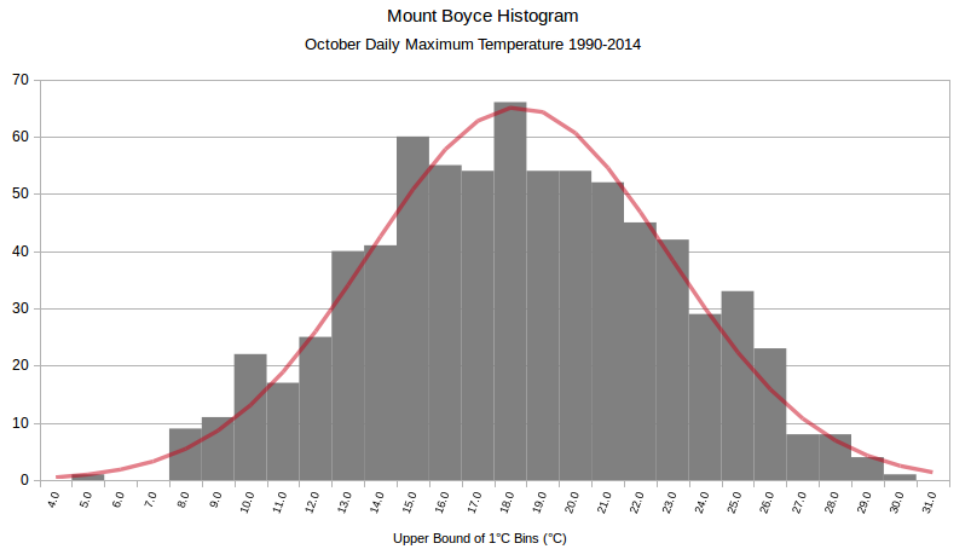

Figure 50. Histogram of daily maximum temperatures at Mount Boyce AWS in October from 1990 to 2004. The red curve is the fitted Gaussian distribution.

Table 3. Fire Perimeters.

Fire	Date	Local time (AEDT)	Estimated burnt area (ha)
State Mine Fire	16 October 2013	4:23 pm	65
		23:09 pm	433
	17 October 2013	10:04 am	578
		11:56 am	1036
		9:56 pm	12436
	18 October 2013	6:41 am	22780
Mount York Road Fire	17 October 2013	6:17 pm	183
		6:53 pm	201
		7:46 am	492
	18 October 2013	9:30 am	623

Table 4. Comparison of mean monthly rainfall and Year 2013 monthly rainfall at Mount Boyce from July to October.

Month	Mean	Year 2013
July	40.6mm	13.0mm
August	54.8mm	15.4mm
September	53.6mm	40.6mm
October	64.5mm	12.8mm

Table 5. Various measures indicating dryness in the Blue Mountains area in October 2013 and the lead up period. All entries are values at Mount Boyce interpolated from the AWRA L4.5 database at 0:25_ horizontal resolution, for October 2013 (unless otherwise stated) against previous and subsequent Octobers. Effective Rainfall is the difference between rainfall and evapotranspiration.

Dryness Measure	Value
Top Layer (0 cm to 10 cm) Soil Moisture	4.29 percentile
Mid Layer (10 cm to 100 cm) Soil Moisture	2.88 percentile
Rainfall	7.31 percentile
Mid Layer (10 cm to 100 cm) Soil Moisture Anomaly	-41.49mm
Effective Rainfall Anomaly	-41.55mm
Effective Rainfall Anomaly August to October 2013	-92.75mm

Fire ground conditions

The several months leading up to October 2013 were quite dry in the Blue Mountains area. Rainfall recorded at Mount Boyce, a hill in the Blue Mountains area with a Bureau of Meteorology maintained Automatic Weather Station (AWS) that is closest to the Blue Mountains fire grounds (Figure 47), had been substantially below average since July (Table 4). The soil dryness from the Australian Water Resources Assessment System (AWRA) (Vase et al. 2013) analysis also shows that it was exceptionally dry in the Blue Mountains area in the lead up of the fires (Table 5). The drought factor at Mount Boyce, interpolated from the Bureau of Meteorology's operational gridded analysis (Finkele et al. 2006), was 8.7.

Figure 48 shows the operational synoptic mean sea level pressure chart at 0000 UTC 17 October 2013. An anticyclone was centred over the Tasman Sea north of New Zealand, and a ridge of high pressure extended from the west, covering Western Australia and South Australia. A deep low pressure system was centred in the Southern Ocean. It extended a cold front northwards which was crossing New South Wales during that day, preceded by a prefrontal trough which extended to northern Australia.

Figure 49 shows the meteogram from 0300 UTC 16 October 2013 to 0900 UTC 18 October 2013 at Mount Boyce. The cold front and the prefrontal trough manifested at Mount Boyce as1 a gradual backing from north-westerlies during the morning to south-westerlies by the afternoon on 17 October 2013. Several smaller wind directional changes were embedded in the larger scale backing of the winds.

The maximum temperature of the day recorded by the Mount Boyce AWS was 22.5°C (recorded at 0311 UTC and 0313 UTC). Considering the altitude of the AWS at 1080m above mean sea level and the time of the year, it was a somewhat hot day. Checking against the climatological data (1990 to 2014, inclusive) for October daily maximum temperature at Mount Boyce AWS, 22.5°C converts to a percentile rank of 82.2. Figure 50 shows the histogram of October daily maximum temperatures at Mount Boyce which fits very well with a Gaussian distribution with a mean of 17.8°C and a standard deviation of 4.6°C. The goodness of fit can be illustrated by checking that 22.5°C is about 1.02 standard deviations above the mean, which is equivalent to a percentile rank of about 84.6.

The humidity on 16 October in the Blue Mountains area was very low, with dew point temperatures recorded at Mount Boyce AWS consistently below -5°C. From about 1800 UTC 16 October 2013, the dew point temperatures increased to around 0°C, before decreasing sharply during the early afternoon of 17 October, to a minimum of -7.6°C around 0300 UTC 17 October. The dew point temperature then increased again, quite abruptly, from about 0°C to 5°C from 0400 UTC to 0900 UTC. It was also windy in the Blue Mountains area on 17 October. The Mount Boyce AWS meteogram shows clearly that winds picked up during the afternoon, with 10-minute mean winds predominantly above 12 m s^{-1} between about 0000 UTC and 0600 UTC. During this time, the maximum 10-minute wind recorded was 17.4 m s^{-1}, the maximum 1-minute wind recorded was 20.6 m s^{-1} at 0243 UTC while the maximum gust recorded was 28.3 m s^{-1} at 0237 UTC. Putting in the drought

factor (8.6), daily maximum temperature (22.5°C), lowest dew point temperature (-7.6°C) and highest 10-minute wind speed during the afternoon (17.4 m s^{-1}) into the calculation of FFDI yields a number of 63.2, which is categorised as a Forest Fire Danger Rating (FFDR) of "Severe".

Numerical simulation

Configuration

The simulation was performed using the Australian Community Climate and Earth-System Simulator (ACCESS), involving a sequence of nested limited-area models embedded within ACCESS global model runs. The numerical weather prediction (NWP) modelling framework is closely related to the Bureau of Meteorology's operational NWP arrangements but run at higher resolutions than is currently employed in real-time operations. The central core of the ACCESS configuration is the UK Met Office's Unified Model (version 8.5). It is a non-hydrostatic model with an Arakawa-C grid in the horizontal and a Charney-Phillips grid in the vertical (Puri et al. 2013).

The simulation was initialised at 0300 UTC 16 October 2013 and run for 54 hours from the initialisation time. Global initial conditions were real-time analyses prepared by the Bureau of Meteorology's National Operations Centre (BNOC). Within the global model run, which has a grid spacing of around 40 km at the equator), three nested limited-area model runs, named "N4000", "N1300" and "N0440", were embedded with grid spacing of approximately 4 km, 1.3 km and 440m respectively (Figure 51). The nested grids were regular latitude-longitude grids, with a larger longitudinal grid spacing (as measured in degrees of longitude) than latitudinal grid spacing, so as to make the central grid rectangles approximately square, accounting for the meridional convergence. Each stage of the simulations (global + nested) used 70 model levels in the vertical direction. These levels are terrain-following in the lowest levels, transitioning smoothly to constant heights above mean sea level in the highest levels. The global model had a model top of 80 km, while the three nested models had model tops of 40 km. The lowest temperature, humidity, and vertical motion levels in the three nested models were 5m above the surface, while the lowest horizontal winds were 2.5m above the surface. The 1.3-km and 440-m models used a three-dimensional sub-grid turbulence scheme of Smagorinsky type, while the 4-km model used the planetary boundary layer parameterisation.

Evaluation of Simulation near the fire ground

Comparison of the NWP model surface and AWS observations at Mount Boyce on 17 October (Figure 49) shows that the conditions at the AWS were reasonably well simulated. The maximum temperature of the day was very well captured, although the cooling associated with the wind change from about 0300 UTC to 0400 UTC in the simulation is not as quick as observed. The simulated minimum temperature on the morning of 18 October was also about 3°C to 4°C warmer than observed. In general, the simulated dew point temperature time series matched quite well with the AWS observations, specifically, the very low humidity on 16 October, an increase of dew point temperatures from about 1800 UTC 16 October, and a period of abrupt drying around 0300 UTC (albeit the simulation

had this early by about 1 hour). A secondary dryness around 0700 UTC in the AWS observations was, however, missed by the simulation.

The biggest inaccuracy is that the model simulated 10-metre winds did not fully capture the high winds observed during the afternoon of 17 October. Between 0000 UTC and 0600 UTC 17 October, when the AWS observations indicated an increase of wind speeds, the NWP model winds showed a slight decline. The simulation did pick up increased windiness above the surface (Figure 55 and Figure 56), which indicates that this error is perhaps due to a near-surface bias in the model winds. Another error in the simulation is the early arrival, by three hours, of an easterly change during the early morning of 18 October (around 2100 UTC 17 October). This is likely due to the complexity of the development of a Southerly Buster off the New South Wales coast reaching the Blue Mountains area during the early morning of 18 October, and the simulation has not captured this complex feature very well. The wind before the arrival of this easterly change was calm, and the simulation did not capture the calm winds between 1700 UTC 16 October and 2100 UTC 16 October, which likely accounts for the error in the simulated minimum temperature through a lack of radiation inversion.

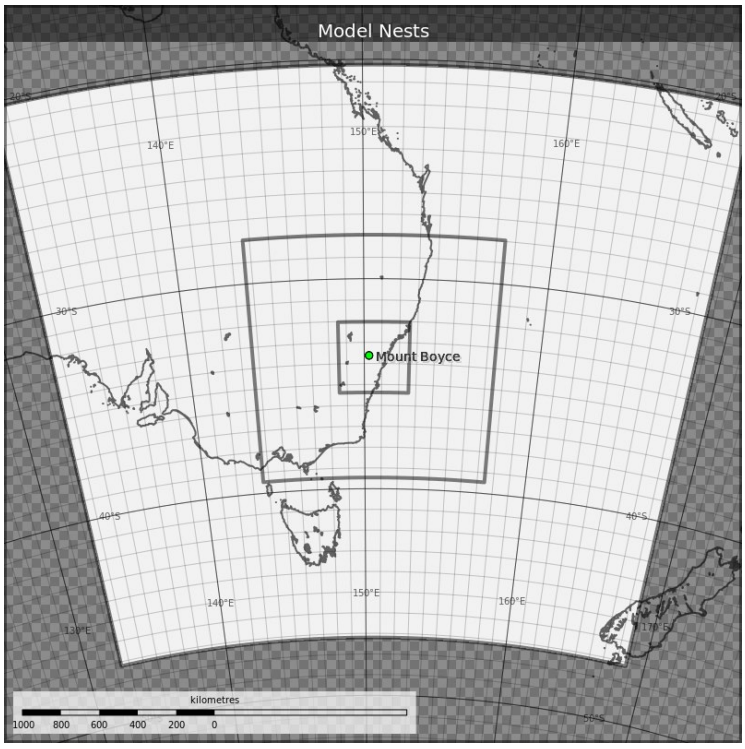

Figure 51. Domain boundaries for the three nested limited-area models, progressively from N4000 (largest) to N1300 and N0440 (smallest). N4000 is nested directly in the global model. Grid spacings are approximately 4:0 km (N4000), 1:3 km (N1300) and 440m (N0440). The dot in the centre marks the approximate location of Mount Boyce in the Blue Mountains.

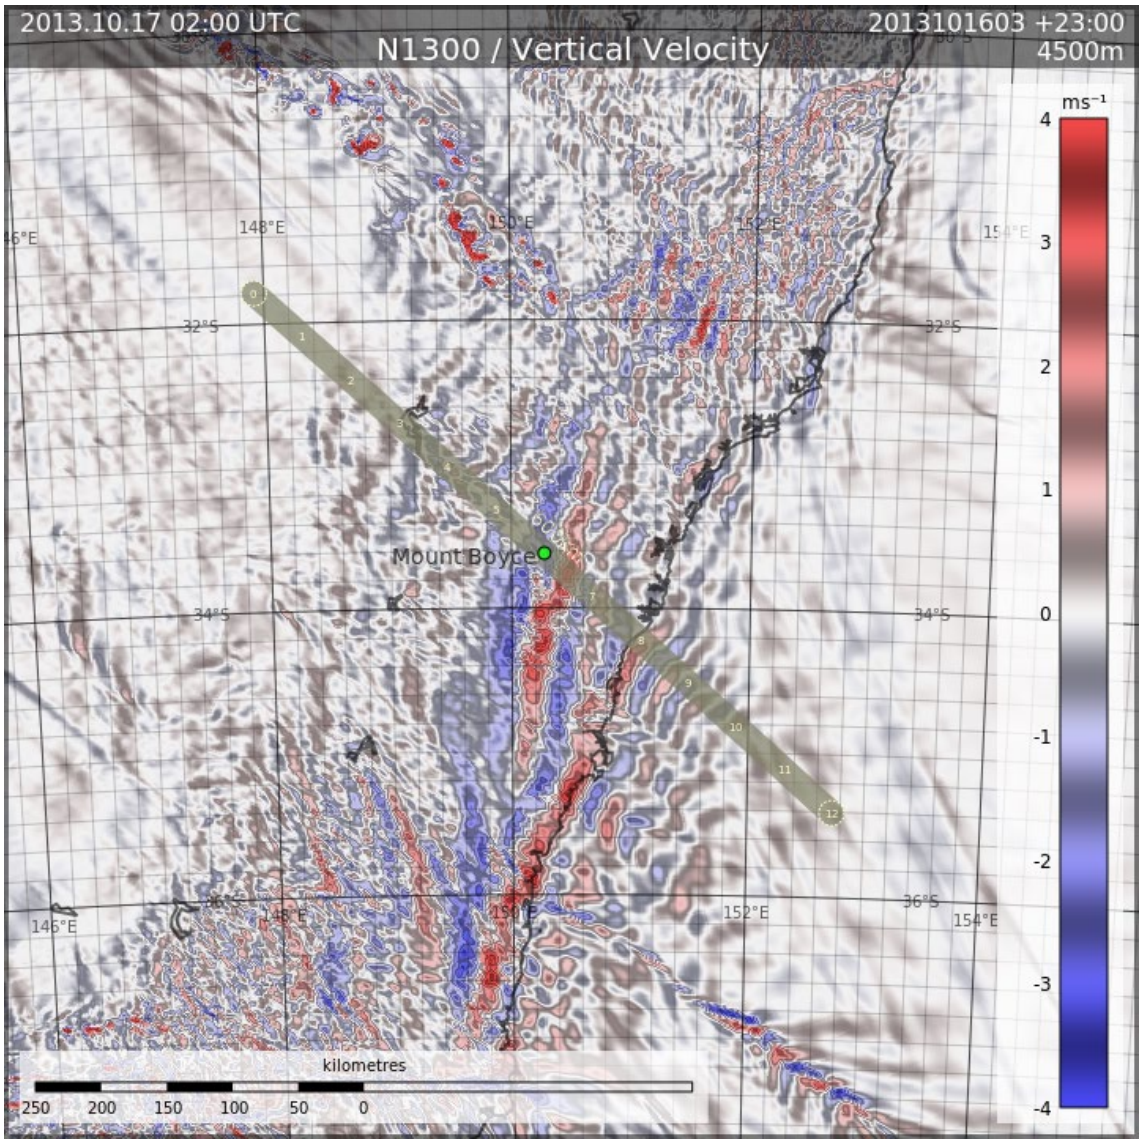

Figure 52. NWP model simulated 4500m above mean sea level vertical velocity at 0200 UTC 17 October 2013 from the N1300 nest. The thick line running northwest to southeast depicts the extent of the vertical cross section in Figure 54, Figure 55, Figure 56, and Figure 57.

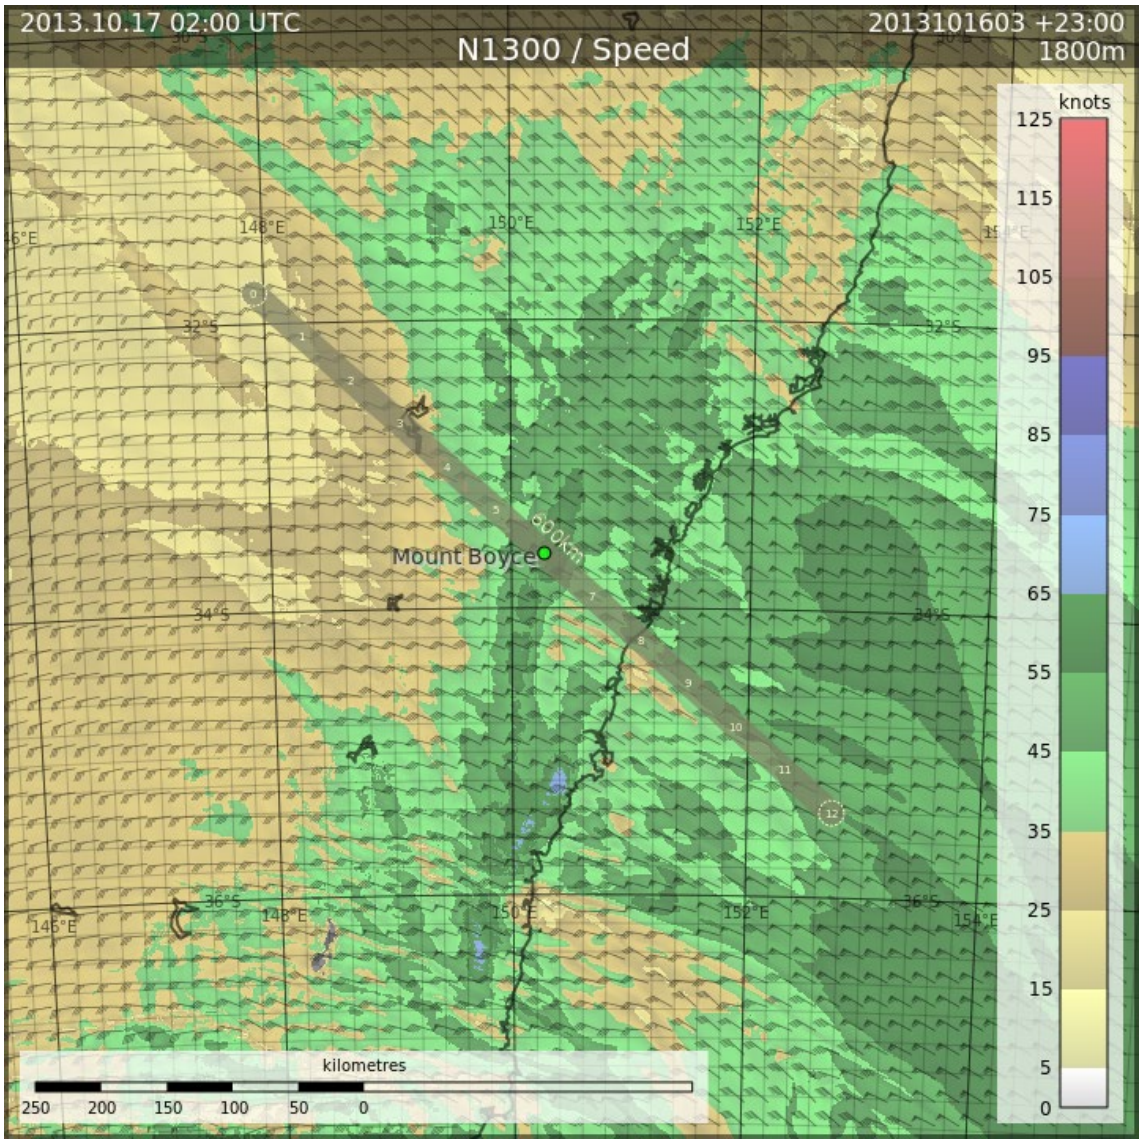

Figure 53: As in Figure 52 but for 1800m above mean sea level horizontal wind speed, also with wind barbs overlaid.

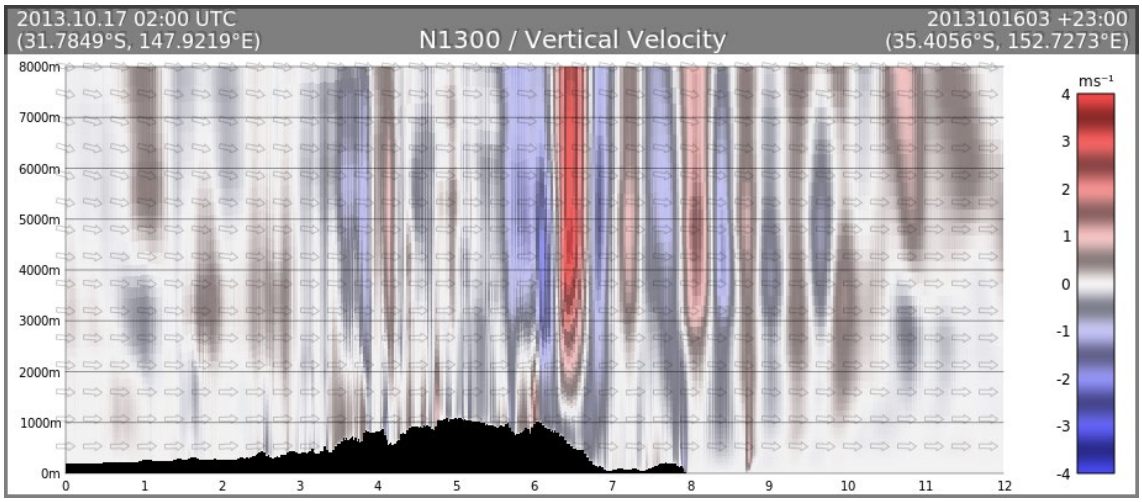

Figure 54. NWP model simulated vertical cross section of vertical velocity at 0200 UTC 17 October 2013 from the N1300 nest. This vertical cross section is 600 km long and is depicted in Figure 52. The numbers 0 to 12 in the horizontal axis are equidistant markers which are 50 km apart. Mount Boyce is located at the position marker 6. Faint arrows indicate the three-dimensional wind direction projected onto the vertical cross sectional plane.

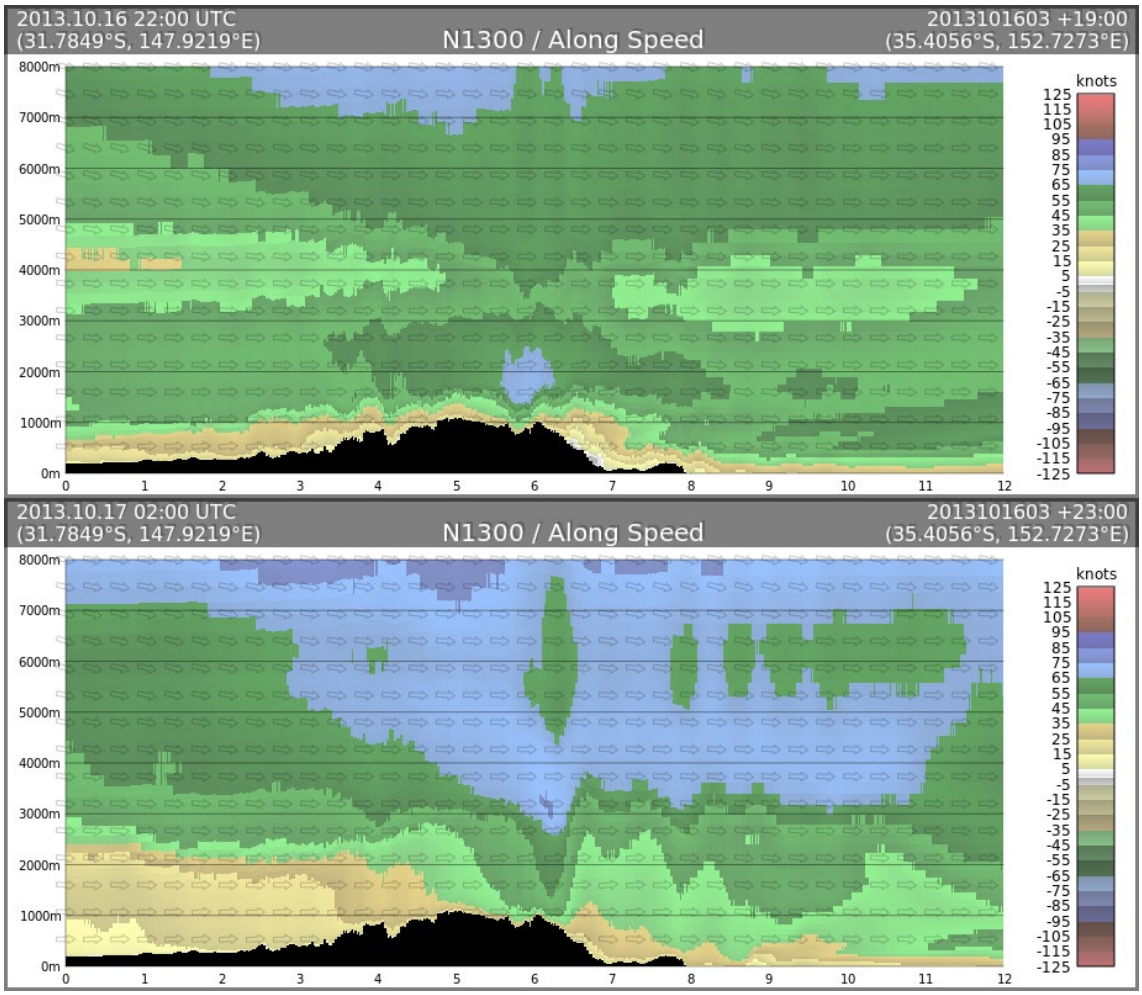

Figure 55. As in Figure 54 but for horizontal winds projected onto the vertical cross section plane for 2200 UTC 16 October and 0200 UTC 17 October.

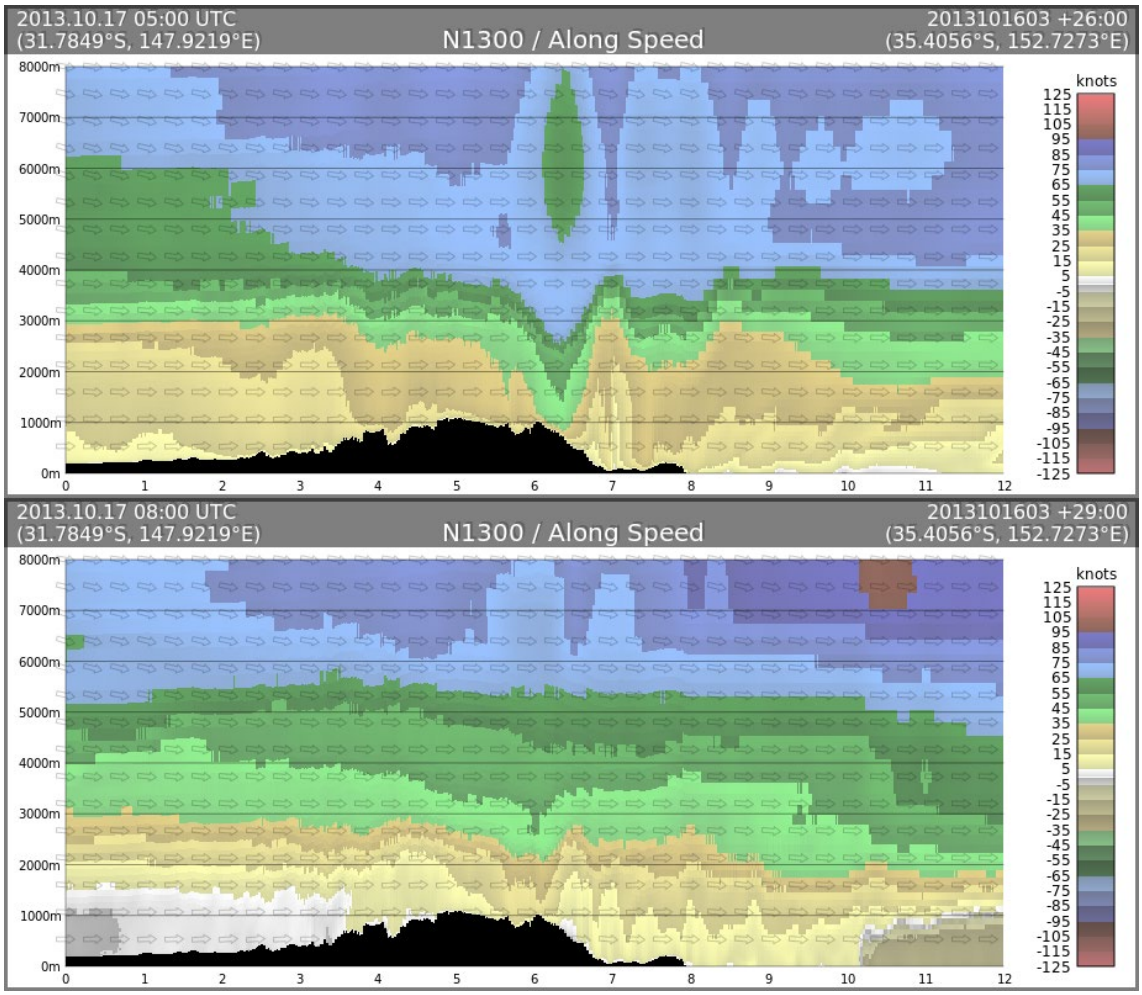

Figure 56. As in Figure 54 but for 0500 UTC 17 October and 0800 UTC 17 October.

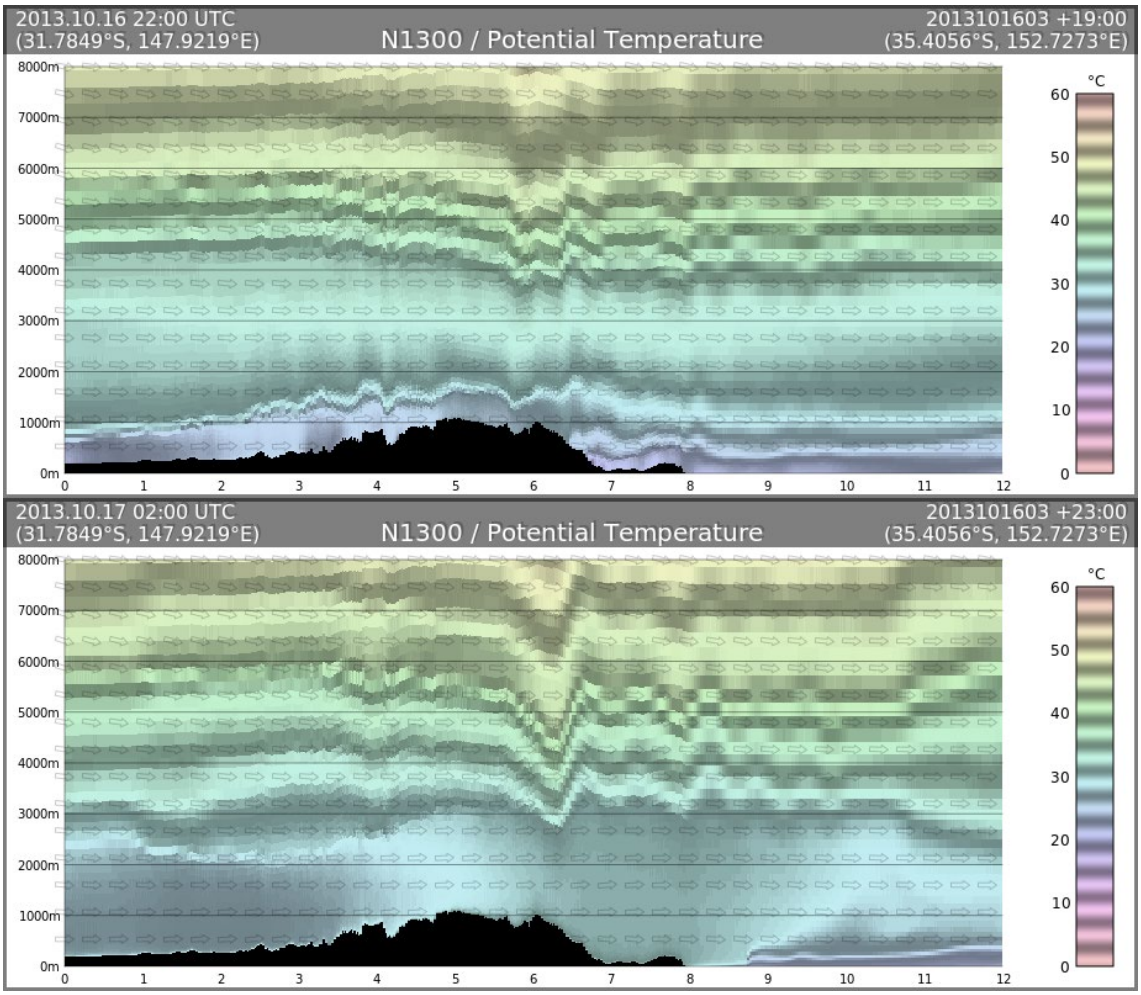

Figure 57. As in Figure 54 but for potential temperature.

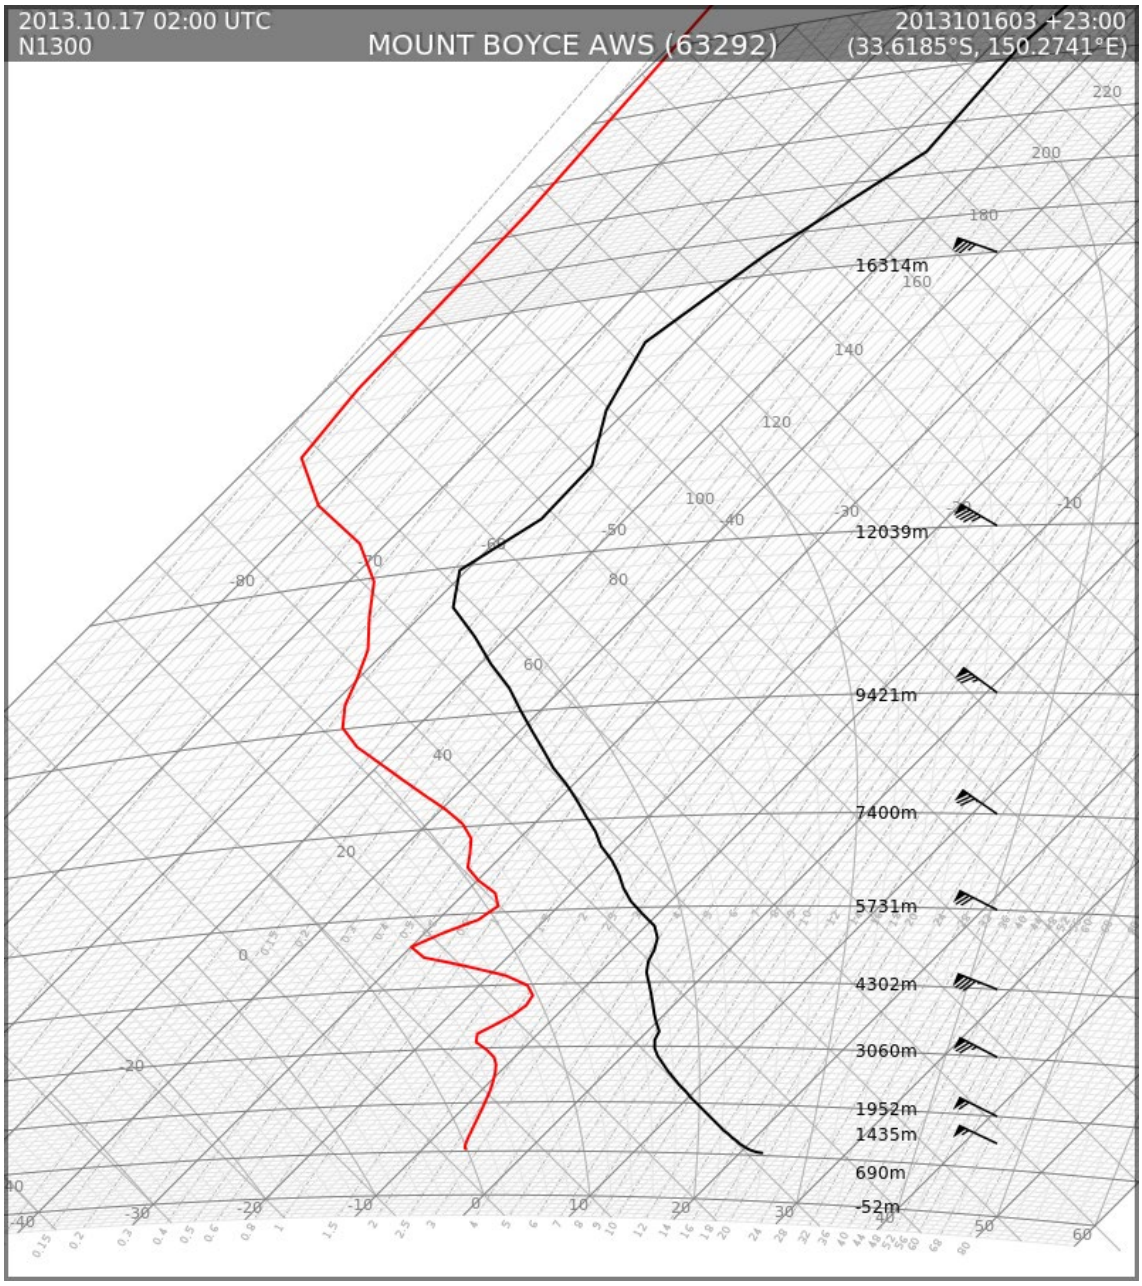

Figure 58. Tephigram of NWP model simulated vertical profile at Mount Boyce at 0200 UTC 17 October 2013 from N1300 nest.

Mesoscale features

Downslope winds

A mountain wave pattern was detected in the Blue Mountains region (Figure 52), and an area of enhanced low level downslope winds was present on the downslope side of the mountain ranges (Figure 53) in the simulation. A vertical cross section of vertical wind speed running northwest to southeast across Mount Boyce at 0200 UTC 17 October (Figure 54) clearly depicts a mountain wave pattern with an approximate horizontal wavelength of 35 km. Figure 55 and Figure 56 show several vertical cross sections of horizontal wind speed projected onto the same cross-sectional plane as in Figure 54. At 2200 UTC 16 October, there was a low level wind max near Mount Boyce (at position marker 6) (Figure 55, top panel) which was related to an area of low level temperature inversion (Figure 57, top panel). By around 0000 UTC 17 October, the low level inversion was broken and the low level wind max relaxed both upstream and downstream of Mount Boyce as the horizontal momentum mixed out by the vertical mixing. However it remained rather windy in the vicinity of Mount Boyce and the immediate areas to the lee of the terrain, forming the pattern of a marked area of enhanced horizontal winds near Mount Boyce and the immediate areas to the lee of the terrain (Figure 55, bottom panel). The model tephigram (Figure 58) located at Mount Boyce at 0200 UTC 17 October indicates that the atmospheric column was dry and adiabatic up to a height of about 3 km above mean sea level. Wind speeds were nearing 60 knots (30.8 m s^{-1}) within this unstable, dry air column.

This pattern persisted until around 0500 UTC 17 October (Figure 56, top panel), then somewhat weakened before mostly disappearing by about 0800 UTC 17 October (Figure 56, bottom panel). It is apparent that the strengthened downslope winds were conducive to the rapid fire progression during the afternoon of 17 October. While in the previous section it was noted that the NWP simulated 10-m winds were much less that the observed winds, the existence of an area of increased winds in the upper levels about and in the lee of the Blue Mountains ranges (as depicted by the form of a marked downward extension of stronger winds from aloft (Figure 55 and Figure 56) was a feature that was likely to have occurred during that afternoon and was responsible for the marked increased in observed surface winds speeds. Moreover, the neutral stratification, coupled with a fiercely burning fire at the surface, would have been likely to promote the formation of pyrocumulonimbus clouds reaching the tropopause about 12 km above mean sea level (Figure 58). At these heights, winds were reaching 60 to 70 knots (31 m s^{-1} to 36 m s^{-1}), and firebrands could be lofted and then transported downstream to great distances. We suggest that these are likely to be responsible for the rapid spread of the Blue Mountains Fires during the afternoon of 17 October.

Dynamically forced downslope winds in the lee of mountain ranges have been observed and well documented at many locations worldwide and have received much research effort. The escalations in fire activities at Margaret River (Kepert and Fawcett 2013) and Aberfeldy (Wells et al. 2014), which were also likely due to enhanced downslope winds, also occurred during the nighttime. However, this Blue Mountains case is of particular interest because the

phenomenon occurred during the day with a deep unstable layer that extended to about 3 km above mean sea level, which is apparent in both the tephigram at Mount Boyce (Figure 58) and the vertical cross section of the potential temperature (Figure 57).

Dry slot

The simulation also revealed the presence of a slot of dry air moving across the Blue Mountains area during the afternoon of 17 October, quite possibly related to the abrupt drying around 0300 UTC as observed in the Mount Boyce AWS. Figure 59, which shows the simulated 2-metre specific humidity at 0145 UTC 17 October, clearly depicts the existence of the dry slot which moved across the Blue Mountains area. A time evolution of the surface specific humidity (Figure 67 and Figure 68) shows that the dry slot was actually not formed within a wetter region of air. Rather, as the moist air-mass from the southwest, associated with the wind change, encountered the moist northerlies over northern New South Wales around 2200 UTC 16 October, the dry area between the two air masses was squeezed into a narrow strip and was maintained until around 0700 UTC 17 October. It is recognised that these moist northerlies are responsible for the increased dew point temperatures observed at Mount Boyce during from 1800 UTC 16 October to 0000 UTC 17 October before the arrival of the dry slot (Figure 49).

Previous studies (Mills 2005, 2008) suggested that such dry slots could be related to middle atmospheric dryness being transported down to the surface ahead of the front by a vertical cross frontal circulation. This theory was supported by middle atmospheric dryness revealed from satellite water vapour imageries in Mills' cases. However, in our study, such middle atmospheric dryness was not seen in the satellite water vapour imagery (Figure 60), and, indeed, the simulated vertical profile do not agree with the notion of middle atmospheric dryness supporting the maintenance of the dry slot (Figure 58). In the vertical profile, there was no marked mid-atmospheric dry layer, but a somewhat moist layer around 3.5 km to 4.5 km, and another around 5 km to 7.5 km, in reasonably good correspondence with the satellite water vapour imagery.

Figure 61 shows a stack of aligned vertical cross sections of simulated specific humidity, vertical velocity and potential temperature, running southwest to northeast at 0145 UTC 17 October, when the dry slot was directly over the Blue Mountains area. The simulated dry slot was about 100 km wide and 3 km deep, with a somewhat moist layer around 3.5 km to 4.5 km, and another around 5 km to 7.5 km, corresponding well with the tephigram (Figure 58). There was upward motion in the north-eastern (right) side of the dry slot and downward motion in the southwestern (left) side of the dry slot, resembling a vertical circulation with the strongest vertical motion between 4 km and 5 km above mean sea level. However it did not appear to be directly related to a frontal invasion as per Mills' cases, but rather, the undulations of the topography. To illustrate this point, Figure 62 shows a similar stack of vertical cross sections several hours earlier at 2030 UTC 17 October. The kink in the potential temperature between the markers 2 and 3 (about 180 km to the southwest of Mount Boyce which is at marker 6) was about 100 km ahead of the moisture discontinuity, and was associated with a skewed dipole of vertical velocity within the layer below 4.5 km from mean sea level

(highlighted by the green ellipse in the middle panel), with the downward motion confined below 3.5 km. Back to Figure 61, the moist layer around 3.5 km to 4.5 km was somewhat distorted by this vertical circulation, but it is not apparent that the dryness within the dry slot below 3 km is related to this vertical circulation. This suggests the notion of vertical cross-frontal circulation dragging down dry middle atmospheric air, as suggested in the papers by Mills (2005, 2008), is not likely to be responsible for the formation and maintenance of the dry slot in this case. Apparently, another mechanism is responsible for this.

To investigate this other mechanism, we note that the north-western portion of the dry slot is much better defined than over the Blue Mountains area (Figure 59), presumably because of the absence of significant terrain and other complicating factors. Therefore, we focus away from the Blue Mountains area, in this north-western portion of the dry slot instead.

Figure 63 and Figure 64 show a sequence of the simulated 10-metre wind direction and 2-metre specific humidity around the north-western portion of the wind change from 1700 UTC 16 October to 0000 UTC 17 October, with a marker X following the leading edge of the wind change. It is clear from the figures that the moisture discontinuity lagged behind the wind change, which resembles an undular bore. The dry slot formed as the dry area was wedged between the two moist air-masses from the southwest and the northeast and was maintained following the convergence of the two air masses. Figure 65 and Figure 66 show a sequence of vertical cross sections of simulated potential temperature and specific humidity following marker X in Figure 63 and Figure 64, with the faint circles corresponding to the position of marker X. The speed of the marker X is shown on the top right-hand corner of the vertical cross section. Before encountering the north-eastern moist air-mass at around 2200 UTC 16 October, marker X travelled at a speed of around 15 m s^{-1} and decelerated to around 6.5 m s^{-1} after the undular bore encountered the north-eastern air-mass around 2230 UTC 16 October. The dry slot maintained its width thereafter.

It is apparent that the dry air between the undular bore and the moisture discontinuity appears like a "dry nose", attached to the head of the southwestern moist air-mass. Upon convergence with the north-eastern moist air-mass, the "dry nose" was trapped within the two air masses, becoming the dry slot, and continued moving north-eastwards for the next several hours. The simulated dry slot at 0000 UTC 17 October was about 50 km in width with asymmetric depths of moist air on both sides, about 1.3 km in the southwestern side and 3 km in the north-eastern side.

Trajectory analysis was performed to illustrate that the slot of dry air was simply advected by the winds during the afternoon of 17 October. A set of 16 parcels (marked as letters A to P) were placed within the dry slot 400m above ground level at 0100 UTC 17 October, and the trajectories of these air parcels were then calculated by integrating both forward and backward for 7 hours.

Figure 67 and Figure 68 show the trajectories of these parcels from 1800 UTC 16 October to 0800 UTC 17 October. It is clear from the figure that the air parcels lined up at around 2300 UTC 16 October, and approximately defined the dry slot from about 0000 UTC 17 October to about 0400 UTC 17 October, before encountering moist convective outflow from the north. Parcels B, C, and D were

lifted by the moist outflow from the north at around 0400 UTC 17 October, with parcels A, E, F, G, and H following in the subsequent hours. By 0800 UTC 17 October, all the parcels were lifted from near the surface and the dry slot was dissipated.

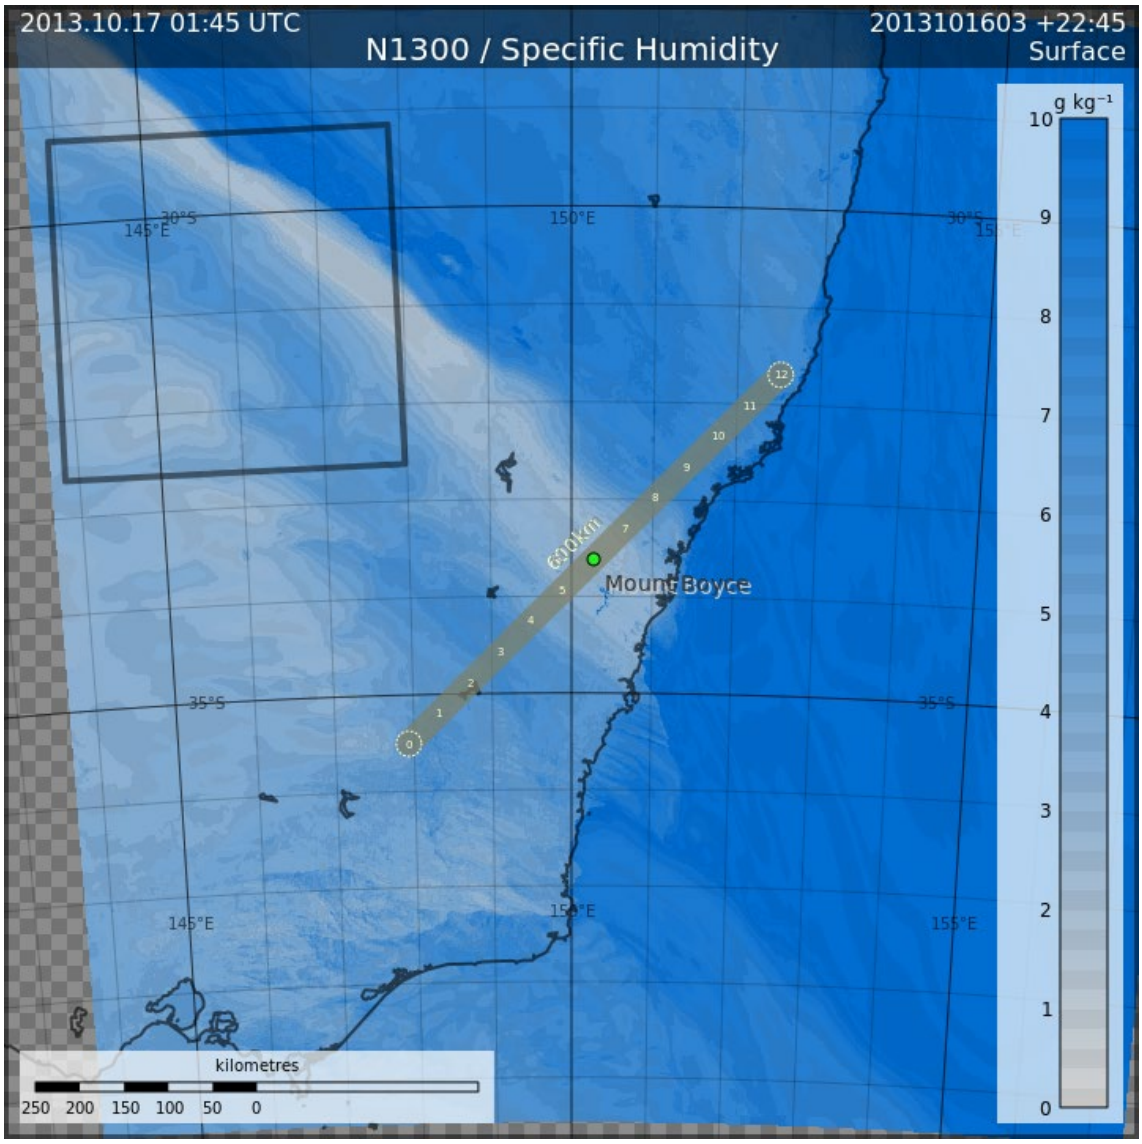

Figure 59. NWP model simulated 2-metre specific humidity at 0145 UTC 17 October 2013 from N1300 nest. The box in the northwest depicts the domain of the charts in Figure 63 and Figure 64. The line running southwest to northeast depicts the extent of the vertical cross section in Figure 61.

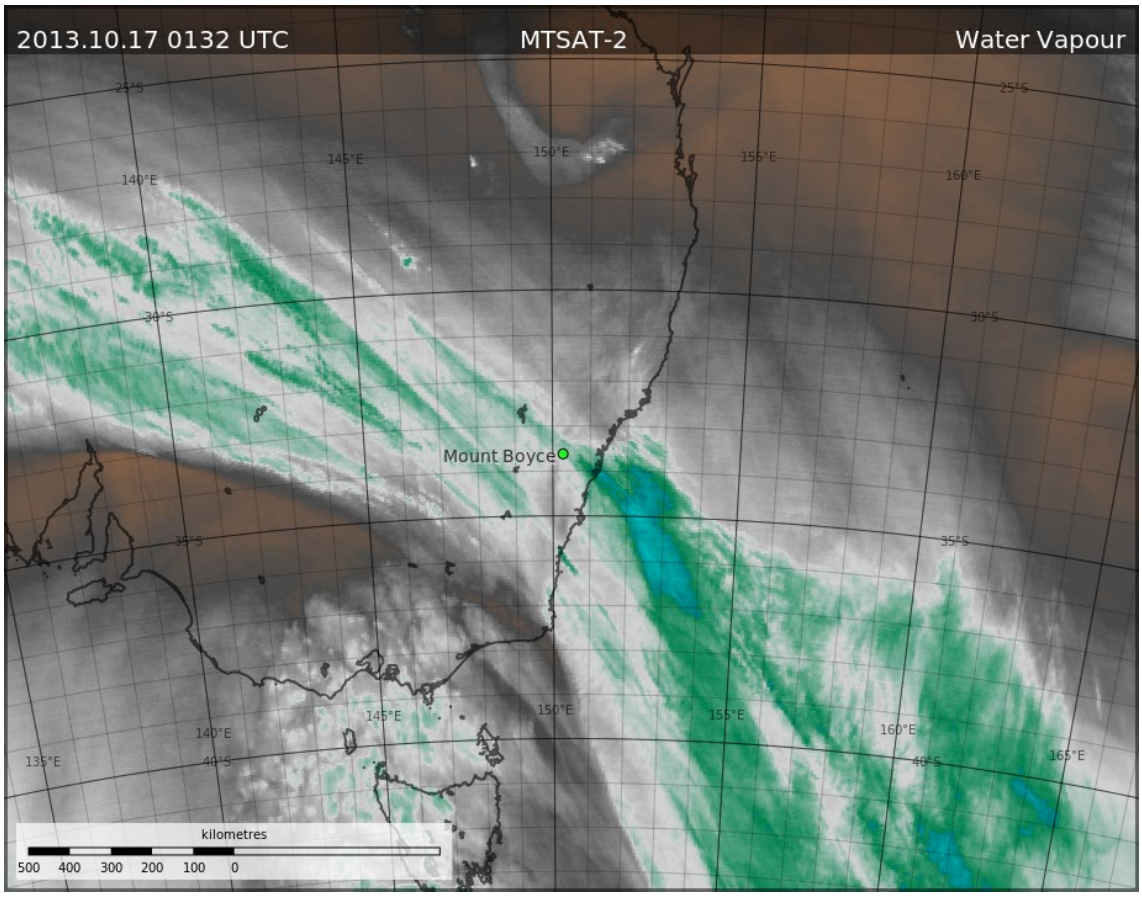

Figure 60. Water vapour imagery at 0132 UTC 17 October 2013 from the MTSAT-2 meteorological satellite. Colour enhancement shows very dry air as brown and more moist air progressively as grey, white, green and cyan.

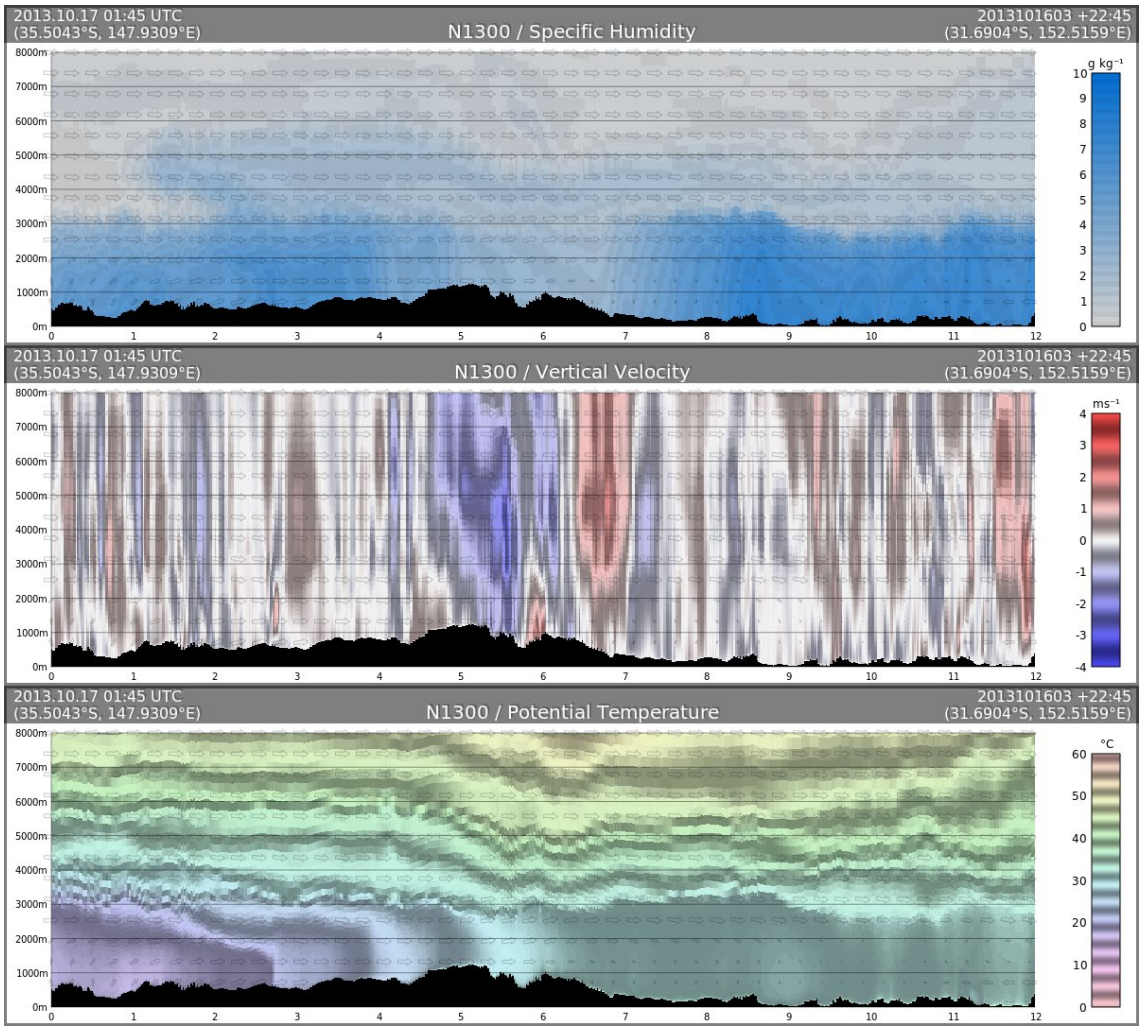

Figure 61. NWP model simulated vertical cross section of specific humidity, vertical velocity and potential temperature at 0145 UTC 17 October 2013 from the N1300 nest. The extent of this vertical cross section is 600 km long and is depicted in Figure 59. Numbers 0 to 12 at the bottom are equidistant position markers which are 50 km apart. Faint arrows indicate the three-dimensional wind direction projected onto the vertical cross-sectional plane.

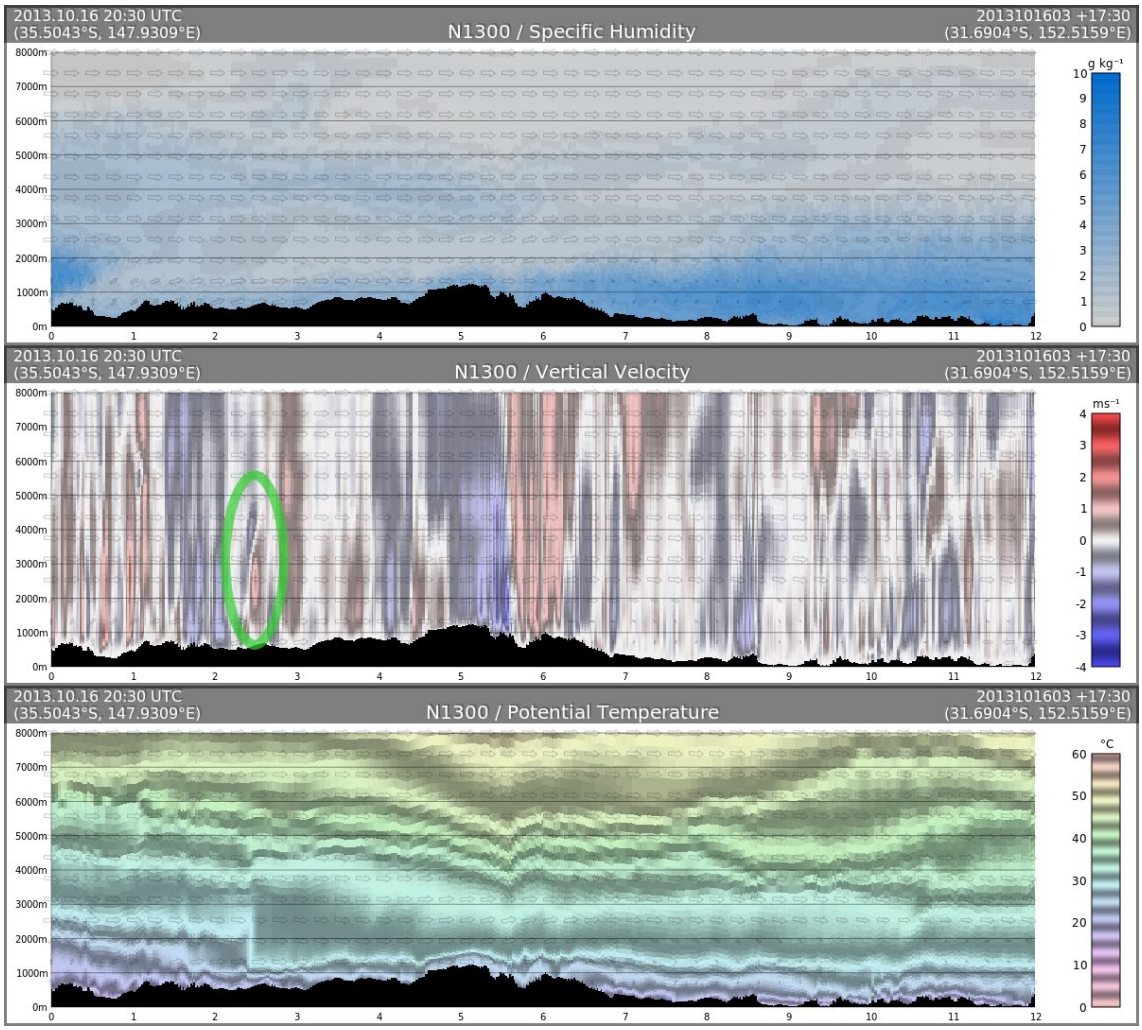

Figure 62. As in Figure 61 but for 2030 UTC 16 October 2013. The green ellipse highlights the vertical velocity dipole associated with the potential temperature discontinuity

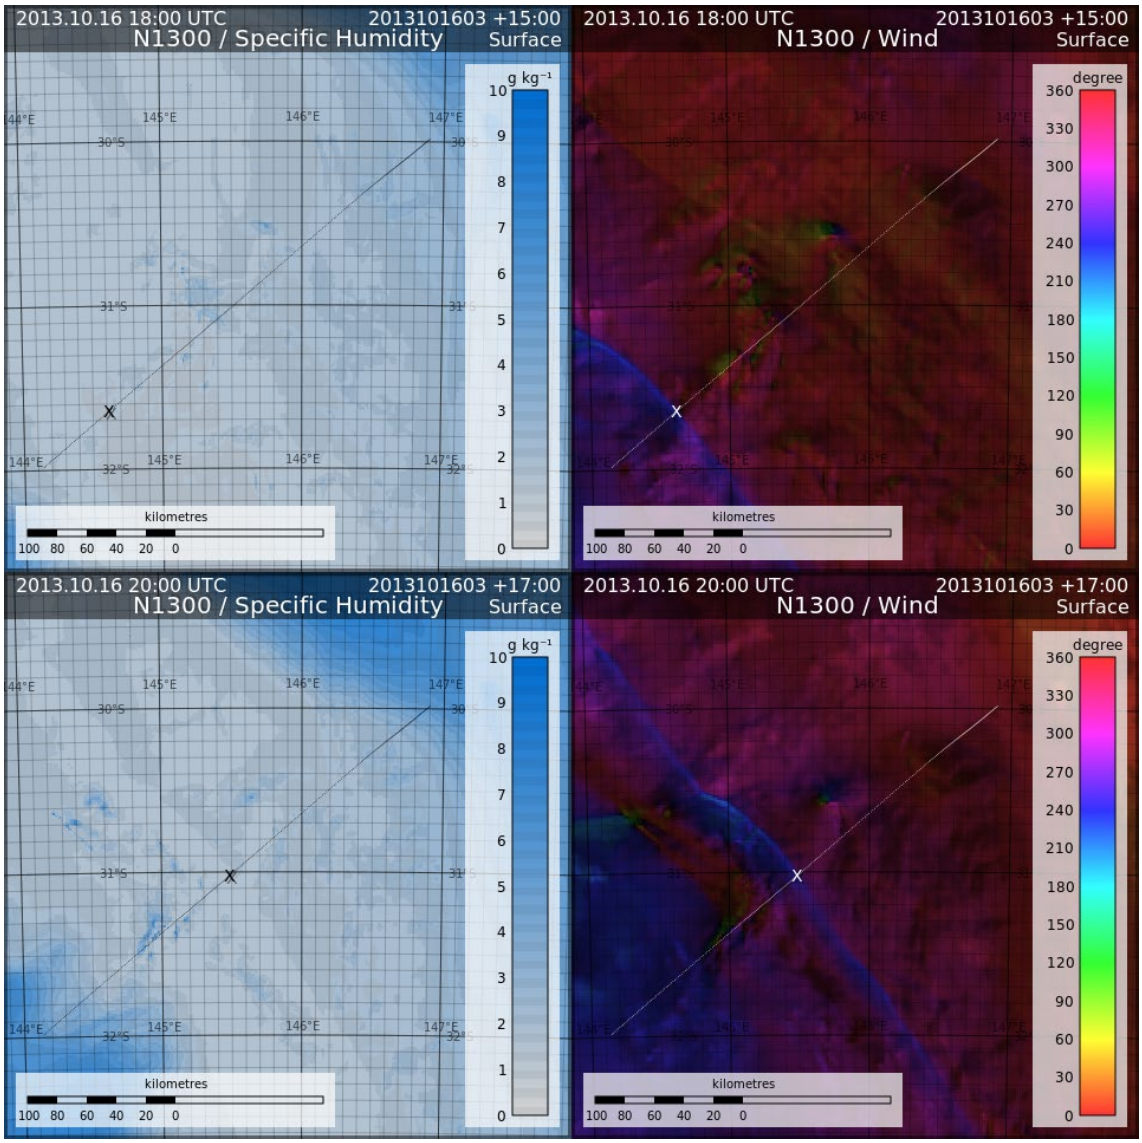

Figure 63. Montage of NWP model simulated 2-metre specific humidity and 10-m wind from 1800 UTC 16 October 2013 to 2000 UTC 16 October 2013. The wind direction is hue coded as per the legend, with brightness varying linearly with wind speed. The dotted line is the track of a marker (X) that moves northeastwards with the wind direction discontinuity. The domain of these charts is depicted by the box over the northwestern region of Figure 59.

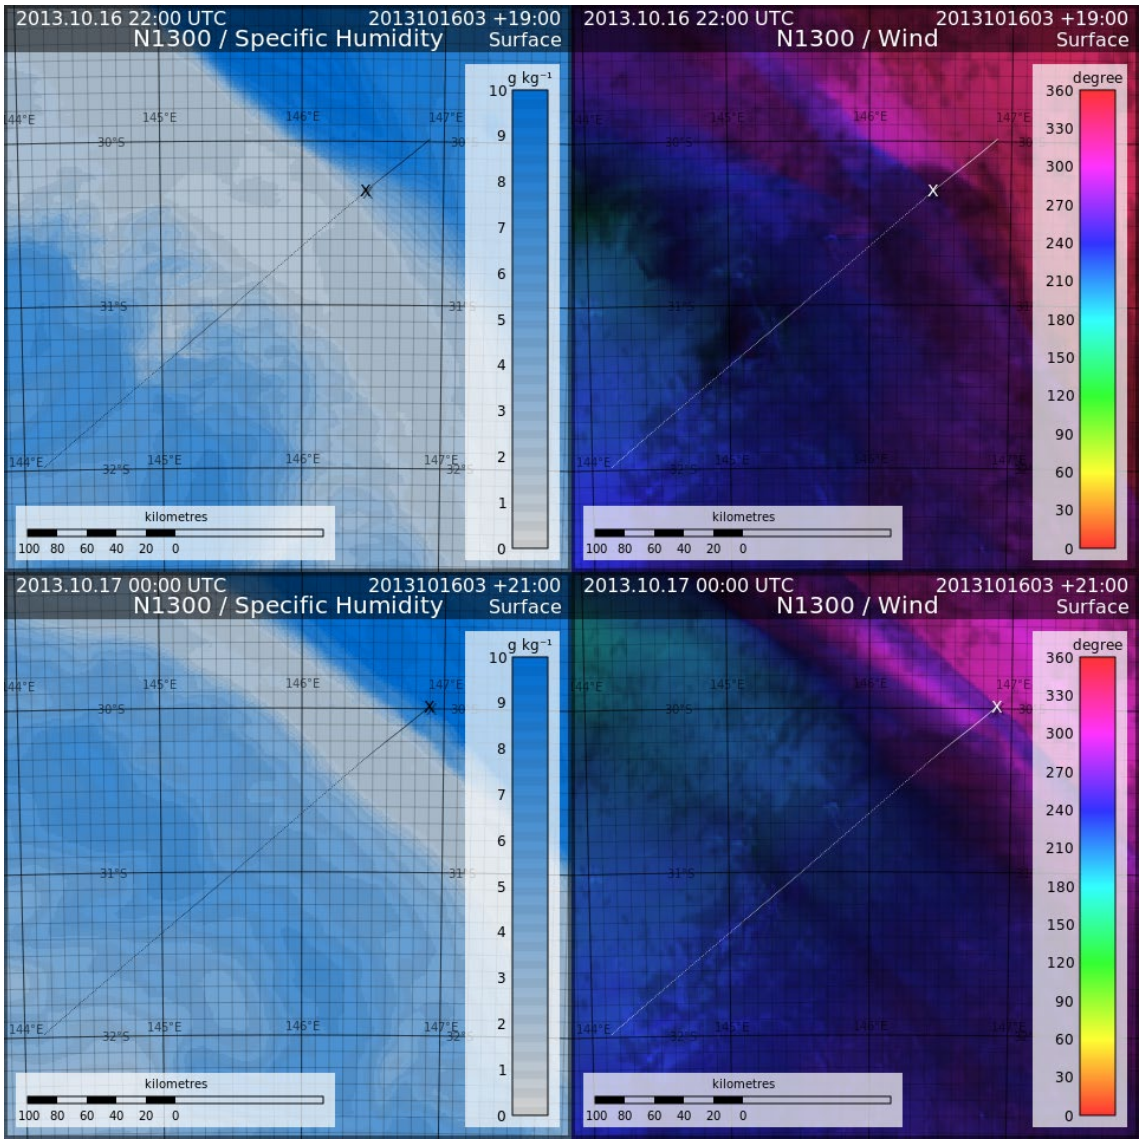

Figure 64. As in Figure 63 but for 2200 UTC 16 October 2013 and 0000 UTC 17 October 2013.

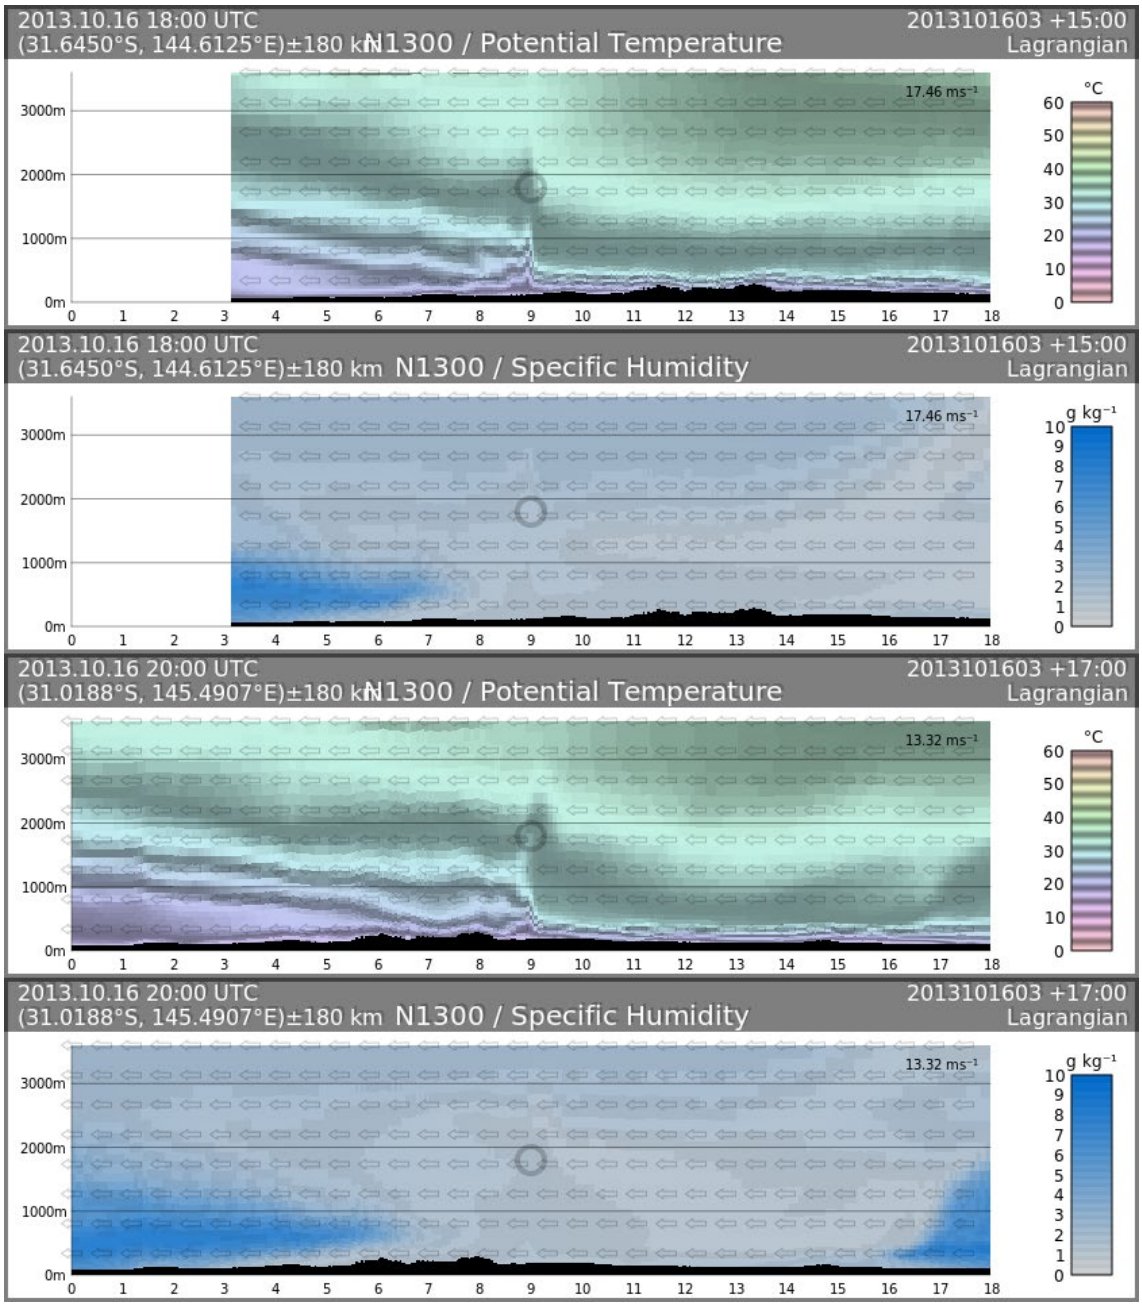

Figure 65. Montage of vertical cross sections of simulated potential temperature and specific humidity at 1800 UTC and 2000 UTC 16 October 2013 from the N1300 nest, centred at the marker X in Figure 63 aligning with the track and spanning 360 km in length. The numbers 0 to 18 in the horizontal axis are equidistant position markers 20 km apart. The faint circle corresponds to the location of X, moving at a speed indicated by the figure marked in the top right-hand corner of the chart. Faint arrows indicate the 3-dimensional wind direction, relative to the motion of X, projected onto the vertical cross-sectional plane.

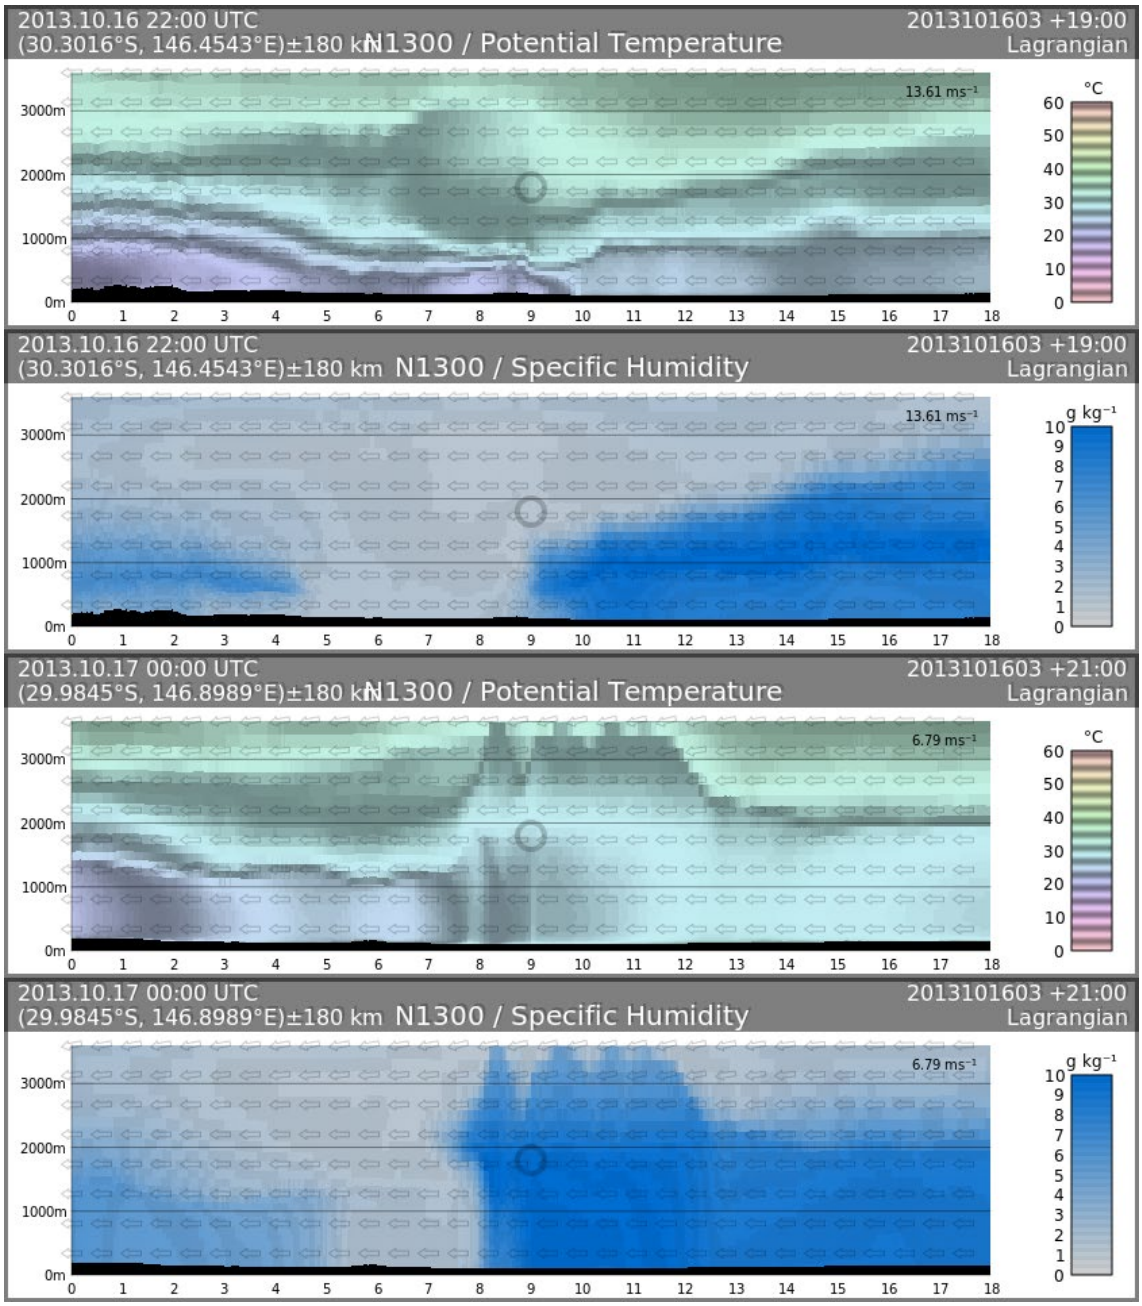

Figure 66. As in Figure 65 but for 2200 UTC 16 October 2013 to 0000 UTC 17 October 2013.

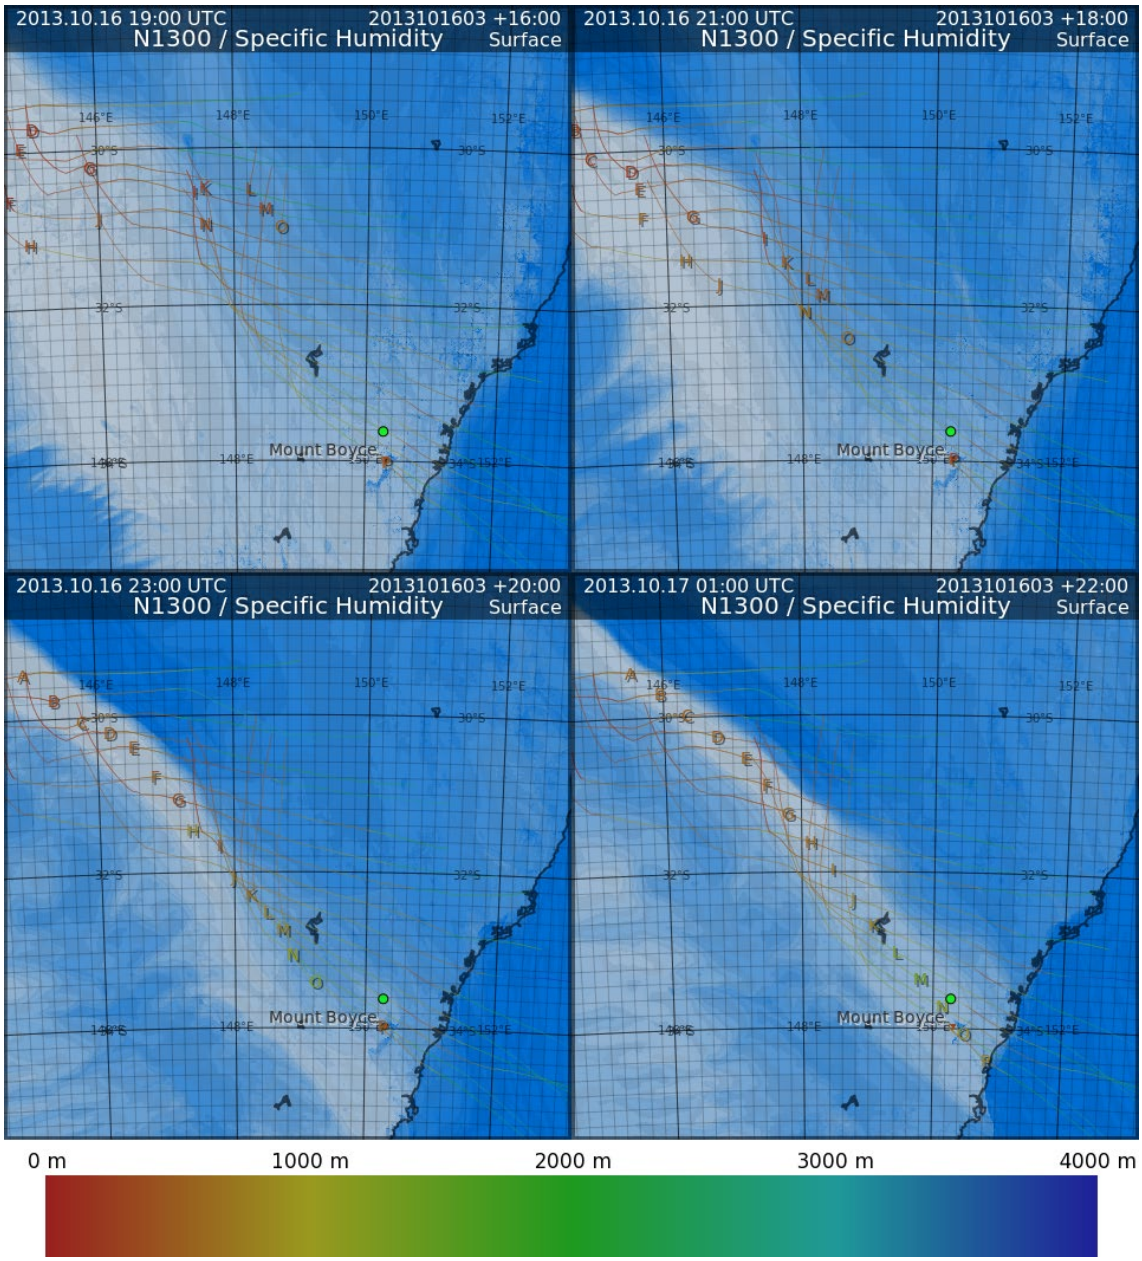

Figure 67. Montage of NWP model simulated 2-metre specific humidity every 2 hours from 1900 UTC 16 October 2013 to 0100 UTC 17 October 2013 from the N1300 nest. The 16 letters A to P depict parcels following trajectories integrated backward for 7 hours initialised at 0100 UTC 17 October 2013, 400m above ground level. Colour coding of specific humidity as in Figure 59. The letters are colour coded with the elevation of the parcel from mean sea level as in the colour legend at the bottom.

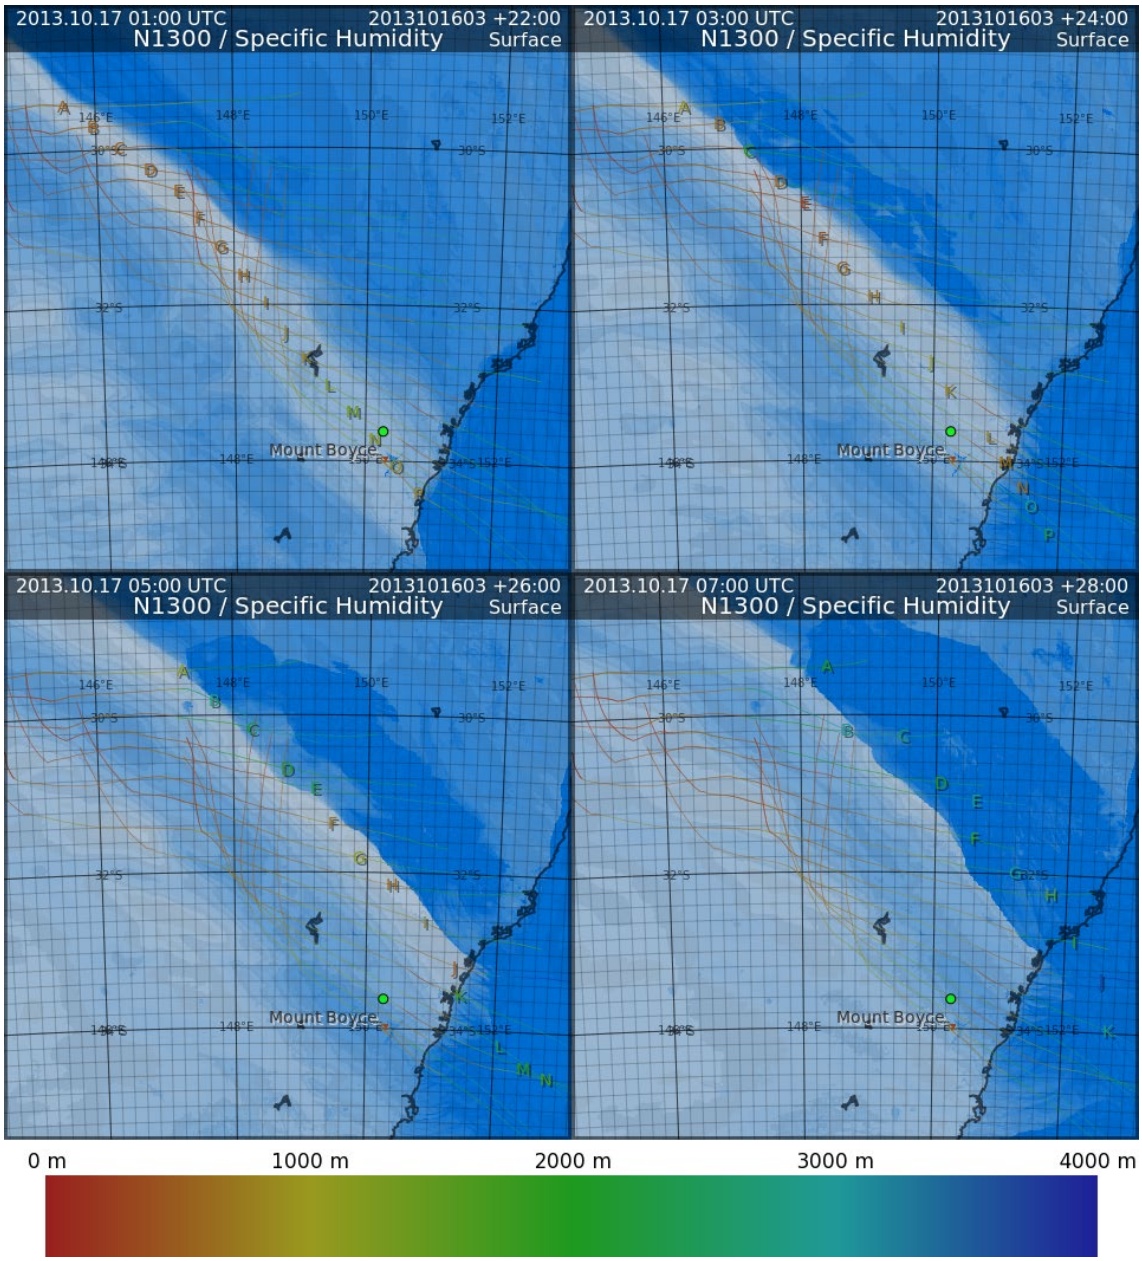

Figure 68. As in Figure 67 but from 0100 UTC to 0700 UTC 17 October 2013 with the alphabets following trajectories integrated forward for 7 hours initialised at 0100 UTC 17 October 2013.

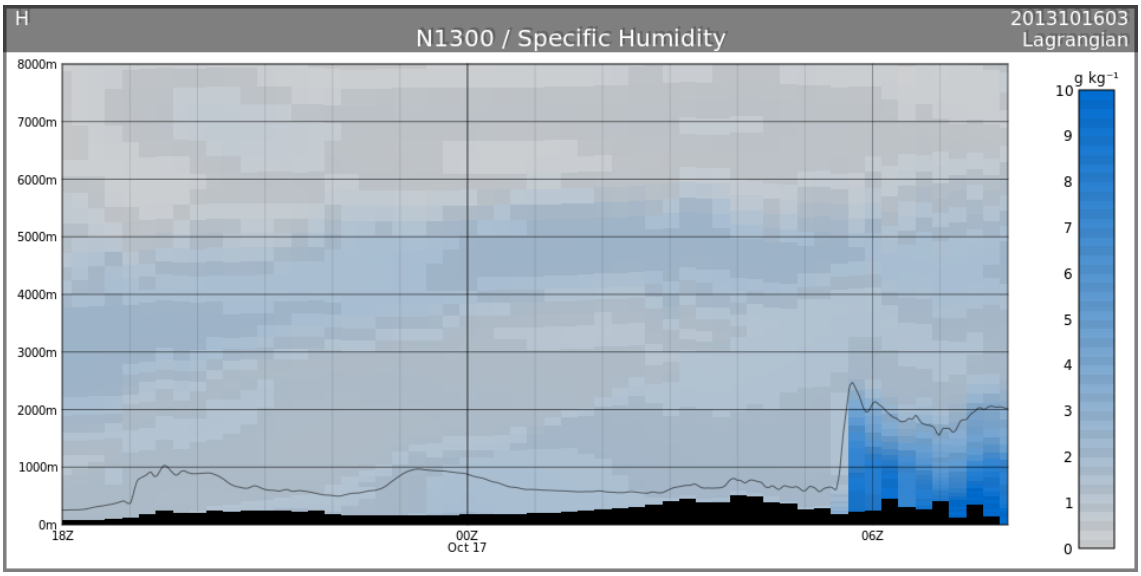

Figure 69. Time cross section of NWP model simulated specific humidity following parcel H in Figure 67 and Figure 68 from 1800 UTC 16 October 2013 to 0800 UTC 17 October 2013 from the N1300 nest. The thin black line depicts the height of parcel H.

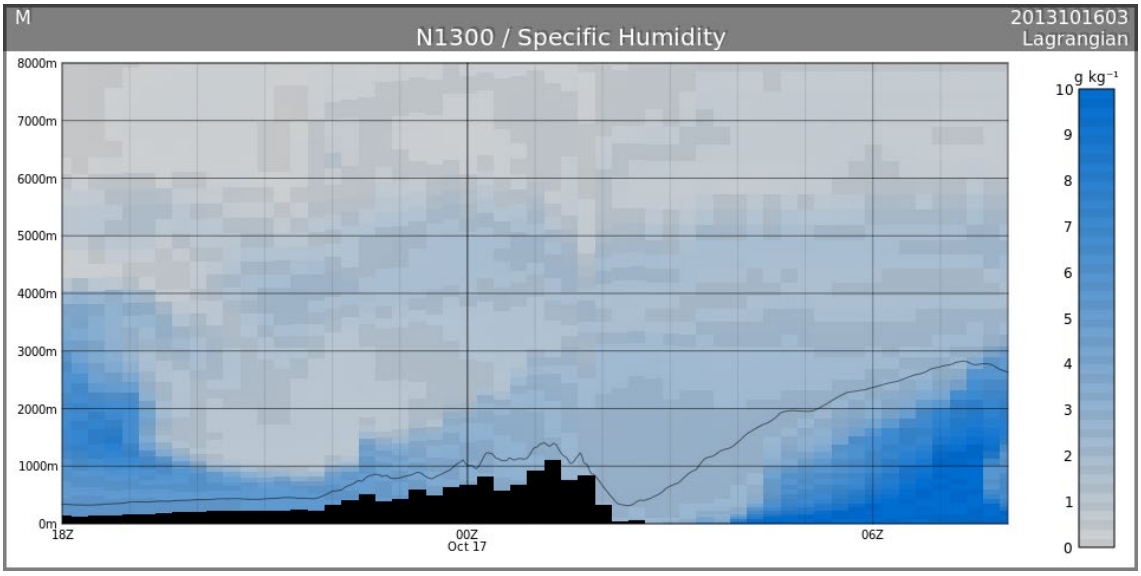

Figure 70. As in Figure 69 but for parcel M in Figure 67 and Figure 68.

Discussion and conclusions

Analysis of the observations at Mount Boyce, a nearby AWS to the Blue Mountains Fires, revealed that the hot, dry and windy conditions are responsible for the severe fire weather conditions on 17 October 2013, leading to the rapid spread of the fires. With substantially below average rainfall in the several months leading up to October 2013, the dryness about the Blue Mountains area was approximately in the top 5 percentile, which laid down the foundation of a potentially significant fire hazard. On 17 October 2013, the daily maximum temperature of Mount Boyce was 22.5°C. Considering the elevation from mean sea level of the AWS site of 1080m, it was actually a rather hot day (82.2 percentile of all October days).

The strong to gale force winds on the day, highlighted by the observation of 28.3 m s^{-1} gust at Mount Boyce, exacerbated by the passage of a slot of very dry air during the afternoon of 17 October 2013, reaching a minimum of -7.6°C, turned out to have elevated the fire risk of the day from high to severe. The NWP simulation suggests that the winds were strengthened by the mountain waves excited by the topography in the form of enhanced downslope winds, and the dry slot moved across the Blue Mountains area coincidentally near the onset time of the increase of wind speed as observed at Mount Boyce. The vertical profile of the atmosphere at Mount Boyce during the afternoon indicated a 3 km column of unstable, dry air with very strong winds (about 30 m s^{-1}), were conducive to the formation of pyrocumulonimbus clouds, and the lofting and transport of fire brands to great distances downstream, which all contributed to the very rapid eastward spread of the State Mine Fire on the afternoon of 17 October 2013.

The formation of the dry slot, according to the simulation, was related to the development of an undular bore which formed upstream the night before, when the cool change slid into the shallow layer of cool air formed by nocturnal cooling. A "dry nose" was formed on the north-eastern periphery of the advancing moist, cool change. Upon collision with the moist air-mass from the north, the "dry nose" became a dry slot of air during the morning of 17 October. The dry slot was then maintained by simple advection and moved across the Blue Mountains area during the afternoon, before dissipating during the early evening when the dry air was lifted from the surface as it encountered the moist outflow of the convective activities further north.

The sparseness of the observational network makes it intrinsically hard to capture mesoscale features, like the ones presented in this paper. Fortunately, good quality, high resolution NWP model simulation data is available, which makes such analyses possible, although certain discrepancies between the observations and the simulation are noted.

Practical implications

This fire event was very high impact, and the mid-October timing was well away from the usual seasonal peak in fire activity. We have identified the meteorological factors responsible, being the mountain waves increasing the surface wind speed, and the dry slot. It appears, based on this research and some of our earlier work, that mountain waves are an underappreciated factor in unexpected fire behaviour. We conducted a major utilisation activity, including a *Hazard Note* and media articles, to raise awareness of this issue.

THE METEOROLOGY OF THE TATHRA BUSHFIRE

Introduction

On the 18th March 2018 a fire started at Reedy Swamp within the Bega Valley Shire on the New South Wales South Coast. Aided by an approaching cold front, the fire burned into the town of Tathra during the mid-afternoon, leading to the evacuation of the township and the destruction of 70 homes and other structures. Weather on the day was characterised by a hot, dry and gusty pre-frontal north-westerly airmass ahead of a cool, gusty coastal change in the late afternoon. Such conditions are notoriously difficult to forecast on the south coast due in-part to the influence of the prominent topography upwind which comprises the highest mountain peaks on the Australian mainland.

The high impact nature of the Tathra bushfire event, coupled with the relatively common incidence of such wind changes and the difficulty they present to forecasters, motivated this case-study in consultation with the project end users.

Methods

To examine the meteorology of the day, two complementary approaches were considered. Firstly, high-frequency observational data analysed from nearby weather stations was compared with available satellite imagery and cross-referenced with significant events in the evolution of the fire. This process distinguishes the key periods for analysis and defines a basic truth which any conceptual understanding of the event must be able to explain. One-minute wind gusts and temperature observations from Bega and Merimbula, the two stations nearest the fireground, are shown in Figure 71.

Note three critical periods:

1) An abrupt increase in temperature and gustiness as winds turned north-westerly at both stations just prior to the first report of the fire at around 12:30 AEDT. The approximate origin of the fire is marked (***O***) at roughly 60-70m AMSL in Figure 72.

2) A re-intensification of conditions at Bega after 14:00 AEDT, with gusts up to 76 km/h and temperatures of around 38°C. Though winds remained highly variable relative humidity dropped as low as 15-20% (not shown) and fire activity peaked. The primary mode of propagation during this period was via mass spotting of embers. The topography inland of Tathra near Vimy Ridge, marked (***V***) in Figure 72 at 150-200m AMSL, provided a key launching point for the fire to cross the Bega River by 14:56 AEDT, and subsequently impact Tathra by 15:34 AEDT. It was also in this location, between Vimy Ridge and the river, that enhanced lee slope fire behaviour was observed.

3) The arrival of the cool change at Merimbula around 16:30 AEDT which reduced temperature by over 10°C, only for conditions to deteriorate once more roughly 20 minutes later. Curiously, the change would not completely establish for a full hour after first incidence. Witness accounts from personnel on the fireground suggest similar behaviour was observed at Tathra. The change would not move through the more inland station at Bega until just after 19:00 AEDT.

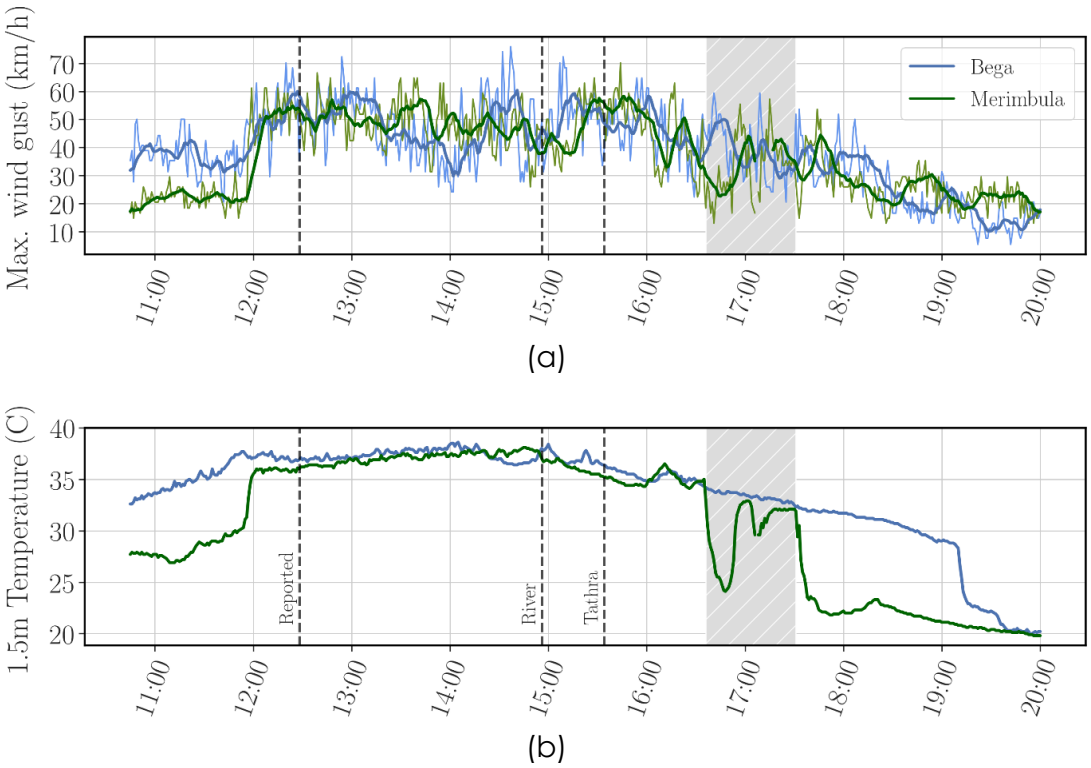

Figure 71: One-minute observations at Bega (blue) and Merimbula (green) automatic weather stations (time in AEDT on March 18th 2018): a) maximum 10m-wind gust (light) with 10-minute rolling average (dark), and; b) 1.5m temperature. The black dashed vertical lines indicate critical periods of the fire evolution: 12:28 AEDT - fire reported; 14:56 AEDT – the fire crosses the Bega River; and 15:34 AEDT - the fire impacts Tathra. The grey shaded region indicates the period of variable conditions from first incidence of the cool change at Merimbula (16:36 AEDT) to the establishment of the south-easterly (17:31 AEDT).

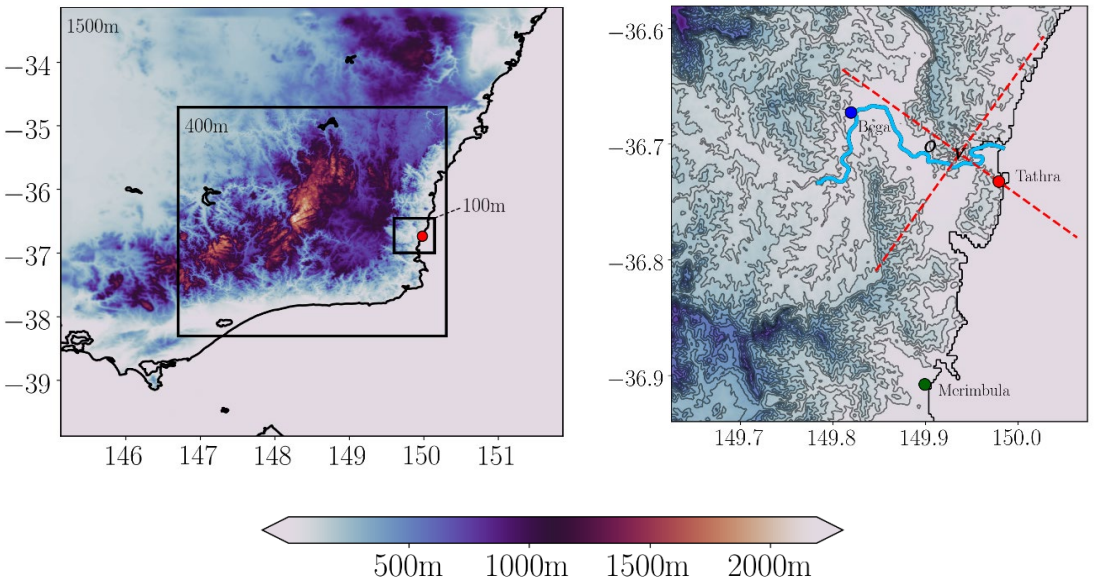

Figure 72: Left: the three highest resolution nested domains centred over the southeast of Australia (1500m, 400m and 100m horizontal resolution) with orography shown. The location of the NSW South Coast town of Tathra is marked in red. Right: a close-up of the orography near Tathra showing the location of the approximate origin of the fire (**O**), Vimy Ridge (**V**), the Bega River (shown in light blue) and the stations at Bega and Merimbula. The red dashed lines show the cross-sections used in Figure 74.

To investigate the above observations, we use a nested configuration of the unified model (UM), the model underpinning the operational ACCESS model used by forecasters at the Bureau of Meteorology. While the highest resolution operational models in Australia currently run at 1.5km horizontal resolution, we were able to run our research simulations to very high resolutions of 400m and 100m. The three highest resolution domains, along with a close-up of the region near the fireground with orography, are shown in Figure 72.

Results

At high resolution, fine-scale details emerge within the simulations which have significant influence on the meteorology of the day. Firstly, the abrupt increase in winds and temperature around the time the fire started coincides with a strengthening of the 400m simulation lee wave pattern, shown as alternating bands of ascending and descending air in Figure 73(a). These trapped lee waves, a form of mountain wave induced by flow over the Great Dividing Range, intensified around this time in response to the onset of diurnal convection, with a strong region of descending, high speed air over the approximate origin of the fire at around the time the fire started. The lee wave pattern, though weakening, persisted throughout the remainder of the afternoon. It is possible that mountain waves may have modulated the developing convection and influenced the severity of surface and near-surface conditions around the time of peak fire activity. The visible image in Figure 73 (b) shows evidence of the lee wave pattern at 13:30 AEDT.

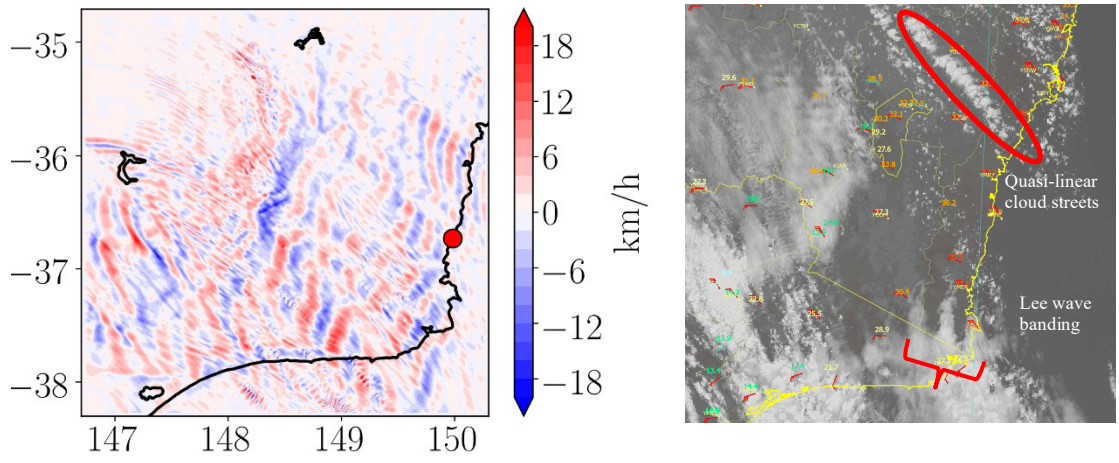

(a) (b)

Figure 73 (a) The 400m simulation vertical velocity at 3.5km above sea level at 11:45 AEDT on March 18th, 2018, shortly before the fire was first reported in Reedy Swamp. Tathra is marked. Note the banded 'lee waves' of ascending (red) and descending (blue) air; (b) a visible satellite image at 13:30 AEDT just prior to the re-intensification of conditions at Bega. The lee waves are visible as cloud bands near the Victorian border, while 'cloud streets' forming to the north are likely to be associated with developing boundary layer rolls.

Secondly, the re-intensification of conditions from around 14:00 AEDT was likely in response to the passage of significant horizontal convective rolls, which the 400m and 100m simulations suggest passed over the fireground around this time. These structures, also termed boundary layer rolls, form in sheared convective boundary layers and consist of counter-rotating gyres with alternating regions of ascending and descending air (Brown, 1980). They are sometimes visible as

"cloud-streets" in satellite imagery, see Figure 73(b), and have previously played a significant role in other extreme fire events, notably the Black Saturday fires (Engel et al., 2013). The rolls act to concentrate the maximum wind energy into streamer-like bands aligned roughly parallel to the mean flow. Their passage generates variable conditions, with dry, high speed air mixed to the surface as the descending flank passes, and typically warmer, less windy conditions with the ascending flank. The observations at Bega during this period are consistent with this picture.

These structures would have had a marked effect on conditions on the fireground. Figure 74 shows vertical cross-sections of the 100m simulation windspeed and vertical motion at approximately 14:10 AEDT as the rolls first moved over the region. The cross-sections show the descending flank of a roll impacting much of the fireground with winds in excess of 90 km/h near ridgetop, and descending air (shown in blue) on the lee slope of Vimy Ridge. However, the conditions associated with the rolls are erratic with considerable variation apparent both spatially and temporally. Intense updrafts (shown in red) are simulated to occur near and on the fireground throughout the period of peak fire activity, with local conditions heavily modulated by the passing rolls. Despite this, there is a clear and prolonged signal of descending air on the lee slope of Vimy Ridge, which would provide favourable conditions for the fire to burn downhill toward the river. This likely explains the enhanced lee slope fire behaviour observed at this location. Furthermore, the updrafts associated with the rolls likely provided a key lofting mechanism, transporting embers ahead of the progressing fire front and aiding in the fires cross of the river and attack on Tathra. Conditions on the fireground during this period were likely similar to observed at Bega, however, for short periods winds may have been significantly stronger, particularly about the exposed ridges where the model indicates that gusts in excess of 90 km/hr were probable. It was not long after the simulation indicated that the rolls first impinged on the fire that it crossed the Bega River (14:56 AEDT) and began to impact the township (15:34 AEDT).

Conditions remained severe until the cool change cleared the South Coast, however, its passage across the region was highly complex. During the period of fluctuation at Merimbula Airport, shown by the grey shading in Figure 71, the 100m simulation indicates that boundary layer rolls were still very much active and even shows some evidence of post-frontal temperature “spikes”; isolated regions behind the change with locally higher temperature and lower humidity. It seems likely, therefore, that the developing cool change, which was initially of a relatively shallow depth, was forced out by residual boundary layer rolls still present above. Based on the observations at Merimbula and the 100m simulation, the change likely first impacted the fireground around 16:30 AEDT. The period following, however, is uncertain. For possibly up to two hours, but more likely around one hour, the fireground was under the influence of two battling airmasses. The first, a building cool change from the southeast, the second characterised by significant vertical convection bringing hot, dry and gusty conditions back to the surface. The orography through the fireground may also have acted to locally impede the south-easterly and encouraged the continuation of severe fire weather for a significant period after the change was first incident.

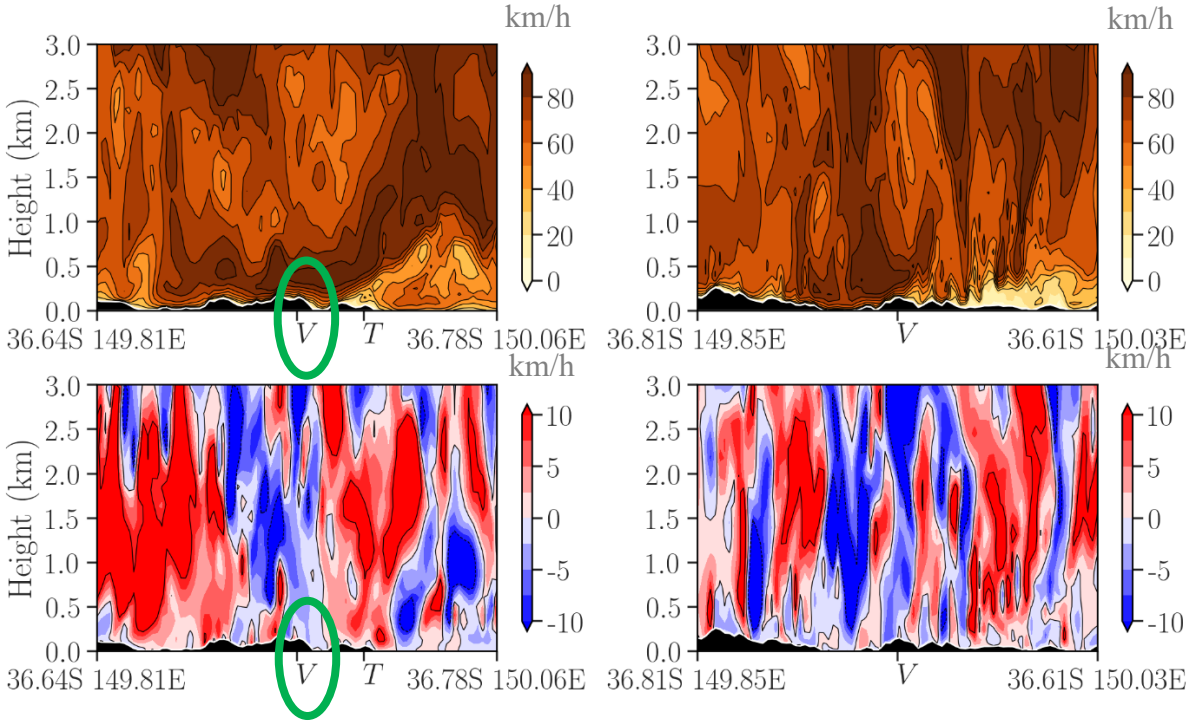

Figure 74: Two vertical cross-sections through Vimy ridge (V), shown in Figure 72, oriented perpendicular to each other at roughly 14:10 AEDT as the boundary layer rolls first pass over the fireground. Left: a cross-section through the path of the fire toward Tathra (T) with the lee slope of Vimy Ridge circled; right: the perpendicular cross-section. Top: horizontal windspeed (km/h) and, bottom; vertical velocity (km/h). Regions of descending air are marked in blue, while red areas show ascending air.

Conclusions

The extreme conditions observed during the Tathra bushfire were a result of interactions between organised convection, internal gravity waves induced by flow over complex terrain, and the passage of a cold frontal airmass boundary. In particular,

- Intensifying mountain waves likely contributed to the strong wind event around the time the fire started and may have influenced the severity of conditions during the period of peak fire activity in the afternoon.
- Boundary layer rolls were the dominant mode of convection and were responsible for strong gradients in wind speed, direction and vertical velocity, generating highly variable conditions across the fireground. It seems probable that the enhanced lee slope fire behaviour observed near Vimy Ridge was a result of these structures passing over the region. Moreover, currents of strongly ascending air likely contributed to the transport of embers which aided in the fires cross of the Bega River, a crucial barrier protecting the town.
- Conditions during the late afternoon were erratic, likely a result of interactions between boundary layer rolls and the passage of the cool change. The change was initially incident in the region around 16:30 AEDT but may have taken a further two hours to fully push over the fireground, with conditions at times oscillating unpredictably between the two extremes; a difficult environment under which to control the spread of a fire.

Practical Implications

These conclusions draw heavily from very high-resolution simulations and illustrate their utility in understanding, and predicting, high impact weather. Though such high-resolution model output is not currently available to forecasters, the results of this study do have implications for operations and fire management.

Firstly, while forecast models do not typically resolve roll-type structures, the preconditions for their formation are understood and they can sometimes be observed; for example, as cloud-streets via satellite and by examining one-minute observations. Meteorologists may therefore be able to advise fire agencies of the associated conditions, for example the risk of enhanced lee slope fire behaviour in hilly terrain and of short-term fluctuations in wind speed and gustiness.

Secondly, the time of a wind change is often forecast, with tolerance, as a simple progression from one regime to another. In the case of the Tathra fire and at least two other South Coast fires that we are aware of, the progression has been highly complex with dangerous conditions redeveloping after an initial reprieve. Forecasting the occurrence of such behaviour at a point is currently very difficult. However, an improved awareness among forecasters and practitioners may allow quick diagnosis when it does occur.

Finally, this case-study highlights the significant spatiotemporal variation in conditions often present during high impact weather. Though the minimum horizontal resolution of data available to forecasters is constrained by the model configuration, the temporal resolution at which this data is output is more flexible. Currently, many fields are output at hourly, or three-hourly intervals. In the case of the Tathra event, much of the detail is lost at this resolution. In order to retain as much information as possible it may be necessary to make select fields, for example the 10m windspeed, available at higher frequency.

For cases in which high-frequency data output is not possible, developing surrogate or proxy fields which encapsulate the sub-sampled distribution, such as a measure of the hourly windspeed variability, may be an alternative. If the numerical model output holds information key to an emergency response, it is crucial that the data made available to forecasters and practitioners does not obscure it.

SECONDARY EYEWALL FORMATION IN TROPICAL CYCLONES

Summary

Roughly half of all intense tropical cyclones experience an eyewall replacement cycle. In these events, a new eyewall forms concentrically around the original one. This secondary eyewall develops its own wind maximum, and both the secondary eyewall and the wind maximum typically intensify and contract, whilst the original eyewall and wind maximum weaken and eventually dissipate. While the evolution of a storm with concentric eyewalls is reasonably well understood, the mechanism or mechanisms by which the outer eyewall forms remain elusive. Understanding secondary eyewall formation is an important problem, for the subsequent eyewall replacement cycle can significantly affect the intensity of the storm, and the formation process and replacement cycle are usually associated with a major expansion of the outer wind field. Both these factors significantly affect the cyclone's impact.

We investigate a high resolution simulation of an eyewall replacement cycle. Boundary layer convergence due to friction substantially influences the evolution of the convection, and we present evidence for a positive feedback involving convection, vorticity and frictional convergence that governs the subsequent evolution of the system. In this feedback, frictional convergence strengthens the convection, stretching of vortex tubes in the buoyant updrafts increases the vorticity, and the vorticity structure of the storm determines the strength and location of the frictional updraft.

Changes in the structure and intensity of tropical cyclones cause difficulties for their management, especially if these changes occur in the last day or two before landfall. Our improved knowledge of these processes will lead to better forecasts and mitigation.

Introduction

In literature, "the eye of the storm" is an apt description for a region of peace and calm, around which tumult and turmoil reign (White, 1973). The analogy is to a tropical cyclone, where the eye, characterised by light winds, little cloud and little or no precipitation, is surrounded by the towering clouds of the eyewall. These eyewall clouds slope outwards from the centre of the storm with height, giving the appearance of a giant amphitheatre. The strongest winds are close to the surface immediately below these clouds, and they generate substantial amounts of rain. Often, the formation of a symmetric eye signifies that the cyclone is entering an intense and destructive phase. Truly, the literary analogy is apt.

About half of all intense tropical cyclones experience an eyewall replacement cycle (ERC). In an ERC, a secondary eyewall forms, concentric about the existing eyewall. Over the next day or so, the new eyewall intensifies and contracts, while the initial eyewall weakens and eventually disappears. During eyewall replacement, the cyclone's intensity is steady or weakens, but intensification typically resumes once the replacement of the initial eyewall is complete. Notable examples of Australian tropical cyclones with an ERC include cyclones Vance (1999), Larry (2006) and Yasi (2011).

Figure 75 illustrates a typical sequence, using satellite images from a variety of passive microwave sensors in the 85 – 91 GHz band on polar-orbiting satellites. These frequencies can see through the dense cirrus overcast that normally obscures our view of the inner workings of tropical cyclones, and hence reveal the rainband and eyewall structure. At the start of the sequence, the primary eyewall is the red circle near the centre, surrounded by the deep convection of the spiral rainbands. Over the next 12 hours, the rainbands become more symmetric and organise into an outer eyewall by 22:43 on 22 August (panel d). The outer eyewall continues to become more symmetric, while the inner one weakens and decays, with the last vestige visible at 12:47 on 28 August (panel i).

The fundamental dynamics of the tropical cyclone eyewall replacement cycle (ERC), after the outer eyewall has formed, have been understood for over three decades (Shapiro and Willoughby 1982; Willoughby et al. 1982). In contrast, the cause of the initial formation of the outer eyewall has proved to be more elusive. Numerous theories have been proposed (see the reviews by Rozoff et al. 2012; Wu et al. 2012) but a consensus has not been achieved. More recently, attention has focussed on the possible role of the boundary layer in secondary eyewall formation (SEF) (Huang et al. 2012; Kepert 2013; Abarca and Montgomery 2013). Therefore, it is of interest to diagnose the boundary-layer processes occurring during SEF and the subsequent evolution of the eyewalls. This report describes our analysis of a SEF/ERC simulated by a high-resolution WRF simulation of a hurricane. An earlier report (Kepert and Nolan 2014) focussed on the boundary layer dynamics. Here we consider the effects of heating and the cloud processes also.

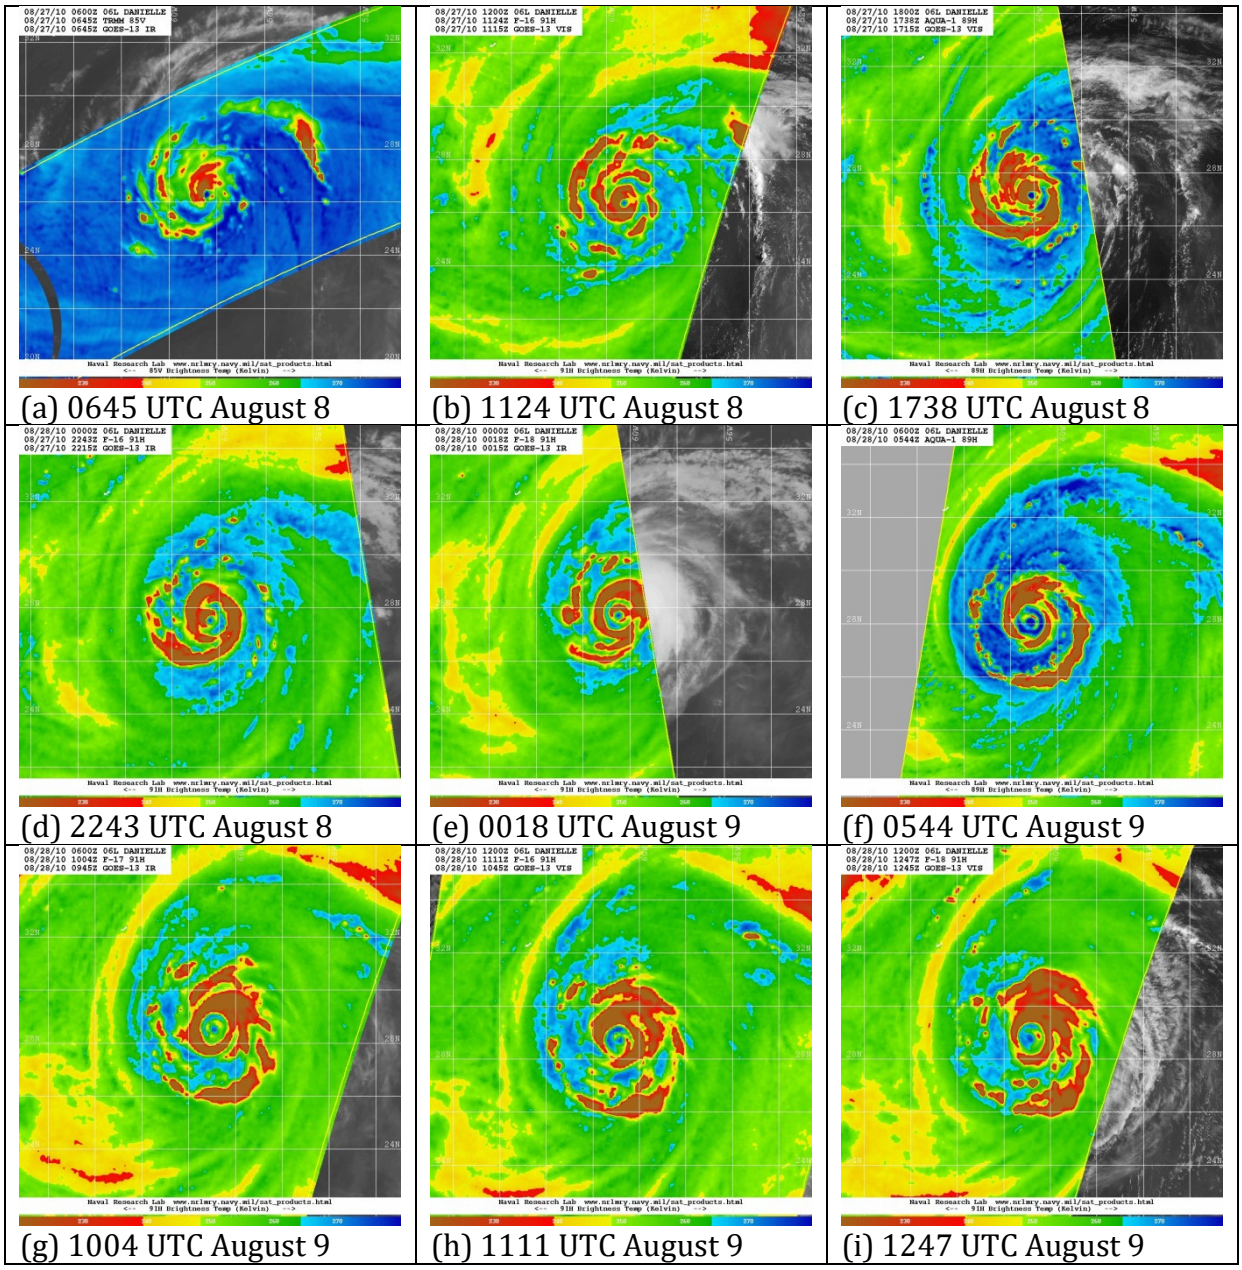

Figure 75: Passive microwave imagery in the 85 – 91 GHz band from TRMM, DMSP and Aqua satellites, for Hurricane Danielle of 2010. Red and brown shading indicate deep convection. Times and dates are indicated beneath each image. Note that the time interval between images varies. Data courtesy of the US Naval Research Laboratories.

The simulation

We use a WRF simulation of a TC that includes a SEF and ERC, prepared as a nature run for data assimilation experiments and described by Nolan et al. (2013). That simulation nested the WRF model (Skamarock et al 2008) from 27km down to 1 km and covered the full life of the hurricane, although we will focus attention on the 48-h period beginning at 0000 UTC 3 August. Details of the simulation, including the choice of the nesting and initial fields, and the model setup including the physical parameterisations, are given by Nolan et al. (2013).

Figure 76 shows the evolution of relevant fields, azimuthally averaged, from the simulation. Panel (a) shows the latent heat release, which is an indication of the occurrence of deep convection. The initial contraction and intensification of the primary eyewall is apparent, as is the formation and subsequent contraction and intensification of the secondary eyewall, indicated by the magenta curve. Panel (b) shows the similar evolution of the near-surface gradient wind, with the formation of the primary wind maximum and its replacement by the secondary one readily apparent. Note also the general outwards expansion of the wind field during the process. Panels (c) and (d) show the convective available potential energy (CAPE), a measure of the favourability of the atmosphere for convective cloud formation, and the vorticity[2] of the gradient wind, respectively, and will be discussed later, as will the green curve.

[2] In an axisymmetric vortex, the vorticity is defined as $\zeta = v/r + \partial v/\partial r$, where v is the azimuthal wind and r is radius. More generally, it is defined as the vector curl of the wind velocity. Given the crucial importance of vorticity to understanding these processes, it is appropriate to give a physical interpretation. Imagine a paddle wheel suspended in the air, with its axis vertical, moving with the wind. It may rotate, for instance if the flow is curved (curvature vorticity), or if the wind is stronger on one side of the paddle wheel than the other (shear vorticity). If observed from a fixed coordinate system, rather than the earth, it also rotates because the earth is rotating (planetary vorticity, accounted for by the Coriolis parameter f), and the absolute vorticity is the sum of the earth-relative and planetary vorticities. In tropical cyclones, the curvature and shear vorticity are of opposite sign outside of the eyewall, and these two terms cancel to some degree, but the net vorticity is mostly of the same sign as f; that is, cyclonic.

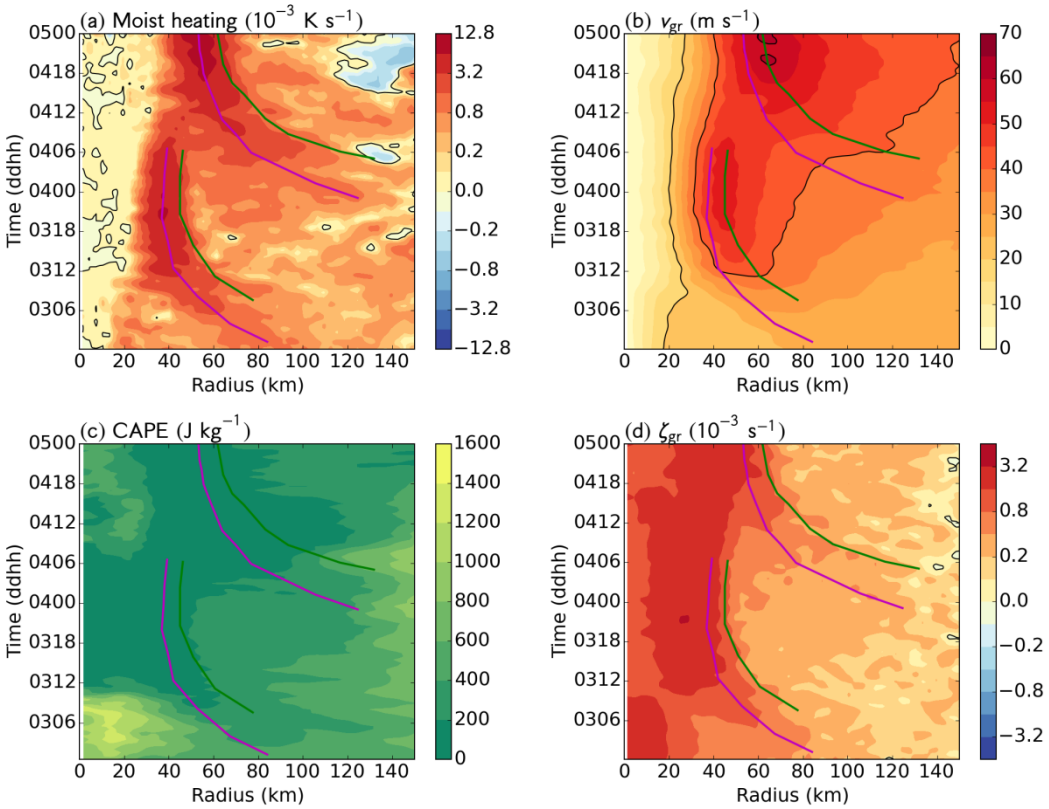

Figure 76: Time-radius diagrams of (a) the moist heating rate, averaged over 1-5 km height; (b) the gradient wind speed at 2-km height; (c) the surface-based CAPE; and (d) the vorticity of the gradient wind. All fields are azimuthally averaged. The magenta curves indicate the approximate location of the two eyewalls, as determined from the moist heating. The green curves indicate the approximate location of the radii of maximum negative radial vorticity gradient.

Analysis tools

The tropical cyclone flow can be divided into two components: the primary circulation, which is the rotational flow, and the secondary circulation, which consists of inflow mainly near the earth's surface, ascent mainly near the cyclone centre, and outflow mainly in the upper troposphere. The secondary circulation is often described as "in, up and out". It is forced by two main mechanisms: surface friction causing the near-surface flow to spiral inwards rather than being purely circular, and latent heat release in the clouds, particularly in and around the eyewall, which causes buoyant ascent. Mathematically, we adopt the convention that a radial flow component directed towards the centre of the storm (i.e. inflow) has a negative sign, whereas outflow is positive.

The frictional component of the secondary circulation, to a first order approximation, does not intensify the cyclone or change its structure, because the inwards advection of absolute angular momentum[3] nearly balances its destruction by surface friction (Kepert 2013). The heating component, in contrast, can intensify the cyclone since its lower branch advects absolute angular momentum inwards, spinning up the storm by the "ice-skater effect".

We assess the heating-induced component by the Sawyer-Eliassen equation, and the frictional component using a diagnostic boundary-layer model.

The Sawyer-Eliassen equation

The Sawyer-Eliassen equation (SEeq) diagnoses the secondary circulation due to latent heat release and other diabatic heating sources. It can also diagnose the effect of momentum sources and sinks, but may be less accurate for boundary layer friction due to the violation of a key assumption within the boundary layer.

We use the form of the SEeq given by Pendergrass and Willoughby (2009), discretised on a uniform grid with 75 grid points in radius, from 4 to 300 km, and 39 in height, from 0 to 19.5 km. The boundary conditions are that the flow perpendicular to the inner, upper and lower boundaries is zero, and that the flow is purely horizontal at r = 300 km. The discretised equation is solved directly using QR decomposition, which is feasible on modern computers with a problem of this size. Although not especially efficient, direct solution avoids the problem of potential nonconvergence of iterative methods should the SEeq be slightly non-elliptic, as sometimes happens in the upper-level outflow at large radius.

In this application, the inputs to the SE equation describing the structure of the storm and the heating are taken to be the azimuthal means, calculated directly from the WRF simulation output. Heating is taken to be the sum of latent heat exchanges and radiation.

[3] Absolute angular momentum is similar to the familiar concept of angular momentum but is calculated from a fixed frame of reference and therefore includes the earth's rotation. If r is the radius and v the azimuthal wind, then the absolute angular momentum is $M_a = rv + 0.5fr^2$, where the Coriolis parameter f is the local vertical component of the earth's rotation.

The boundary-layer diagnostic model

The boundary-layer diagnostic model has been described by Kepert and Wang (2001), and Kepert (2012, 2017). It solves the equations of fluid motion, with high-quality parameterisations of friction and turbulence, for a prescribed fixed pressure field representative of a tropical cyclone. It is run forward in time to a steady state, and the resulting flow represents the equilibrium boundary-layer flow given that pressure forcing. One day of model time usually gives a sufficiently steady state. For convenience, the pressure field is specified in terms of the gradient wind,

$$\frac{v^2}{r} + fv = -\frac{1}{\rho}\frac{\partial p}{\partial r},$$

where v is the azimuthal wind, r is radius, f is the Coriolis parameter, ρ is density and p is pressure.

In contrast to the SEeq, where the domain includes the full depth of the troposphere, here the domain is only 2.25 km deep, sufficient to contain the boundary layer.

As with the SE equation, the necessary input to this model of the gradient wind is calculated from the azimuthal mean of the WRF simulation output, at hourly intervals. The boundary-layer model's turbulence and friction parameterisations were configured to be reasonably consistent with those in WRF.

The boundary-layer model has been shown to be able to accurately reproduce the distribution of vertical motion in tropical cyclones when used in this way (Kepert and Nolan 2014, Zhang et al 2017).

Dynamics of the eyewall replacement cycle

Does frictional convergence influence the convection?

Figure 77a shows the azimuthal-mean vertical velocity at 1-km height from the WRF simulation. The evolution of the two eyewalls is clear, and closely follows that of the latent heat release in Figure 76a, as is usual in tropical cyclones. Figure 77b shows the diagnosed frictional updraft from the boundary-layer model, which clearly reflects a very similar pattern. There are two significant differences: the updraft in the boundary-layer model is consistently weaker and located at slightly larger radius. We will show shortly that the difference in strength of the updraft is due to the absence of buoyant convection in the boundary-layer model.

To confirm that boundary-layer frictional convergence is the cause of the convection, we have to eliminate other possible causes, in particular thermodynamic factors. Convective available potential energy (CAPE), shown in Figure 76c, is a widely used measure of the favourability of the atmosphere for convection, with higher numbers being more favourable. The values shown there are somewhat, but not strongly, favourable. There are no features in the CAPE field which would promote convection at the time and location of the secondary eyewall. Indeed, the CAPE at this time is amongst the lowest in the figure and decreases as the eyewall strengthens. This decrease is due to two reasons (not shown): cold downdrafts from convection cooling and drying the lower troposphere, and the broadening of the upper warm core reducing the equilibrium level for deep convection.

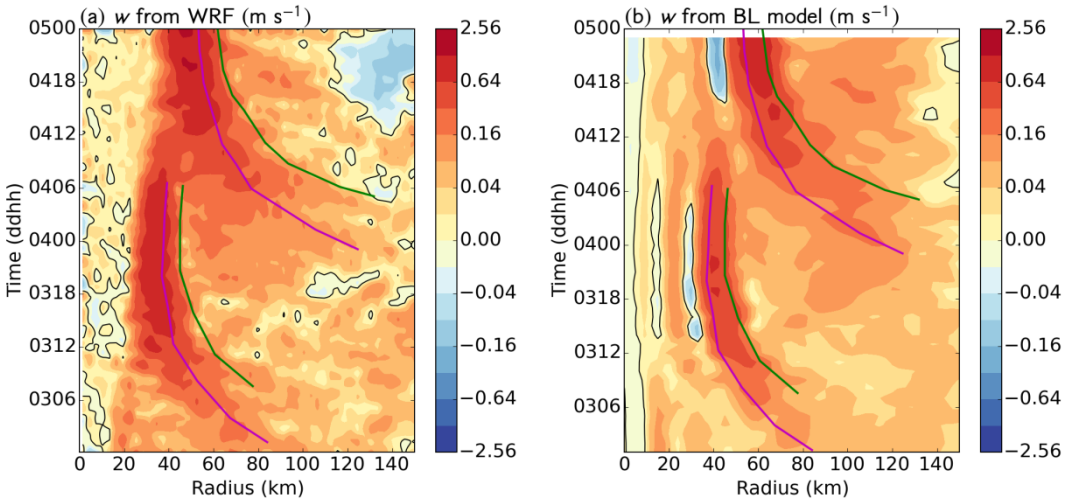

Figure 77: Time-radius plots of the azimuthal-mean vertical velocity at 1-km height, from (a) the WRF simulation and (b) the diagnostic boundary-layer model.

How does the convection alter the cyclone structure and intensity?

The modelled secondary circulation at four key times before and during the ERC are shown in Figure 78, together with the corresponding diagnosed fields from a preliminary calculation with the SEeq. There is generally quite good agreement. One systematic difference is that the diagnosed main updrafts are often too weak in the lowest 2 km of the atmosphere, likely due to the absence of friction in this calculation. Indeed, the updraft in this region appears to be forced by both latent heat release and friction, since both diagnostic methods underestimate it. The diagnostic calculation also has inflow around 3 km height and 80 km radius, whereas WRF has outflow. Again, this may be due to the absence of friction. There are also a number of technical issues still to be resolved with the SEeq calculation, whose contribution is presently unknown – we note that this is a preliminary calculation.

Nevertheless, there is quite good agreement with the WRF simulation. We may therefore appeal to earlier studies with the SEeq (Shapiro and Willoughby 1982, Willoughby et al 1982) to note that the expected storm evolution from this pattern would be for the outer wind maximum to intensify and contract, qualitatively consistent with the evolution in Figure 76b. Future work will examine the extent to which quantitative agreement is obtained.

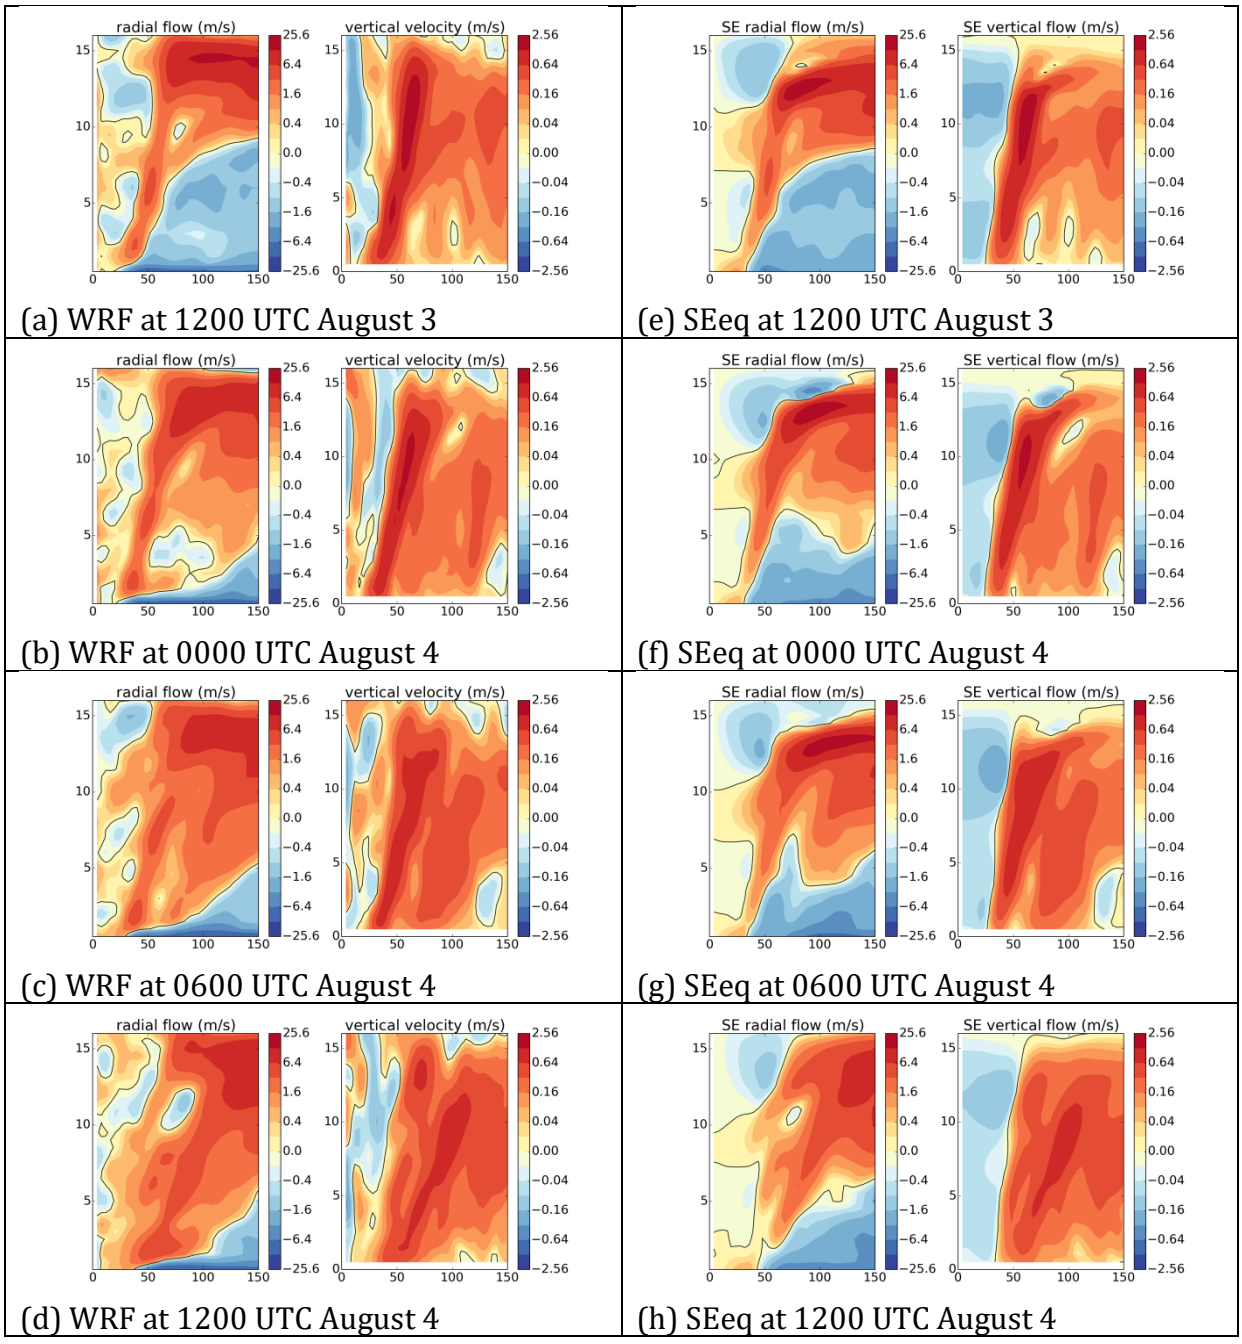

Figure 78: Secondary circulation during the ERC and SEF. Panels (a – d) show the azimuthal-mean radial wind (left, positive is outwards) and vertical wind (right, positive is upwards) from the WRF simulation at times as indicated. Panels (e – h) show the corresponding diagnosed secondary circulation using the SEeq.

How does the cyclone's pressure field affect the frictional convergence?

We have seen above that the diagnosed frictional convergence determines the location and intensity of the convective heating. The only data passed from WRF to the diagnostic model in this calculation is the pressure field, in the form of the gradient wind. The question, then, is what characteristic of the pressure field leads to localised updrafts outside of the primary eyewall?

In Figure 77, we saw that the diagnostic BL model reproduces the location and relative strength of the two eyewall updrafts reasonably accurately, although underpredicts the strength because in the WRF simulation that is enhanced by the additional forcing from buoyant convection. It is also apparent in this figure that the changes in the gradient wind are spatially smoother than the vertical velocity response and are relatively subtle. Pressure is the radial integral of the gradient wind equation, so the changes in pressure must be even smoother than for gradient wind[4]. Nevertheless, these subtle changes must be responsible for the changes in the frictional forcing of the updraft, because they are the only information passed to the boundary-layer model. In other words, we know **that** the pressure field affects the frictional convergence, and we need to determine **how**.

Kepert (2001) developed a simplified diagnostic tropical cyclone boundary layer model. Compared to the model used here, the simplifications included a linearisation, and adoption of less realistic representations of turbulent diffusion and the air-sea momentum transfer. While calculations with this model are expected to be less accurate than from the full diagnostic model, they offer the great benefit that an analytic solution is available. That is, we can directly examine the equations to understand how the vertical velocity relates to the pressure (or gradient wind) structure.

Kepert (2013) used that model to show that near the eyewall(s) of typical tropical cyclones, the updraft is approximately proportional to the radial gradient of the vorticity of the gradient wind, multiplied by the drag the wind exerts on the sea surface (approximately proportional to the square of the wind speed), divided by the square of the absolute vorticity. This equation suggests that we can expect to find enhanced updrafts where there is a locally strong negative radial vorticity gradient. Such gradients are especially effective if they occur where the vorticity is relatively low, because of the division by the square of the vorticity. The surface friction has a lesser effect on locating the updraft, because it varies relatively slowly with radius.

Discussion and conclusions

Figure 76b,d and Figure 77a compare the joint evolution of the vorticity of the gradient wind, the frictional updraft, and the moist heating. These figures show the mutual contraction of the region of strong vorticity gradient and the region of strong convective latent heat release. Indeed, a similar relationship is also apparent earlier in the cyclone's life, during the initial contraction and

[4] Differentiation acts as a high-pass filter, emphasising the small scales, as can be easily shown using the Fourier transform.

intensification of the primary eyewall between 0000UTC and 1200UTC on August 3.

Clearly, these features are strongly correlated. However, correlation is not causation. Two phenomena, A and B may be correlated may occur for several reasons: A causes B, B causes A, or that both are caused by some third factor C.

In this chapter, we have provided additional information that does allow us to attribute cause. In particular, using the diagnosed frictional convergence, which depends only on the cyclone's pressure field, we have shown that the distribution of vertical velocity at the top of the boundary layer is determined largely by frictional processes, although friction is insufficient to explain the full magnitude of the ascent. By comparing the evolution of the diagnosed frictional convergence to that of the convective latent heat release, together with the absence of any features sufficient to explain the localisation of the convection in the stability or moisture fields, we show that the frictional updraft is largely determining the location and strength of the convection. Calculating that part of the secondary circulation induced by heating, using the Sawyer-Eliassen equation largely accounts for that part of the low-level updraft missing from the frictional calculation. The evolution of the vortex structure explained by advection of absolute angular momentum by this secondary circulation is largely consistent with the evolution of the cyclone; that is, it leads to changes in the gradient wind that are similar to those in the WRF simulation. Most importantly, these include an inwards migration and strengthening of the vorticity features that we have theoretically linked to the evolving frictional updraft.

To summarise, we have confirmed that the positive feedback mechanism hypothesised by Kepert (2013) and further discussed by Kepert and Nolan (2014) and Kepert (2017) is indeed operating in this case. Here, vorticity-induced boundary-layer convergence acts to promote convection, provided that the stability and moisture are also favourable. Vortex-stretching in convective updrafts increases the local vorticity. We have analysed the combined, cyclone-scale, effect of many individual clouds by applying their combined heating to the Sawyer-Eliassen equation, but analyses at the cloud scale (not shown) similarly show that convective updrafts are acting to increase the vorticity beneath the developing secondary eyewall. These vorticity changes in turn further strengthen the frictional updraft. There is a further important subtlety in all of this, in that the relative location of the various processes is important for the precise details of the interaction, particularly the initial rapid contraction of the outer eyewall, followed later by slower contraction and intensification. These can be largely explained as a contribution of nonlinearity in the boundary layer, as detailed in Kepert (2017).

Practical implications

Forecasting ERCs is challenging, because it is clear from the work described here that the changes in the early stages are quite subtle, and therefore difficult to detect. This is especially true in the harsh environment of tropical cyclones, where observations are difficult to take. Modelling ERCs is likewise challenging, for this work implies that the interaction between friction, clouds and pressure in the cyclone must be represented with sufficient fidelity. Nevertheless, the success of

the simulation used here and others shows that that fidelity has been achieved in current NWP systems.

The initialisation of such simulations is challenging, especially since the early signs of an ERC are subtle. Small errors in the initial pressure field could either completely remove the local vorticity perturbation and hence the frictional updraft, or add a spurious one. The necessary precision in the initialisation will be beyond the reach of our observing and data assimilation systems, at least in the absence of aircraft reconnaissance, for some time. Ensemble prediction methods provide the only presently viable means of dealing with this uncertainty.

In spite of these considerable difficulties, forecasting ERCs is important. They represent a substantial additional difficulty for intensity forecasting, because of their large impact on the intensity evolution of the storm. They also strongly affect the ocean response. A substantial import of vorticity to the storm in the region of the developing outer eyewall seems to be an inherent part of the ERC process, and this import of vorticity explains the wind field expansion. The wind field expansion affects not just the width of the damage swath, but also the timing and duration of damaging winds. It also profoundly increases the ocean hazard, because applying strong winds to a larger area of the ocean's surface greatly increases the severity and extent of storm surge, damaging waves and coastal erosion. With the high concentration of vulnerable populations and infrastructure near the coast in Australia, improving our ability to predict ERC will clearly help mitigate tropical cyclone impacts.

FINDINGS

The previous chapter presented detailed summaries of the numerous subprojects within this project, each of which had its own set of findings. Rather than repeating those here, in this chapter we aim to consolidate and synthesise findings across subprojects and present a "big picture" view of the project's achievements and significance.

HIGH-RESOLUTION NUMERICAL WEATHER PREDICTION (NWP) AND ENSEMBLES

We conducted four case studies of severe weather events, using the Bureau's ACCESS system at the finest resolution of between 400 m and 1.5 km, depending on the event. A fifth study used the WRF model on a tropical cyclone. In each case, we were able to discern and analyse details within the simulated meteorology which plausibly contributed to the severity of the event. These events included two fires (the State Mine fire, and the Tathra fire), and two severe mid-latitude weather events (the Dungog East Coast Low, and the South Australian tornado outbreak) as well as the tropical cyclone. The two midlatitude events were studied using ensemble simulation.

For the State Mine fire, several mesoscale features were identified which could account for the rapid spread of the fire. The first was a mountain wave structure with very strong downslope winds developing over the fire ground during the afternoon of 17 October 2013, accompanied by hot and dry atmospheric conditions. The second was a slot of dry air with an approximate width of 50 km that passed over the fire ground around and shortly after midday which likely exacerbated the fires. The third was an undular bore that formed during the night of 16 October, which was related to the formation of the dry slot. The formation and evolution of these mesoscale features are described, and existing theories are considered to explain these phenomena.

The Tathra fire study combined observations with very high-resolution modelling to understand the mechanisms behind the disaster. At several points in the fire progression, atmospheric phenomena were active which greatly increased the risk. These were (a) mountain waves causing strong surface winds during the ignition, (b) the onset of boundary layer rolls in conjunction with weaker mountain waves interacting to intensify the fire and promote ember transport as it jumped the Bega river on its approach into Tathra, and (c) a complex interaction between the boundary layer rolls in the north-westerly flow and the wind change propagating up the NSW south coast leading to erratic conditions around Tathra in the late afternoon. The study points to the value of such modelling in understanding these processes, and to the potential for better managing such events in real time if we can deploy such modelling along with the necessary tools and training to interpret it in real time into operations.

The Dungog East Coast low simulations captured the event very well, accurately predicting the rainfall extent and the timing and duration of the strong winds. Probabilistic analysis of the output showed the highest probabilities of extreme rain where it occurred, but also lower risk in adjacent regions. This much more

detailed depiction of the risk would have helped emergency services pre-plan, including indicating safer areas for potential evacuations.

Our initial non-ensemble simulation of the South Australian event correctly captured the time and location of the thunderstorm mesocyclone that spawned one of the observed tornadoes. Events at such small scales as tornadoes are not normally regarded as being predictable in a deterministic framework, so we extended that simulation to be a 6-member ensemble. These showed a fascinating range of intensity and structure of the event, in which all had at least some tornado risk, but the level and extent of the risk varied.

Many potential users of such simulations are aware that the computational demands are high, such that they are not yet available in real time. However, our studies also highlight the difficulty of extracting data in a timely fashion and putting it into a form that can be communicated to users and improve outcomes. Often, the conditions that lead to adverse outcomes are small in scale and locating them within the voluminous output takes time and skill. The Tathra fire was a good example of this.

However, some of this modelling capability is available in operations now. The Bureau has recently begun running operational limited-area ensembles on a 2.2 km grid in parts of eastern Australia. The configuration is similar in many respects to that which we used for the Dungog East Coast Low simulation (which was the first high-resolution ensemble simulation run in the Bureau). We expect that further growth in our analysis capabilities will continue to improve the ability to deliver excellent real-time outcomes from this resource.

The simulations of the Dungog event were also used in the CRC Impact Prediction project, a joint project between the Bureau and Geoscience Australia which aimed to predict physical impacts on structures and the environment from the weather event, thereby providing a more user-centric, outcome-oriented forecasting service. The ACCESS nesting suite used in most of these case studies is part of the overall ACCESS infrastructure maintained by the Bureau, and we are grateful to our colleagues in Earth System Modelling and at the UK Met Office for their wonderful work. It also forms the underpinning framework for the CRC *Coupled fire-atmosphere* project.

BUSHFIRE PLUME DYNAMICS

We conducted an extended research program exploring bushfire plume dynamics, with the main focus on two important problems: forecasting pyrocumulonimbus clouds and predicting ember transport.

For each target problem, we quickly realised that crucial underpinning knowledge was missing. We used large eddy simulation (LES) to address this gap. LES allowed us to simulate the turbulent structure of large bushfire plumes, spanning the full depth of the atmospheric boundary layer (or more) and extending many kilometres downwind from the fire. We used these simulations firstly to drive an ember transport model, and later to simulate the development of pyroconvective clouds.

Our ember transport simulations showed the crucial role of turbulence, which could more than double the predicted maximum transport distance as well as driving lateral dispersion and thereby increasing the width of the ember swath. Importantly, turbulence is not explicitly included in any of the current operational spotfire prediction techniques. Turbulence is also a key component of entrainment into the plume, the process by which the plume engulfs environmental air, causing it to expand and cool, thereby greatly affecting its vertical development. Our analysis of entrainment, in conjunction with a consideration of the thermodynamics of combustion and of cloud formation, demonstrated that nearly all the moisture in pyroconvective clouds comes from the atmospheric environment, not from combustion or the fuel. This crucial fact, although controversial, pointed the way towards predicting the formation of these clouds.

The challenge in predicting pyrocumulonimbus (pyroCb) formation is that it depends on both the fire and the atmosphere, and information from both scientific disciplines needs to be combined. We developed a paradigm, the PFT or pyrocumulonimbus fire-power threshold, that framed the question in this way. Specifically, we ask what fire properties will lead to pyroCb in any given atmospheric environment. Guided by our LES and thermodynamic work, we proposed that the use of a fire power threshold was the answer. Fires hotter than the threshold would lead to pyroCb, but not fires that were colder. We also demonstrated how to use a simplified plume model to estimate an approximate PFT from meteorological data (either model or observed soundings). We developed techniques to present maps of these data to users, and after extensive testing on historical events, ran a real-time trial over the summer of 2019-20. The prototype system performed extremely well over that season, and rapidly became widely used during the summer. Of the 30-odd events, it gave warning of nearly all. Analysis of the exceptions has led to further improvement of the tool, and it was awarded CRC utilisation funding to take it the next step closer to operations.

Our work on a simple model, or parameterisation, of ember transport similarly used a simplified plume model, albeit a more sophisticated one than for the PFT. The model consists of four parts: the plume model, a model of turbulence within the plume, a probabilistic model of ember transport in the plume and a model of ember falling and transport beneath the plume. Importantly, the scheme is computationally very fast, sufficiently so to be suitable for use in coupled and uncoupled fire-spread models, as well as for standalone use. Our results support recent suggestions, based on an analysis of line-scan observations, that there are (at least) two regimes in long-range ember transport, and we are optimistic that it will be possible to develop a tool to diagnose which regime – or perhaps where on the spectrum – a fire is in.

We explored the sensitivity of predictions from the technique to environmental conditions, showing that inversion height (or boundary layer depth) is a crucial parameter. Extreme-range transport requires a deep boundary layer. The parameterisation was implemented to run in conjunction with the CSIRO Spark fire-spread model, greatly improving predictions of the notorious Kilmore East fire of 2009 and correctly predicting the observed maximum spotting distance in this event.

Each of these activities, directed towards improving predictions of aspects of fire-atmosphere behaviour, relied crucially on LES as a tool to develop understanding of plume dynamics. We developed and documented techniques to facilitate this, which we are confident will continue to prove useful into the future.

BRINGING THE TWO STRANDS TOGETHER

These two strands of work on bushfire plumes also points to a broader truth about future operational solutions to problems. High-resolution NWP is a wonderful resource but is not the whole answer. Extracting useful information in a timely fashion requires sophisticated tools, some of which need new scientific knowledge to underpin them. The PFT and the ember transport parameterisation each unlock new information from NWP, providing answers that are implicit within the resource, but require work to realise. We look forward, in the future, to considering how these tools will extend value in an ensemble context.

Often, the dangerous conditions occur in short windows, or over very limited areas, or require a fire ignition to turn the potential danger into reality. They may be present in just a few ensemble members. Improving our interpretative abilities will continue to yield dividends.

KEY MILESTONES

The project has achieved all its milestones. This section briefly discusses the major accomplishments.

We conducted five case studies of severe weather events, two fires, two mid-latitude events and one tropical cyclone. The latter two events incorporated ensemble simulations, two of the first high-resolution ensemble simulations conducted within the Bureau. The results were widely communicated in a variety of formats and to a variety of audiences. The forms included academic journal publications; conference extended abstracts, oral presentations, and posters; and other talks and seminars. These communications focussed on the lessons learned, perhaps most strongly in the case of the Tathra fire where our research results were discussed by RFS personnel at the coronial enquiry. Audiences included specialists in meteorology, emergency services and fire.

Our idealised studies of plume dynamics using large eddy simulation (LES) resulted in the first such simulations of long-range ember transport, and some of the first of pyrocumulus formation. Both were the subject of comprehensive communication with similar methods and audiences to the case studies.

These two strands of the plume dynamics work have each led to substantial milestones in the form of simpler models, informed and calibrated by the LES work. For pyroconvection, we have achieved the ambitious goal of a forecasting tool, the pyrocumulonimbus firepower threshold (PFT). This has been extensively communicated, with highlights including three articles in premium academic journals, and extensive training of staff from the Bureau of Meteorology and fire agencies in its use. It was also the subject of a highly successful real-time trial, which again has led to a comprehensive report.

The second strand, ember transport, similarly led to a practical forecasting tool, implemented in a fire-spread model and also suitable for other applications. This work completed later than the PFT, and the COVID pandemic has limited our ability to widely communicate it. Nevertheless, a journal article has been prepared and several conference presentations given. Further communication will follow as restrictions ease.

The final major deliverable of the project is, or course, this report. We thank readers who have read to this point and encourage those who may have skipped over the detailed results in previous sections to backtrack.

UTILISATION AND IMPACT

SUMMARY

We conducted a highly successful real-time trial of the pyrocumulonimbus firepower threshold (PFT) as a major utilisation exercise over the 2019-20 fire season, which is discussed in detail in the Results section of the report. Several of our user testimonials are based on that trial.

A detailed report on that trial was prepared, which provides a substantial learning and policy development resource for further use of the technique.

Our pyroCb outreach articles and blogs have garnered significant media interest, including one blog post having global reach. The PFT has also featured in a science magazine article, and another Bureau of Meteorology blog.

We have begun a utilisation project to take the PFT work further, detailed below.

Working with CSIRO and AFAC, we are putting the necessary things in place to be able to convert the prototype implementation of the ember transport parameterisation in Spark, into an operationally robust and supported system. This will both facilitate further research, and provide the framework to allow users to learn, understand and apply the system.

Much of the output of our research is knowledge, and we have worked towards presenting that in a form that is accessible to practitioners as well as researchers. To this end, we have given numerous seminars and presentations to operational folk, as well as preparing targeted documentation, particularly on the PFT and the Tathra case study. Several of the presentations have been recorded and are available as a permanent record.

David Wilke spent 5 months over the recent fire season working in operations, greatly improving the transfer of knowledge between the research and operational communities in both directions. A summary of that work is also provided below.

PRE-OPERATIONALISATION OF THE PYROCUMULONIMBUS FIREPOWER THRESHOLD (PFT) PYROCB PREDICTION DIAGNOSTIC.

Following the success of the real-time trial of a PFT forecasting diagnostic during the 2019/2020 southern Australian fire season, there has been overwhelming support from fire-weather forecasters and fire-behaviour analysts to improve usability for the next fire season. Planned improvements include: increasing the output frequency, optimising display features and improving access to forecast products. Additional fine-tuning and adjustments to the suite of products are also planned.

A utilisation proposal was submitted late in 2019, and an initial project agreement between the Bureau of Meteorology and the CRC signed in early 2020. In March, following the real-time trial of the PFT and consultation with end-users, an extension to the project was proposed, which was signed off between the two organisations in June 2020. Work began in October.

Output Description

This utilisation project aims to implement, as much as possible, the feedback development ideas raised by users. This will include:

- Substantial improvements in computing efficiency to enable more forecast domains to be run, including the routine running of high-resolution forecasts from operational numerical weather prediction models.
- Improved labelling to more easily identify threat location and timing
- Password-protected web access to forecast products
- Tuning to eliminate false triggering of the PFT flag during cold outbreaks.
- Incorporation of buoyancy losses from entrainment in the cloudy part of the plume.

Additional improvements will be added in consultation with end-users as the project progresses.

PFT Media

Pyrocumulonimbus events and the PFT gathered significant media interest, given the very active fire season recently experienced. One blog post went global. Notable activity includes:

- BoM Blog: "Predicting fire-generated thunderstorms". http://media.bom.gov.au/social/blog/2355/predicting-fire-generated-thunderstorms/
- BoM pyroCb explainer video. https://www.facebook.com/watch/?v=2709867255956548
- BoM pyroCb blog post went global during the peak of the fires. http://media.bom.gov.au/social/blog/1618/when-bushfires-make-their-own-weather/
- Article in the science magazine "Particle": When bushfires create their own weather. https://particle.scitech.org.au/earth/australias-bushfires-creating-firestorms/
- A radio interview with Paris Brooke from 3MGB a Mallacoota public radio station.

PFT Invited presentations

NSW RFS Fire-behaviour analyst workshop, August 2019

Bureau of Meteorology Annual Research and Development workshop, 25—29 November 2019.

Naval Research Laboratory in Monterey California, USA*

National Interagency Fire Center in Boise Idaho, USA*

Missoula Fires Science Laboratory in Missoula Montana, USA*

*Scheduled for May 2020, postponed due to COVID-19 Lockdown

International Impact

Early in 2020, Nick Nauslar from the Storm Prediction Center in Oklahoma USA (at that time) discovered the PFT training video and began using the manual PFT calculation method to assess pyroCb potential during the 2020 USA fire season. Nick was seen by his peers to have developed an uncanny insight into pyroCb formation, and he began sharing the PFT concept among fire-weather forecasters and fire-behaviour analysts throughout the USA. Further exposure of the PFT in the USA included seminars Kevin delivered to fire-weather staff from the Naval Research Laboratory, and the 2020 Predictive Services Annual Meeting. There has been strong interest in setting up the PFT to run in the USA, with negotiations underway with people from the National Severe Storms Laboratory to add the PFT to their suite of experimental forecast products.

OPERATIONAL SUPPORT

Over the summer David Wilke was recalled to operations in the New South Wales State Forecasting Centre for a significant period to aid in the provision of severe weather services. This work was primarily as an embedded meteorologist in the NSW Rural Fire Service, a role which requires frequent liaison with fire behaviour analysts. The experience, whilst initially detrimental to the project milestone completion, has produced several clear benefits for the CRC. In particular, it has fostered deeper connections between the team and the RFS, facilitating a number of discussions regarding CRC projects. Information gathered relating to the Tathra fire has been crucial to the teams understanding of the key fire behaviour of the event and has greatly informed the method of analysis. The period of operational forecasting also provided a perfect environment in which to discuss the PFT with fire behaviour analysts and meteorologists. The diagnostic was regularly included in routine briefings throughout the real-time trial conducted over the summer. This constant assessment of the diagnostic, with information fed back to the team, significantly enhanced the CRC-projects connection to end-users.

EMBER TRANSPORT PARAMETERISATION

The transport model has been implemented to run in conjunction with CSIRO Spark fire spread model, and is producing excellent results, albeit on limited testing so far, as detailed elsewhere in this report. This success has gathered significant interest from the AFAC Fire Prediction Services. Development of an operationally robust version will require suitable agreements with CSIRO and funding. We have been working with the business manager of AFAC Fire Prediction Services, John Bally, and with CSIRO to enable this next step.

TRAINING AND COMMUNICATION WITH PRACTITIONERS

Much of the output of our research is knowledge, and we have worked towards presenting that in a form that is accessible to practitioners as well as researchers.

To this end, we have given numerous seminars and presentations to operational folk, as well as preparing targeted documentation.

Hazard Notes

We undertook a significant utilisation activity in 2016 with the publication of a *Hazard Note* on the effects of mountain waves on bushfires. This work drew on our analysis of the State Mine fire, as well as previous work on the Ellensbrook and Aberfeldy fires for the former Bushfire CRC. In it, we discussed how mountain waves can lead to strong winds on the lee slope of a hill or mountain, which are often stronger at night, contrary to the normal diurnal cycle. This *Hazard Note* drew considerable attention, and the material has subsequently been adapted for the Bureau of Meteorology's blog, for a feature article in *Fire Australia*, and for an article in the *Weekly Times*. It can be difficult to foresee where a piece of research will lead, but we are very pleased that this work has clearly struck a chord with people exposed to or responsible for managing fire risk, and we trust it will lead to improved mitigation in the future.

We published two additional *Hazard Notes* that year, on next-generation fire modelling and on monitoring and predicting natural hazards.

SUPPORT OF ENQUIRIES

The study of the Tathra fire event was quoted in evidence to the coronial inquiry by two RFS personnel, helping to provide a rich and scientifically rigorous discussion of the meteorology of the event to the coroner.

UTILISATION AND IMPACT EVIDENCE

Please refer to the product user testimonials on page 12 of this report.

CONCLUSION

The project has studied a wide range of hazardous meteorological phenomena, hazardous either on their own or in conjunction with bushfires. We have used state-of-the-art high resolution modelling for this purpose, including ensemble modelling and large-eddy simulations. We have also used simplified models of bushfire plumes and the tropical cyclone circulations, and direct and remotely-sensed observations of the phenomena of interest.

We have studied long-range ember transport and found that turbulence within the bushfire plume is an important factor, roughly doubling the maximum distance that embers can be carried and substantially increasing the spread of the landing distribution. The landing distribution has substantial crosswind spread in light winds, but little in strong winds. Ember fallout is strongly "clumpy" in character, due to the turbulence within the plume.

We built on these findings to develop a computationally cheap model of bushfire ember transport. This model was implemented as a parameterisation within the CSIRO Spark fire-spread simulator and applied to the Kilmore East fire of 7 February 2009 (Black Saturday), greatly improving the simulations. It was also used to examine the sensitivity of long-range transport to atmospheric and fire conditions, where it highlighted the importance of inversion height (or boundary layer depth) in limiting the transport distance.

We studied the formation of pyrocumulus clouds in two frameworks, one involving detailed simulation and the other a consideration of the underlying thermodynamics. These have shown that, most of the time, the moisture from the combustion process plays a negligible role in the formation of pyrocumulus. We have also identified the rare situations in which this is not so. This outcome was most encouraging, for it shows that forecasting methods for pyrocumulus can mostly ignore the combustion moisture, which is an exceedingly difficult parameter to quantify.

We used this new knowledge to develop a new forecasting technique for pyrocumulonimbus formation, the pyrocumulonimbus firepower threshold or PFT. Firstly, we defined the PFT as a suitable paradigm for pyroCb forecasting, in which fires where the power is greater than the PFT initiate pyroCb, but cooler fires do not. This is an operationally useful framework in which to work since it separates meteorological and fire conditions. Secondly, we developed a simple and approximate way of calculating PFT. The third step was to derive the PFT-flag, which indicates when weather conditions are favourable for both pyroCb and for intense fires.

The PFT was tested in a real-time trial over the summer of 2019-20, which was an outstanding success. Of the 30-odd pyroCb observed, the tool gave warning of nearly all. It also gave good indication of when pyroCb were not expected, even though intense fires were present at the time. The few unsuccessful forecasts were analysed and will lead to further improvements in the technique in the future. We are planning additional work to use and extend this tool to examine the causes of the marked worldwide increase in observed pyrocumulonimbus in recent years, with funding from other sources.

The PFT is currently being bought closer to operations, with the assistance of utilisation funding from the CRC.

We prepared a detailed case study of the east coast low of April 2015, which caused severe wind damage and flooding to the Hunter Valley and surrounding regions. This study was novel in its use of a 24-member high-resolution ensemble, rather than just a single simulation. The use of the ensemble facilitated certain diagnostics, as well as providing an excellent example of some of the advantages of ensemble prediction, including increased accuracy, probabilistic forecasts and reduced likelihood of missing significant developments.

We similarly used high-resolution ensemble prediction to explore the SA tornado outbreak of September 2016, which destroyed critical electricity infrastructure and led to the state blacking out. The ensemble simulation showed an intriguing spread of convective mode, severity, and extent, but all members had at least some indication of tornadic potential on the diagnostics used. One member very closely resembled the fine details of the observed event, including the location of one destructive tornado.

Our first bushfire case study was the Blue Mountains fires of October 2013, focusing on the day of the most damaging fire activity, on which some 200 houses were lost. We identified two meteorological factors that likely contributed to this disaster, the development of mountain waves causing strong surface winds, and a dry slot of unusual origin. This is the third case we have studied with a clear mountain wave signature in the fire escalation, so we prepared a CRC *Hazard Note* as well as other publications, to communicate these findings. By raising awareness, we hope that weather forecasters and fire fighters will be better able to recognise these events and respond appropriately.

We also analysed the Tathra fire of March 2018, using a single high-resolution simulation. Through this analysis, we found several mesoscale features which likely contributed to the extreme fire event. Key of these was the development of coherent horizontal convective rolls which emanated from inland and aided the fires spread toward Tathra. The rolls interacted with the terrain to produce complex regions of ascending and descending air, likely accelerating the lofting of firebrands and potentially contributing to the significant lee slope fire behaviour observed. Mountain waves, specifically trapped lee waves, occurred on the day and are hypothesised to have contributed to the strong winds around the time the fire began, and may have influenced conditions during the period of peak fire activity when the fire spotted across the Bega River and impacted Tathra. Finally, the passage of the cool change through the fireground was complex, with frontal regression observed at a nearby station and likely also through Tathra. We postulate that interactions between the rolls and the initially weak change resulted in highly variable and dangerous fire weather across the fireground for a significant period after the change was initially incident

Our final case study was the analysis of a simulation of a tropical cyclone eyewall replacement cycle. We have developed a clear understanding of the processes involved, a complex and subtle interaction between clouds, surface convergence and the local rotation of the flow. The critical phase is the formation of the outer eyewall, and our new knowledge indicates the critical factors to being able to predict this. Eyewall replacement cycles are important,

for they affect several aspects of the tropical cyclone's impact, including the intensity, extent of damaging winds, and storm surge.

Numerous factors are placing an increased burden on the ability of the Australian community to cope with natural hazards. These include greater exposure to risk, increased societal expectations, and climate change. Advanced modelling has come to be an indispensable tool in meteorology and oceanography for supporting forecasts, warnings and action, and is gaining a place in fire prediction also. Ensuring that the modelling reflects the best possible science, makes efficient use of expensive computer resources, and is appropriately communicated to and used by those responsible for mitigation, is a critical task.

However, the advanced numerical weather prediction systems along do not provide the complete answer. Often, advanced diagnostics are needed to extract the vital information from the model run. With the PFT, we developed a new tool to aid such endeavours. Our analyses of the case studies exhibited other sophisticated tools, including updraft helicity, the Okubo-Weiss parameter, and ensemble-based diagnostics.

NEXT STEPS

The project has produced two significant operational decision support tools. The first, the PFT, is presently the subject of a utilisation project. There are several avenues in which the second, the simple ember transport model, could be bought into regular use. One is through incorporation into a fire spread simulator such as CSIRO's Spark. We have a demonstration version of this, but work would be needed to achieve an operationally robust version. Similarly, ember transport could be incorporated into ACCESS-Fire.

These two activities, as well as the 5 case studies, have produced a significant volume of learnings. We are currently ensuring that these are captured and made useable. Journal articles are in the publication pipeline, and seminars and conference presentations will be given. Some of the material has been offered for inclusion in the Firepedia, and some has been used for training both meteorologists and fire practitioners. These activities will continue.

We have also developed a significant amount of software to support the analysis of high-resolution model simulations. We are presently in the early stages of case studies of 5 fire events from the 2019-20 season using ACCESS-Fire, and will use this exercise to further refine our display and analysis suite. There would be significant value in taking this a step further, to a semi-automated, quick turnaround analysis capability for post-analysis of severe weather events, but this step would require specific support.

PUBLICATIONS LIST

PEER-REVIEWED JOURNAL ARTICLES

Kepert, J. D. and D. S. Nolan, 2014: Reply to "Comments on 'How Does the Boundary Layer Contribute to Eyewall Replacement Cycles in Axisymmetric Tropical Cyclones?'". *J. Atmos. Sci.*, **71**, 4692–4704, doi:10.1175/JAS-D-14-0014.1

Kepert, J. D., 2017: Time and space scales in the tropical cyclone boundary layer, and the location of the eyewall updraft. *J. Atmos. Sci.*, **74**, 3305–3322, doi:10.1175/JAS-D-17-0077.1

Kepert, J. D., 2018: The boundary layer dynamics of tropical cyclone rainbands. *J. Atmos. Sci.*, **75**, 3777–3795, doi:10.1175/JAS-D-18-0133.1.

Kepert, J. D., K. J. Tory and W. Thurston, 2000: A simple, physically based model of ember transport within bushfire plumes. *J. Atmos. Sci.,* in internal review.

Peace, M., T. Mattner, G. A. Mills, J. D. Kepert, and L. McCaw, 2015: Fire modified meteorology in a coupled fire-atmosphere model. *J. Appl. Met. Clim.*, **54**, 704–720, doi:10.1175/JAMC-D-14-0063.1

Peace, M., T. Mattner, G. A. Mills, J. D. Kepert, and L. McCaw, 2016: Coupled fire-atmosphere simulations of the Rocky River fire using WRF-SFIRE. *J. Appl. Met. Clim.*, **55**, 1151–1168, doi:10.1175/JAMC-D-15-0157.1.

Peace, M., L. McCaw, J. D. Kepert, G. A. Mills and T. Mattner, 2016: WRF and SFIRE simulations of the Layman fuel reduction burn. *Australian Meteorol. Oceanogr. J.*, **65**, 302–317.

Peace, M., L. McCaw, B. Santos, J. D. Kepert, N. Burrows, and R. J. B. Fawcett, 2017: Meteorological drivers of extreme fire behaviour during the Waroona bushfire, Western Australia, January 2016. *J. S. Hem. Earth Sys. Sci.*, **67**, 79–106, doi:10.22499/3.6702.002

Thurston, W., R. J. B. Fawcett, K. J. Tory and J. D. Kepert, 2015: Simulating boundary-layer rolls with a numerical weather prediction model. *Quart. J. Roy. Meteor. Soc.*, **142**, 211–223.

Thurston, W., K. J. Tory, R. J. B. Fawcett and J. D. Kepert, 2015: Large-eddy simulations of pyro-convection and its sensitivity to environmental conditions. Submitted to *Austral. J. Emergency Mgt.* special issue for AFAC Conference. Not published since the journal special issue was cancelled.

Thurston, W., Kepert, J. D., Tory, K. J. and Fawcett, R. J. B., 2017: The contribution of turbulent plume dynamics to long-range spotting. *Int. J. Wildland Fire*, **26**, 317–330, doi:10.1071/WF16142

Tory, K. J., W. Thurston and J. D. Kepert, 2018: Thermodynamics of pyrocumulus: A conceptual study. *Mon. Wea. Rev.*, **146**, 2579–2598, doi:10.1175/MWR-D-17-0377.1.

Tory, K. J., 2019: Pyrocumulonimbus firepower threshold: A pyrocumulonimbus prediction tool. Australian Journal of Emergency Management Monograph No. 5, December 2019, pp 21—27. Extended abstracts from *the Bushfire and Natural*

Hazards Cooperative Research Centre Research Forum. https://www.aidr.org.au/media/7379/monograph-no5-extended-abstracts-final.pdf

Tory, K. J., 2019: Predicting Fire Thunderstorms. Asia Pacific Fire Magazine. Issue 72, pp36—38. https://apfmag.mdmpublishing.com/predicting-fire-thunderstorms-2

Tory, K. J., and J. D. Kepert, 2020: Pyrocumulonimbus Firepower Threshold: A pyroCb prediction method. *Weather and Forecasting,* In revision.

Wilke, D. J., J. D. Kepert and K. J. Tory, 2020: The meteorology of the Tathra bushfire. Mon. Wea. Rev., internal review.

Zovko-Rajak, D., K. J. Tory, R. J. B. Fawcett, J. D. Kepert and L. J. Rikus, 2018: High-resolution ensemble prediction of the Australian East Coast Low of April 2015. *J. South. Hemisph. Earth Sys. Sci.,* 165 – 183.

Zovko-Rajak, D., K. J. Tory and J. D. Kepert, 2020: A case study of South Australia's severe thunderstorm and tornado outbreak 28 September 2016. *J. Southern Hemisphere Earth System Science,* internal review.

PEER-REVIEWED JOURNAL ARTICLES (INDEPENDENT OF THIS PROJECT)

Aijaz, S., J. D. Kepert, H. Ye, Z. Huang and A. Hawksford, 2019: Bias correction of tropical cyclone parameters in the ECMWF ensemble prediction system in Australia. *Mon. Wea. Rev.*, **147**, 4261– 4285, doi:10.1175/MWR-D-18-0377.1.

Davidson, N. E., Y. Xiao, Y. Ma, H. E. Weber, X. Sun, L. J. Rikus, J. D. Kepert, P. J. Steinle, G. S. Dietachmayer, C. Lok, J. Fraser, J. Fernon and J. Shaik, 2014: ACCESS-TC: Vortex Specification, 4D-VAR Initialization, Verification and Structure Diagnostics. *Mon. Wea. Rev.*, **142**, 1265–1289, doi:10.1175/MWR-D-13-00062.1.

Kepert, J. D., J. Schwendike and H. Ramsay, 2016: Why is the tropical cyclone boundary layer not "well mixed"? *J. Atmos. Sci.,* **73**, 957–973.

Stern, D. P., J. D. Kepert, G. H. Bryan and J. D. Doyle, 2020: Understanding atypical midlevel wind speed maxima in hurricane eyewalls. *J. Atmos. Sci.*, **77**, 1531 – 1557, doi:10.1175/JAS-D-19-0191.1

Vaughan, A., K. J. E. Walsh and J. D. Kepert, 2020: The stationary banding complex and secondary eyewall formation in tropical cyclones. *J. Geophys. Res.*, **125**, e2019JD031515, doi:10.1029/2019JD031515.

Yu, C.-L., Didlake, A.C.,Jr., Kepert, J.D. and Zhang, F., 2020: Investigating axisymmetric and asymmetric signals of secondary eyewall formation using observations-based modeling of the tropical cyclone boundary layer. J. Geophys. Res., in revision.

Zhang, F., D. Tao, Y. Q. Sun and J. D. Kepert, 2017: Dynamics and predictability of secondary eyewall formation in sheared tropical cyclones. *J. Adv. Model. Earth Sys.*, **9**, 89–112, doi:10.1002/2016MS000729

Zieger, S., D. J. M. Greenslade and J. D. Kepert, 2018: Wave ensemble forecasts for tropical cyclones in the Australian region. *Ocean Dynamics*, 23pp, doi:10.1007/s10236-018-1145-9.

Zieger, S., Greenslade, D., Aijaz, S., Kepert, J.D. and Burton, A., 2020: Hindcasting of tropical cyclone winds and waves. *Ocean Engineering*, in revision.

Zieger, S., Kepert, J.D. Greenslade, D., and Aijaz, S., 2020: Assessment of tropical cyclone wave models for engineering applications. *Ocean Engineering*, submitted.

CONFERENCE PROCEEDINGS ARTICLES AND EXTENDED ABSTRACTS

W. Thurston, K.J. Tory, R.J.B. Fawcett and J.D. Kepert, 2016: Large-eddy simulations of pyro-convection and its sensitivity to moisture. Proceedings of the 5th International Fire Behaviour and Fuels conference, Melbourne, April 2016.

W. Thurston, K.J. Tory, R.J.B. Fawcett and J.D. Kepert, 2016: Long-range spotting by bushfire plumes: The effects of plume dynamics and turbulence on firebrand trajectory. Proceedings of the 5th International Fire Behaviour and Fuels conference, Melbourne, April 2016.

S.E. Ching, R.J.B. Fawcett, W. Thurston, K.J. Tory and J.D. Kepert, 2016: Mesoscale Features Related to the Blue Mountains Fires of 17 October 2013 Revealed by High Resolution Numerical Weather Prediction (NWP) Modelling. Proceedings of the 5th International Fire Behaviour and Fuels conference, Melbourne, April 2016.

Tory, K.J. and J.D. Kepert: Insights from the development of a pyrocumulonimbus prediction tool. Proceedings, Bushfire and Natural Hazards CRC & AFAC conference Perth, 5 – 8 September 2018, 277 – 288.

Kevin Tory, 2019: Pyrocumulonimbus Firepower Threshold: A pyrocumulonimbus prediction tool. *AFAC19 powered by INTERSCHUTZ Extended abstracts from the Bushfire and Natural Hazards CRC Research Forum (non-peer reviewed); Australian Journal of Emergency Management, Monograph No. 5, 2019.*

Dragana Zovko-Rajak, Kevin J. Tory and Jeffrey D. Kepert, 2019: A case study of South Australia's severe thunderstorm and tornado outbreak (28 September 2016). *AFAC19 powered by INTERSCHUTZ Extended abstracts from the Bushfire and Natural Hazards CRC Research Forum (non-peer reviewed); Australian Journal of Emergency Management, Monograph No. 5, 2019.*

TECHNICAL REPORTS

Kevin J. Tory and William Thurston, 2015: Pyrocumulonimbus: A literature review. CRC Report.

Will Thurston, Kevin Tory, Jeff Kepert, Robert Fawcett, 2015: Briefing note: Methodology for pyrocumulus large-eddy simulations – a rationale. BNHCRC Internal Report.

Tory, K.J., 2018: *Models of Buoyant Plume Rise*, BNHCRC, https://www.bnhcrc.com.au/publications/biblio/bnh-5267, 26pp.

HAZARD NOTES

Next generation fire modelling. BNHCRC Hazard Note 21. https://www.bnhcrc.com.au/hazardnotes/21.

Monitoring and Predicting Natural Hazards. BNHCRC Hazard Note 22. https://www.bnhcrc.com.au/hazardnotes/22.

Fire escalation by Downslope Winds. BNHCRC Hazard Note 24. https://www.bnhcrc.com.au/hazardnotes/24.

Two additional hazard notes are currently in the final stages of preparation, on the Tathra Fire and SA Tornado Outbreak case studies.

INVITED CONFERENCE PRESENTATIONS

J.D. Kepert, M. J. Naughton and J. Bally. Managing Severe Weather – Progress and Opportunities. Invited presentation at BNHCRC and AFAC Conference September 2014.

J. D. Kepert. The role of the boundary layer in tropical cyclone dynamics. Eighth International Workshop on Tropical Cyclones, Jeju, Sth Korea, December 2014.

J. D. Kepert. Why use ensemble prediction? Invited keynote presentation for BNHCRC and AFAC Conference September 2016.

Zovko-Rajak, D., Improving predictions of severe weather and saving lives. Invited presentation at the Science at the Shine Dome Annual Symposium – Predict, Respond, Recover: science and natural disasters, Canberra, 22 May 2018. Invited presentation.

Kevin Tory, Jeff Kepert, David Wilke and Zach Porter, 2019: The PyroCb Firepower Threshold: A pyrocumulonimbus prediction tool. Burea of Meteorology Research and Development Workshop, Melbourne, Australia, Nov 25 – 28. **Keynote presentation.**

Jeff Kepert, Kevin Tory, Mika Peace, Will Thurston, Harvey Ye, Jesse Greenslade, 2020: Fire-Atmosphere Interaction. Presented at France-Australia Bushfire Science Workshop, 15 September 2020

Jeff Kepert, Kevin Tory and Will Thurston, 2020: A fast, physically based scheme for predicting long-range ember transport in bushfire plumes. Presented at the Bureau of Meteorology Annual Research and Development Workshop, Melbourne, Australia, 23 November 2020.

Kevin Tory, Harvey Ye, and Gilbert Brunet , 2020: Tropical cyclone formation regions in CMIP5 models: A global performance assessment and projected changes. Presented at the Bureau of Meteorology Annual Research and Development Workshop, Melbourne, Australia, 23 November 2020.

OTHER CONFERENCE PRESENTATIONS

AFAC Conference 2014

The effect of fire-plume dynamics on the lateral and longitudinal spread of long-range spotting by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert. This talk presented some early results from the ember transport project.

Improved Predictions of Severe Weather to Help Reduce Community Impact by Jeff Kepert, Kevin Tory, Robert Fawcett and Will Thurston. This poster gave an overview of the project.

The Effects Of Fire-Plume Dynamics On The Lateral And Longitudinal Spread Of Long-Range Spotting by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert. This poster presented further information on the ember transport work.

Managing severe weather – progress and opportunities by Jeff Kepert, Mike Naughton and John Bally. This was a review paper focussing largely on ensemble prediction.

The heatwaves of the 2013 – 2014 Australian summer by Robert Fawcett and John Nairn. This presentation and paper discussed the heat wave events of he previous summer, and how the Bureau's pilot heatwave forecasting service performed.

Modelling the fire weather of the Coonabarabran fire of 13 January 2013 by Robert Fawcett, Claire Yeo, Will Thurston, Jeff Kepert and Kevin Tory. We began this work under the old Bushfire CRC as a side project, and completed it in the first months of this project. Conference presentation and paper.

Meteorological and fire behavioural lessons learned from the Aberfeldy Fire, Victoria, 17 January 2013 by Tim Wells, Claire Yeo and Robert Fawcett.

Fire Danger Indices: Current Limitations and a Pathway towards Better Indices Claire Yeo, Jeff Kepert and Robin Hicks. This poster reported on work which will inform the project aiming to build a new Fire Danger Rating system for Australia.

Mitigating the Effects of Severe Fires, Floods and Heatwaves Through The Improvements Of Land Dryness Measures and Forecasts Imtiaz Dharssi, Vinod Kumar, Claire Yeo, John Bally and Jeff Kepert.

An enhanced national Fire Danger Rating System for Australia – where to from here? Panel discussion led by Liam Fogarty, Simon Heemstra, Mike Rumsewicz and Jeff Kepert.

UK MetOffice Fire Symposium, December 2014

Will Thurston attended this symposium and presented the following papers.

Oral: The effect of fire-plume dynamics on the lateral and longitudinal spread of long-range spotting by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert.

Oral: Modelling the fire weather of the Coonabarabran (NSW) fire of 13 January 2013 by Robert Fawcett, Claire Yeo, Will Thurston, Jeff Kepert, and Kevin Tory.

Poster: Meteorological and fire behavioural lessons learned from the Aberfeldy Fire (Victoria) on 17 January 2013 by Robert Fawcett, Tim Weills, Claire Yeo and Will Thurston.

Poster: Fires and convection: a comparison of two extreme cases in southeast Australia by Andrew Dowdy and Will Thurston.

American Meteorological Society 11th Symposium on Fire and Forest Meteorology, May 2015

Oral: The effect of fire-plume dynamics on the lateral and longitudinal spread of long-range spotting by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert.

Oral: Modelling the fire weather of the Coonabarabran (NSW) fire of 13 January 2013 by Robert Fawcett, Claire Yeo, Will Thurston, Jeff Kepert, and Kevin Tory

Poster: Meteorological and fire behavioural lessons learned from the Aberfeldy Fire (Victoria) on 17 January 2013 by Robert Fawcett, Tim Wells, Claire Yeo and Will Thurston.

Unified Model Users GroupWorkshop, June 2015

Oral: Fire modelling at the Bureau of Meteorology by Jeff Kepert, Will Thurston, Kevin Tory and Robert Fawcett.

Australian Meteorological and Oceanographic Society Annual Conference, July 2015

Will Thurston attended the annual conference of the Australian Meteorological and Oceanographic Society (Brisbane) and presented the following papers.

Talk: Long-range spotting by bushfire plumes: The effects of in-plume turbulence on firebrand trajectory by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert.

Poster: Large-eddy simulations of pyro-convection and its sensitivity to environmental conditions by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert.

Poster: Modelling the fire weather of the Blue Mountains fires of October 2013 by Simon E Ching, Robert J B Fawcett, William Thurston, Kevin J Tory and Jeffrey D Kepert.

CRC and AFAC Conference, September 2015

We submitted a number of abstracts to the 2015 AFAC/CRC conference in Adelaide, and Jeff Kepert, Will Thurston and Simon Ching attended.

Oral: William Thurston, Kevin J. Tory, Robert J. B. Fawcett and Jeffrey D. Kepert, 2015: Large-eddy simulations of pyro-convection and its sensitivity to environmental conditions. Proceedings of the Research Forum, AFAC 2015 Conference

Poster: Long-range spotting by bushfire plumes: The effects of in-plume turbulence on firebrand trajectory, by Will Thurston, Kevin Tory, Robert Fawcett and Jeff Kepert.

CAWCR Annual Modelling Workshop, October 2015

Kepert, J.D., R.J.B. Fawcett, E. Ching, W. Thurston and K.J. Tory. High-resolution ensemble prediction of an east coast low. Paper presented at the Centre for Australian Weather and Climate Research, 19 – 22 October 2015.

Modelling the Fire Weather of the Blue Mountains Fires of 17 October 2013. Simon E. Ching, Robert J. B. Fawcett, William Thurston, Kevin J. Tory and Jeffrey D. Kepert. Paper presented at the Centre for Australian Weather and Climate Research, 19 – 22 October 2015.

Bushfire Workshop, Monash University, January 2016

The effect of turbulent plume dynamics on ember transport. William Thurston, Kevin J Tory, Robert J B Fawcett and Jeffrey D Kepert

The importance of fire moisture to PyroCumulus development. William Thurston, Kevin J Tory, Robert J B Fawcett and Jeffrey D Kepert

Australian Meteorological and Oceanographic Society Annual Conference, February 2016

High-resolution ensemble prediction of an East Coast Low. Jeff Kepert, Robert Fawcett, Simon Ching, William Thurston and Kevin Tory

Time and space scales in the tropical cyclone boundary layer. Jeffrey D. Kepert

Kepert, J.D., R.J.B. Fawcett, E. Ching, W. Thurston and K.J. Tory. High-resolution ensemble prediction of an east coast low. Presented at the AMOS Annual Conference, Melbourne, 8 – 11 February 2016.

5th International Fire Behaviour and Fuels Conference, April 2016

Large-eddy simulations of pyro-convection and its sensitivity to moisture, W. Thurston, K.J. Tory, R.J.B. Fawcett and J.D. Kepert

Long-range spotting by bushfire plumes: The effects of plume dynamics and turbulence on firebrand trajectory, W. Thurston, K.J. Tory, R.J.B. Fawcett and J.D. Kepert

Mesoscale Features Related to the Blue Mountains Fires of 17 October 2013 Revealed by High Resolution Numerical Weather Prediction (NWP) Modelling by S.E. Ching, R.J.B. Fawcett, W. Thurston, K.J. Tory and J.D. Kepert. Paper presented at the Fire Behaviour and Fuels conference, Melbourne, April 2016.

Bureau of Meteorology Research and Development Seminar Series

Why is the tropical cyclone boundary layer not well mixed? Jeffrey Kepert, 29th June 2016

The effect of turbulent plume dynamics on ember transport. William Thurston, 6th April 2016

Large-eddy simulations of pyro-convection and its sensitivity to environmental conditions. William Thurston, 4th November 2015

High resolution modelling of the meteorology of the Blue Mountain bushfires of October 2013. Simon Ching, 16th October 2015.

CRC and AFAC Conference, September 2017

Zovko-Rajak, D., J. D. Kepert, K.J. Tory and R.J.B. Fawcett, 2018: Ensemble prediction of the East Coast Low of April 2015. Poster presented at the CRC Research Showcase, Adelaide, July 2017 and AFAC Conference, Sydney, September 2017.

Zovko-Rajak, D., K. J. Tory, R. J. B. Fawcett, J. D. Kepert and L. J. Rikus, 2018: Dynamics and sources of the striated delta cloud observed during the east coast low of April 2015. Poster presented at the AMOS conference, Sydney, February 2018.

European Geophysical Union, April 2018

Tory, K. J., W. Thurston, J. D. Kepert: Thermodynamics of pyrocumulus formation. Presented at the EGU General Assembly, Vienna, Austria, April 2018.

Tory, K. J., W. Thurston, J. D. Kepert: The modest role of pyrogeneous moisture in pyrocumulus development. Presented at the EGU General Assembly, Vienna, Austria, April 2018.

33rd American Meteorological Society Conference on Hurricanes and Tropical Meteorology, April 2018

Kepert, J.D., 2018: The Boundary Layer Dynamics of Tropical Cyclone Rainbands. Paper presented at the 33rd American Meteorological Society Conference on Hurricanes and Tropical Meteorology, Ponte Vedra, FL.

S. Aijaz, J. D. Kepert, H. Ye, Z. Huang and A. Hawksford, 2018: Bias correction of tropical cyclone structure in global ECMWF-ensemble prediction system for NW Australia. Paper presented at the 33rd American Meteorological Society Conference on Hurricanes and Tropical Meteorology, Ponte Vedra, FL.

Kevin J. Tory, and Harvey Ye, 2018: Projected changes of tropical cyclone formation regions in the southern hemisphere. Paper presented by Jeff Kepert at the 33rd American Meteorological Society Conference on Hurricanes and Tropical Meteorology, Ponte Vedra, FL.

American Meteorological Society Symposium on Fire and Forest Meteorology, April 2018

Peace, M., J. D. Kepert and H. Ye, The ACCESS coupled fire-atmosphere model. Paper presented at the American Meteorological Society Symposium on Fire and Forest Meteorology.

Mika Peace, Jeffrey D Kepert, Lachlan McCaw, Neil Burrows, Bradley Santos, Robert Fawcett, 2018: Extreme fire behaviour at the Waroona fire, Western Australia January 2016. Paper presented at the American Meteorological Society Symposium on Fire and Forest Meteorology

Annual Meeting of the Asia-Oceania Geophysics Society, June 2018

Zovko-Rajak, D., K. J. Tory, R. J. B. Fawcett, and J. D. Kepert: High-resolution ensemble prediction of the Australian East Coast Low of April 2015. Presented at the AOGS 15th Annual Meeting, Honolulu, Hawaii, June 2018.

CRC and AFAC Conference, September 2018

Zovko-Rajak, D., K. J. Tory, J. D. Kepert, J. Fisher and M. Bass: A case study of South Australia's severe thunderstorm and tornado outbreak (28 September 2016). Poster presented at AFAC Conference.

CRC and AFAC Conference, September 2019

Jeff Kepert, Serena Schroeter, Kevin Tory and Will Thurston, 2019: A fast, physically based scheme for predicting long-range spotting potential. Presented at the AFAC-2019 Conference, August 27 – 30, Melbourne, Australia.

Kevin Tory and Jeff Kepert, 2019: Pyrocumulonimbus Firepower Threshold: A pyrocumulonimbus prediction tool. Presented at the AFAC-2019 Conference, August 27 – 30, Melbourne, Australia.

Dragana Zovko-Rajak, Kevin Tory and Jeff Kepert, 2019: A case study of South Australia's severe thunderstorm and tornado outbreak (28 September 2016). Presented at the AFAC-2019 Conference, August 27 – 30, Melbourne, Australia.

Annual Conference of the Australian Meteorological and Oceanographic Society, February 2020

Jeff Kepert, Serena Schroeter, Kevin Tory and Will Thurston, 2020: A fast, physically based scheme for predicting long-range ember transport in bushfire plumes. Presented at the 2020 Conference of the Australian Meteorological and Oceanographic Society, Perth, WA, February 10 – 13.

Jeff Kepert, 2020: Boundary layer dynamics of tropical cyclone rainbands. Presented at the 2020 Conference of the Australian Meteorological and Oceanographic Society, Perth, WA, February 10 – 13.

CRC and AFAC Conference, Indeterminate date 2020

Jesse Greenslade, Mika Peace, Jeffrey D Kepert, Harvey Ye, 2020: ACCESS-Fire: a case study. Poster presented at the 2020 CRC and AFAC Conference.

Jeff Kepert, Will Thurston, and Kevin Tory, 2020: Improved Predictions of Severe Weather: Ember Transport. Poster presented at the 2020 CRC and AFAC Conference.

Kevin Tory and Jeff Kepert, 2020: Real time trial of the Pyrocumulonimbus Firepower Threshold: A prediction tool for deep moist pyroconvection. Poster presented at the 2020 CRC and AFAC Conference.

David Wilke, Jeff Kepert, Kevin Tory, Paul Fox Hughes, 2020: Improved predictions of severe weather: The meteorology of the Tathra bushfire. Poster presented at the 2020 CRC and AFAC Conference.

Other

Kepert, J.D., S Aijaz, H. Ye, Z. Huang, D. Greenslade, S. Zieger and A. Hawksford, Bias correction of tropical cyclone intensity and structure in an ensemble prediction system. Presented at NOAA Hurricane Research Division, Miami, FL, 23 April 2018.

Jeff Kepert, Kevin Tory, Dragana Zovko Rajak, and Serena Schroeter, 2019: Improved predictions of severe weather to reduce community risk. Presented at CRC RAF, Perth, 30 July 2019.

David Wilke, 2020: Results of the Tathra bushfire study. Presentation to fire weather forecasters in the New South Wales State Forecasting Centre

TEAM MEMBERS

RESEARCH TEAM

Jeff Kepert

Jeff began his career with the Bureau as a forecaster, before moving into training and then research. As a researcher, Jeff has spent most of his career on tropical cyclones and bushfire meteorology, but also worked for a while on data assimilation. He is prominent nationally and internationally, having convened major conferences, edited journals and served on the Australian Standards Wind Loading committee.

Jeff particularly values working as a researcher for the Bureau, because it not only provides an environment to do high standard research on interesting weather, but also the opportunity and support to see that research used to improve our ability to manage severe weather, thereby saving lives, property and money.

Kevin Tory

Kevin is a member of the High Impact Weather Research team from the Bureau of Meteorology. He began his research career with the Bureau 20 years ago, and in that time has worked in a wide range of areas including: weather model verification, air quality modelling, tropical cyclone dynamics with a focus on their formation, tropical cyclone projections in a future warmer world, and bushfire meteorology. In this time, he has published regularly in prominent international journals, and has internationally recognised expertise in tropical cyclone formation and tropical cyclone detection.

Kevin enjoys working as a researcher in the Bureau, surrounded by experts in a wide range of weather-related fields, such as modelling, forecasting, observing and dynamical understanding. Having this highly valuable resource at hand is a major contributor to the quality of his research and other researchers in the organisation.

Eng (Simon) Ching

Eng Ching completed a MPhil in meteorology at the City University of Hong Kong before joining the Bureau of Meteorology as a forecaster. He is renowned within forecasting circles for his deep understanding of atmospheric dynamics, and for its application to forecasting problems. He also has very strong IT skills, and a remarkable ability to interrogate meteorological output data, especially from models, to discern what is happening within the atmosphere.

Robert J. B. Fawcett

Robert Fawcett has worked at the Bureau for more than twenty years, mostly in climate monitoring / analysis. He participated in the development of the Bureau's first national heatwave forecast system and contributed to two recent Bureau reports on changes in fire weather across Queensland and New South Wales. He

developed great expertise in running ACCESS at very high resolution and was one of the first in the Bureau to successfully achieve sub-kilometre runs, of the meteorology of Black Saturday. He used these skills to prepare the first high-resolution ensemble simulation in the Bureau, of the East Coast Low presented in this report.

Serena Schroeter

Serena has a PhD in Quantitative Marine Science from the University of Tasmania's Institute for Marine and Antarctic Studies and has examined climate interactions with Antarctic sea ice in global coupled climate models to determine the cause of discrepancies between observed and modelled sea ice trends, to aid in future model development.

William Thurston

Will is a senior scientist in the Atmospheric Dispersion & Air Quality group at the UK Met Office, where he develops and deploys forecast systems using the NAME atmospheric dispersion model, specifically for predicting the spread of windborne animal and plant pathogens. Prior to joining the Met Office in 2017, Will spent six years in the High Impact Weather team at the Bureau of Meteorology, using high-resolution simulations to conduct fire weather research. Specific interests included atmospheric boundary-layer--fire interaction, pyro-convective cloud development and ember transport. Will has a BSc (Hons.) in Physics from the University of Bristol and MRes and PhD degrees from the University of Leeds, specialising in boundary-layer processes in both the atmosphere and ocean.

David Wilke

David Wilke is currently on an assignment of duties as a researcher in the High Impact Weather research team from the Bureau of Meteorology. David completed a PhD in fluid dynamics at the University of Adelaide in 2016 with a project considering blood flow in the umbilical cord. Since 2015 he has been working as a meteorologist in the NSW regional forecasting centre.

David's primary focus is running very high-resolution simulations of the Tathra bushfire in order to better understand the event and characterise the severe fire weather that developed. He is also part of a Bureau and Geoscience Australia team developing a prototype workflow for an impact forecast. This is an endeavour to combine numerical weather prediction output and information regarding the built environment into a spatial representation of impacts to inform forecasters and stakeholders of the effects of severe weather.

Dragana Zovko Rajak

Dragana Zovko Rajak was born in Mostar, Bosnia and Herzegovina. She completed undergraduate studies in physics and geophysics at the University of Zagreb, Croatia, and her PhD at the School of Earth Sciences at the University of Melbourne in 2016. Her PhD thesis focused on examining possible mechanisms and sources of turbulence generation around thunderstorms using a three-dimensional numerical model. This work has important implications not only for

the aviation industry but also for fundamental understanding of turbulence and cloud–environment interactions. Since 2016, Dragana has been a part of the High Impact Weather team at the Bureau of Meteorology. Her research involves application of numerical weather prediction systems to high impact weather, as well as understanding the physical processes that contribute to severe weather events in Australia, including east coast lows and severe thunderstorms.

END USERS

End-user organisation	End-user representative	Extent of engagement (Describe type of engagement)
AFAC	John Bally	Closely engaged with all work, especially fire and ensembles.
Bureau of Meteorology	Mike Bergin	Closely engaged with all work, especially tropical cyclones. Retired 2017.
SA Powernet	Frank Crisci	Deeply interested in SA tornado outbreak study because of powerline impact
Bureau of Meteorology	Paul Fox-Hughes	Closely engaged with all aspects of the project.
Rural Fire Service, NSW (Bureau of Meteorology from August 2020)	Simon Heemstra	Closely engaged with fire weather aspects of the project.
NSW SES	Caroline Ortel	Closely engaged with east coast low part of the project. Retired 2017.
Department of Environment and Conservation, WA	Lachie McCaw	Closely engaged with fire weather aspects of the project.
Rural Fire Service, NSW	Laurence McCoy	Closely engaged with fire weather aspects of the project.

REFERENCES

Albini FA (1979) Spot fire distance from burning trees: a predictive model. USDA Forest Service, Intermountain Forest and Range Experiment Station, Technical Report INT-56. (Ogden, UT)

Albini FA (1981*a*) A model for the wind-blown flame from a line fire. *Combustion and Flame* **43**, 155–174. doi:10.1016/0010-2180(81) 90014-6

Albini FA (1981*b*) Spot fire distance from isolated sources: extensions of a predictive model. USDA Forest Service, Intermountain Forest and Range Experiment Station, Research Note INT-309. (Ogden, UT)

Albini FA (1983*a*) Potential spotting distance from wind-driven surface fires. USDA Forest Service, Intermountain Forest and Range Experiment Station. Research Paper INT-309. (Ogden, UT)

Albini FA (1983b) Transport of firebrands by line thermals. Combustion Science and Technology 32, 277–288. doi:10.1080/00102208308923662

Albini, Frank A., Martin E. Alexander and Miguel G. Cruz, 2012: A mathematical model for predicting the maximum potential spotting distance from a crown fire. *International Journal of Wildland Fire,* **21(5)**, 609-627. https://doi.org/10.1071/WF11020.

Baum HR, McCaffrey BJ (1989) Fire induced flow field – theory and experiment. In 'The Second International Symposium on Fire Safety Science', 13–17 June 1988, Tokyo, Japan. (Ed. T Wakamatsa) pp. 129–148. (International Association for Fire Safety Science: London, UK)

Bednarczyk, C. N. and B. C. Ancell, 2015: Ensemble Sensitivity Analysis Applied to a Southern Plains Convective Event. Mon. Wea. Rev., 143, 230-249, doi:10.1175/MWR-D-13-00321.1.

Berlad A.L. and S.-L. Lee, 1968: Long Range Spotting. *Combustion and Flame*, **??**, 172 – 173.

Bhutia, Sangay, Mary Ann Jenkins and Ruiyu Sun, 2010: Comparison of Firebrand Propagation Prediction by a Plume Model and a Coupled–Fire/Atmosphere Large–Eddy Simulator, 2010: *J. Adv. Model. Earth Syst.*, **2**, Art. #4, 15 pp., DOI:10.3894/JAMES.2010.2.4.

Blanchi R, Leonard JE (2005) Investigation of bushfire attack mechanisms resulting in house loss in the ACT bushfire 2003. Report for Bushfire CRC, CSIRO Manufacturing and Infrastructure Technology, Highett, VIC, Australia G.A. Davidson, Atmos. Environ. 20 (1986) 471–478.

Bowler, N. E., A. Arribas, K. R. Mylne, K. B. Robertson, and S. E. Beare, 2008: The MOGREPS short-range ensemble prediction system. Q. J. R. Meteorol. Soc., 134, 703-722, doi:10.1002/qj.234.

Brinkmann,W. A. R., 1974: Strong DownslopeWinds at Boulder, Colorado. Mon. Wea. Rev., 102, 592–602.

Brown, A. R., S. H. Derbyshire, and P. J. Mason, 1994: Large-eddy simulation of stable atmospheric boundary layers with a revised stochastic subgrid model. Quarterly Journal of the Royal Meteorological Society, 120, 1485–1512.

Ching, S.E., R.J.B. Fawcett, W. Thurston, K.J. Tory and J.D. Kepert, Mesoscale features related to the Blue Mountains Fires of 17 October 2013 revealed by high resolution numerical weather prediction modelling, submitted to *Monthly Weather Review*.

Cruz, M.G., A.L. Sullivan, J.S. Gould, N.C. Sims, A.J. Bannister, J.J. Hollis, and R.J. Hurley, 2012: Anatomy of a catastrophic wildfire: The Black Saturday Kilmore East fire in Victoria, Australia, *Forest Ecology and Management*, **284**, 269 – 295.

Cruz M.G., Gould J.S., Alexander M.E., Sullivan A.L., McCaw W.L., Matthews S., 2015: *A Guide to Rate of Fire Spread Models for Australian Vegetation*. CSIRO Land and Water Flagship, Canberra, ACT, and AFAC, Melbourne, Vic, 123pp.

Cunningham, P., S. L. Goodrisk, M.Y. Hussain and R.R. Linn, 2005: Coherent vortical structures in numerical simulations of buoyant plumes from wildland fires. *International Journal of Wildland Fire*, **14**, 61 – 75.

Cunningham, P. and M. J. Reeder, 2009: Severe convective storms initiated by intense wildfires: Numerical simulations of pyro-convection and pyro-tornadogenesis. *Geophys. Res. Lett.*, **36**, L12812, doi:10/1029/2009GLO39262.

Devenish, B. J., and J. M. Edwards, 2009: Large-eddy simulation of the plume 246 generated by the fire at the Buncefield oil depot in December 2005. Proceedings of the Royal Society A, 465, 397–419.

Devenish, B. J., G. G. Rooney, and D. J. Thomson, 2010: Large-eddy simulation of a buoyant plume in uniform and stably stratified environments. Journal of Fluid Mechanics, 652, 75–103.

Dowdy, A. J., G. A. Mills, and B. Timbal, 2013: Large-scale diagnostics of extratropical cyclogenesis in eastern Australia. Int. J. Climatol., 33, 2318-2327, doi:10.1002/joc.3599.

Durran, D. R., 1990: Mountain waves and downslope winds. Atmospheric Process over Complex Terrain., Meteor. Monogr., Vol. 23, Amer. Met. Soc., 59–83.

Durran, D. R., and J. B. Klemp, 1987: Another look at downslope winds. Part II: Nonlinear amplification beneath wave-overturning layers. J. Atmos. Sci., 44, 3402–3412.

Ellis, P. F. M., 2010: The effect of the aerodynamic behaviour of flakes of jarrah and karri bark on their potential as firebrands. *J. Roy. Soc. West. Aust.*, **93**, 21–27.

Engel, C. B., Lane, T. P., Reeder,M. J., and Rezney, M. (2012). The meteorology of Black Saturday. *Quarterly Journal of the Royal Meteorological Society*, **139**, 585–599.

Finkele, K., G. A. Mills, G. Beard, and D. A. Jones, 2006: National gridded drought factors and comparison of two soil moisture deficit formulations used in prediction of Forest Fire Danger Index in Australia. Australian Meteorological Magazine, 55, 183–197.

Fromm, M. D., A. Tupper, D. Rosenfeld, R. Servranckx and R. McRae, 2006: Violent pyro-convective storm devastates Australia's capital and pollutes the stratosphere. *Geophys. Res. Lett.*, **33**, L05815, doi:10.1029/2005FL025161.

Fromm, M. D.,R. H. D. McRae, J. J. Sharples and G. P. Kablick III, 2012: Pyrocumulonimbus pair in Wollemi and Blue Mountains National parks, 22 November 2006. *Aus. Met. Ocean J.*, **62**, 117—126.

Goens, D. W. and P. L. Andrews, 1998: Weather and fire behavior factors related to the 1990 Dude Fire near Payson, Arizona. In: Proceedings: 2nd symposium on fire and forest meteorology. Boston, MA: American Meteorological Society: 153-158.

Gray, M. E. B., J. Petch, S. H. Derbyshire, A. R. Brown, A. P. Lock, H. A. Swann, and P. R. A.Brown, 2001: Version 2.3 of the Met Office large eddy model: Part II. Scientific documentation. Turbulence and Diffusion Note 276, UK Met Office, Exeter, United Kingdom, 49 pp.

Griffiths, D., 1998: Improved Formulae for the McArthur Forest Fire Danger Meter. Meteorological Note 214, Bureau of Meteorology.

Hakim, G. J. and R. D. Torn, 2008: Ensemble Synoptic Analysis. Meteorol. Monogr., 33, 147-162, doi:10.1175/0065-9401-33.55.147.

Hall, James, Peter F. Ellis, Geoffrey J. Cary, Glenys Bishop and Andrew L. Sullivan 2015: Long-distance spotting potential of bark strips of a ribbon gum (Eucalyptus viminalis). *International Journal of Wildland Fire* **2015**, *24*, 1109–1117. http://dx.doi.org/10.1071/WF15031.

Hill, A. J., C. C. Weiss, B. C. Ancell, A. J. Hill, C. C. Weiss, and B. C. Ancell, 2016: Ensemble Sensitivity Analysis for Mesoscale Forecasts of Dryline Convection Initiation. Mon. Wea. Rev., 144, 4161-4182, doi:10.1175/MWR-D-15-0338.1.

Hilton, J., W. Swedosh, L. Hetherton, A. Sullivan and M. Prakash, 2019: *Spark User Guide 1.1.2*. CSIRO, Australia, 51pp.

Himoto K, Tanaka T (2005) Transport of disk-shaped firebrands in a turbulent boundary layer. In 'the Eighth International Symposium on Fire Safety Science', 18–23 September 2005, Beijing, China. (Eds DT Gottuk, BY Lattimer) pp. 433–444. (International Association for Fire Safety Science: Baltimore, MD)

Hodgson, A., 1967. Fire management in Eucalypt forest. In: Proceedings of 6th Annual Tall Timbers Fire Ecology Conference, March 6-7, 1967, Tallahassee, Florida. Charied by E.V. Komarek, Tall Timbers Research Station, Tallahassee, Florida.

Holland, G. J., A. H. Lynch, and L. M. Leslie, 1987: Australian East-Coast Cyclones. Part I: Synoptic Overview and Case Study. Mon. Weather Rev., 115, 3024-3036, doi:10.1175/1520-0493(1987)115<3024:AECCPI>2.0.CO;2.

Johnson, R. H., R. S. Schumacher, J. H. Ruppert Jr., D. T. Lindsey, J. E. Ruthford and L. Kriederman, 2014: The role of convective outflow in the Waldo Canyon fire. *Mon. Wea. Rev.*, **142**, 3061—3080.

Kepert, J. D., 2001: The dynamics of boundary layer jets within the tropical cyclone core. Part I: Linear theory. *J. Atmos. Sci.*, **58**, 2469–2484, doi:10.1175/1520-469(2001)058,2469:TDOBLJ.2.0.CO;2.

Kepert, J.D., 2013: How does the boundary layer contribute to eyewall replacement cycles in axisymmetric tropical cyclones? *J. Atmos. Sci.*, **70**, 2808–2830, doi:10.1175/JAS-D-13-046.1.

Kepert, J.D., 2017: Time and space scales in the tropical cyclone boundary layer, and the location of the eyewall updraft. Accepted for publication, *J. Atmos. Sci.*

Kepert, J. D., and R. J. B. Fawcett, 2013: Meteorological Aspects of the Margaret River Fires of November 2011. 20th International Congress on Modelling and Simulation, Adelaide, Australia, The Modelling and Simulation Society of Australia and New Zealand Inc., 180–186.

Kepert, J.D. and D.S. Nolan, 2014: Reply to ''Comments on 'How Does the Boundary Layer Contribute to Eyewall Replacement Cycles in Axisymmetric Tropical Cyclones?''' *J. Atmos. Sci.*, **71**, 4692-4704, DOI: 10.1175/JAS-D-14-0014.1

Kepert, J.D., and Y. Wang, 2001: The dynamics of boundary layer jets within the tropical cyclone core. Part II: Nonlinear enhancement. *J. Atmos. Sci.*, **58**, 2485–2501, doi:10.1175/1520- 469(2001)058,2485:TDOBLJ.2.0.CO;2.

Klemp, J. B., and D. R. Lilly, 1975: The dynamics of wave-induced downslope winds. J. Atmos. Sci., 32, 320–339.

Koo E (2006) Wildfire modeling: firebrands and contiguous fire spread. PhD thesis, University of California – Berkeley.

Koo, Eunmo , Rodman R. Linn, Patrick J. Pagni and Carleton B. Edminster, 2012: Modelling firebrand transport in wildfires using HIGRAD/FIRETEC. *International Journal of Wildland Fire* **2012**, *21*, 396–417. http://dx.doi.org/10.1071/WF09146

Koo E, Pagni PJ, Linn RR (2007) Using FIRETEC to describe firebrand behavior in wildfires. In 'The Tenth International Symposium of Fire and Materials', 29–31 January 2007, San Francisco, CA. (Interscience Communications: London, UK)

Koo, E., P. J. Pagni, D. R. Weise and J. P. Woycheese, 2010: Firebrands and spotting ignition in large-scale fires. *Int. J. Wild. Fire*, **19**, 818—843.

Lareau, N. P., and C. B. Clements, 2016: Environmental controls on pyrocumulus and pyrocumulonimbus initiation and development. Atmospheric Chemistry and Physics, 16, 4005–4022.

Lareau, Neil P. and Craig B. Clements 2017: The Mean and Turbulent Properties of a Wildfire Convective Plume. *J. Appl. Meteorol. Clim.,* **56,** 2289–2299. DOI: https://doi.org/10.1175/JAMC-D-16-0384.1

Lee SL, Hellman JM (1969) Study of firebrand trajectories in a turbulent swirling natural convection plume. *Combustion and Flame* **13**, 645–655. doi:10.1016/0010-2180(69)90072-8

Lee SL, Hellman JM (1970) Firebrand trajectory study using an empirical velocity-dependent burning law. *Combustion and Flame* **15**, 265–274. doi:10.1016/0010-2180(70)90006-4.

Leslie, L. M. and M. S. Speer, 1998: Short-range ensemble forecasting of explosive Australian east coast cyclogenesis. Wea. Forecast., 13, 822-832, doi:10.1175/1520-0434(1998)013<1205:COSREF>2.0.CO;2.

Locatelli, J. D., M. T. Stoelinga, P. V. Hobbs, and J. Johnson, 1998: Structure and evolution of an undular bore on the high plains and its effects on migrating birds. Bull. Amer. Meteor. Soc., 79,1043–1060.

Long, R. B., 1955: Some aspects of the flow of stratified fluids. III. Continuous density gradients. Tellus, 22, 471–479.

Luderer, G., J. Trentmann and M. O. Andreae, 2009: A new look at the role of fire released moisture on the dynamics of atmospheric pyro-convection. *Int. J. Wild. Fire*, **18**, 554—562.

Luke, R. H., and A. G.McArthur, 1986: Bushfires in Australia. Australian Government Publishing Service, Canberra.

Martin, Jonathan and Thomas Hillen, 2016: The Spotting Distribution of Wildfires. *Appl. Sci*., 6, 177; doi:10.3390/app6060177.

Mason, P. J., and D. J. Thomson, 1987: Large-eddy simulations of the neutral-static-stability planetary boundary layer. Quarterly Journal of the Royal Meteorological Society, 113, 413–443.

McArthur, A. G., 1966: *Weather and grassland fire behaviour*. Department of National Development, Forestry and Timber Bureau, Leaflet 100, 23 pp.

McArthur, A. G., 1967: Fire behaviour in eucalypt forests. Department of National Development Forestry and Timber Bureau, Canberra, Leaflet 107.

McArthur, A.G., 1969. The fire control problem and fire research in Australia. In: Proceedings of the 1966 Sixth world forestry congress, vol. 2. Spanish Minister of Agriculture, Madrid, pp. 1986–1991.

McRae, R. H. D., J. J. Sharples, S. R. Wilkes, and A. Walker, 2013: An Australian pyro-tornadogenesis event. Nat. Hazards, 65, 1801–1811.

Mills, G.M., 2005: On the subsynoptic scale meteorology of two extreme fire weather days during the eastern Australian fires of January 2003. Australian Meteorological Magazine, 54, 265–290.

Mills, G. M., 2008: Abrupt surface drying and fire weather. Part I: overview and case study of the South Australian fires of 11 January 2005. Australian Meteorological Magazine, 57, 299–309.

Mills, G. A., R. Webb, N. E. Davidson, J. Kepert, A. Seed, and A. D., 2010: The Pasha Bulker east coast low of 8 June 2007. Technical Report 023.

Muraszew A, Fedele JB (1977) Trajectory of firebrands in and out of fire whirls. Combustion and Flame 30, 321–324. doi:10.1016/0010-2180 (77)90081-5

New South Wales Rural Fire Service, 2014: Bush Fire Bulletin. Number 2/2014.

Noble, I. R., G. A. B. Barry, and G. A. M., 1980: McArthur's fire-danger meters expressed as equations. Australian Journal of Ecology, 5, 201–203.

Nolan. D.S., R. Atlas, K. T. Bhatia, and L. R. Bucci, 2013: Development and validation of a hurricane nature run using the joint OSSE nature run and the WRF model. *J. Adv. Model. Earth Syst.*, **5**, 382–405, doi:10.1002/jame.20031.

Peace, M., L. McCaw, B. Santos , J.D. Kepert, N. Burrows, and R.J.B. Fawcett, 2017: Meteorological drivers of extreme behaviour during the Waroona bushfire, Western Australia, January 2016. Submitted to *J. SH. Earth System Science.*

Page WG, Wagenbrenner NS, Butler BW, Blunck DL (2019) An analysis of spotting distances during the 2017 fire season in the Northern Rockies, USA. *Canadian Journal of Forest Research* **49**, 317–325. doi:10.1139/CJFR-2018-0094

Pendergrass, A. G., and H. E. Willoughby, 2009: Diabatically induced secondary flows in tropical cyclones. Part I: Quasisteady forcing. *Mon. Wea. Rev.*, **137**, 805–821.

Pereira, José C. F., José M. C. Pereira, André L. A. Leite, and Duarte M. S. Albuquerque, 2015: Calculation of Spotting Particles Maximum Distance in Idealised Forest Fire Scenarios, Journal of Combustion **2015**, Article ID 513576, 17 pages, http://dx.doi.org/10.1155/2015/513576.

Pepler, A. S. and C. S. Rakich, 2010: Extreme inflow events and synoptic forcing in Sydney catchments. IOP Conference Series: Earth and Environmental Science, 11, 012010.

Petch, J. C., and M. E. B. Gray, 2001: Sensitivity studies using a cloudresolving model simulation of the tropical west Pacific. Quarterly Journal of the Royal Meteorological Society, 127, 2287–2306.

Peterson, D. A., E. J. Hyer, J. R. Campbell, M. D. Fromm, J. W. H. C. F. 274 Butler, and M. A. Fenn, 2015: The 2013 Rim Fire: Implications for predicting extreme fire spread, pyroconvection, and smoke emissions. Bulletin of the American Meteorological Society, 96, 229–247.

Peterson, D. A., E. J. Hyer, J. R. Campbell, J. E. Solbrig, and M. D. Fromm, 2017: A conceptual model for development of intense pyrocumulonimbus in western North America. Mon. Weather Rev., 145, 2235–2255.

Potter, B. E., 2005: The role of released moisture in the atmospheric dynamics associated with wildland fires. *Int. J. Wild. Fire*, **14**, 77—84.

Puri, K., G. Dietachmayer, P. Steinle, M. Dix, L. Rikus, L. Logan, M. Naughton, C. Tingwell, Y. Xiao, V. Barras, I. Bermous, R. Bowen, L. Deschamps, C. Franklin, J. Fraser, T. Glowacki, B. Harris, J. Lee, T. Le, G. Roff, A. Sulaiman, H. Sims, X. Sun, Z. Sun, H. Zhu, M. Chattopadhyay, and C. Engel, 2013: Implementation of the initial ACCESS numerical weather prediction system. Aust. Meteorol. Oceanogr. J., 63, 265-284.

Rabenhorst, S., D. N. Whiteman, D. L. Zhang, and B. Demoz, 2014: A Case Study of Mid-Atlantic Nocturnal Boundary Layer Events during WAVES 2006. J. Appl. Meteor. Climatol., 53, 2627–2648.

Rawson RP, Billing PR, Duncan SF (1983) The 1982–83 forest fires in Victoria. *Australian Forestry* **46**, 163–172. doi:10.1080/00049158. 1983.10674395

Rosenfeld, D., M. D. Fromm, J. Trentmann, G. Luderer, M. O. Andreae and R. Servanckx, 2007: The Chisolm firestorm: observed microstructure, precipitation and lightning activity of a pyro-cumulonimbus. *Atmos. Chem. Phys.*, **7**, 645—659.

Rudlosky, S. D and H. E. Fuelberg, 2011: Seasonal, Regional, and Storm-Scale Variability of Cloud-to-Ground Lightning Characteristics in Florida. *Mon. Wea. Rev*., **139**, 1826–1843.

Sanders, F., 1986: Explosive Cyclogenesis in the West-Central North Atlantic Ocean, 1981-84. Part I: Composite Structure and Mean Behavior. Mon. Wea. Rev., 114, 1781-1794, doi:10.1175/1520-0493(1986)114<1781:ECITWC>2.0.CO;2.

Sardoy N, Consalvi JL, Porterie B, Fernandez-Pello AC (2007) Modeling transport and combustion of firebrands from burning trees. *Combustion and Flame* **150**, 151–169. doi:10.1016/J.COMBUSTFLAME.2007.04.008

Sardoy, N., J.L. Consalvi, A. Kaiss, A.C. Fernandez-Pello, B. Porteriea, 2008: Numerical study of ground-level distribution of firebrands generated by line fires. *Combustion and Flame* **154,** 478–488. doi:10.1016/j.combustflame.2008.05.006.

Scorer, R., 1949: Theory of waves in the lee of mountains. Quart. J. Roy. Meteor. Soc., 75, 41–56.

Shapiro, L. J. and H. E. Willoughby, 1982: The response of balanced hurricanes to local sources of heat and momentum. *J. Atmos. Sci*., **39**, 378–394.

Simpson, J. E., 1998: Gravity Currents: In the Environment and the Laboratory. Cambridge University Press, 262 pp.

Skamarock, W. C., J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, M. G. Duda, X.-Y. Huang, W. Wang, and J. G. Powers (2008), A description of the advanced research WRF version 3, NCAR Tech. Note, 4751STR, Boulder, Colorado, pp. 113.

Smith, R. B., 1985: On severe downslope winds. J. Atmos. Sci., 42, 2597–2603.

Sofiev, M.; Ermakova, T.; Vankevich, R. Evaluation of the smoke-injection height from wild-land fires using remote-sensing data. Atmos. Chem. Phys. 2012, 12, 1995–2006.

Speer, M. S., P. Wiles, and A. S. Pepler, 2009: Low pressure systems of the New South Wales coast and associated hazardous weather: establishment of a database. Aust. Meteorol. Oceanogr. J., 58, 29-39.

Storey, Michael A, Owen F Price, Ross A Bradstock and Jason J Sharples, 2020: Analysis of Variation in Distance, Number, and Distribution of Spotting in Southeast Australian Wildfires. *Fire*, **2020**, 3, 10; doi:10.3390/fire3020010.

Storey, Michael A., Owen F. Price, Jason J. Sharples and Ross A. Bradstock, 2020: Drivers of long-distance spotting during wildfires in south-eastern Australia. *International Journal of Wildland Fire* https://doi.org/10.1071/WF19124

Su, C.-H., Eizenberg, N., Steinle, P., Jakob, D., Fox-Hughes, P., White, C. J., Rennie, S., Franklin, C., Dharssi, I., and Zhu, H., 2019: BARRA v1.0: the Bureau of Meteorology Atmospheric high-resolution Regional Reanalysis for Australia, *Geosci. Model Dev*., **12**, 2049-2068, https://doi.org/10.5194/gmd-12-2049-2019.

Swann, H., 1998: Sensitivity to the representation of precipitating ice in CRM simulations of deep convection. Atmos. Res., 47–48, 415–435.

Tarifa CS, del Notario PP, Moreno FG (1965a) On flight paths and lifetimes of burning particles of wood. In 'Tenth Symposium on Combustion', 17–21 August 1964, Cambridge, UK. pp. 1021–1037. (The Combustion Institute: Pittsburgh, PA)

Tarifa CS, del Notario PP, Villa AR, Martinez L, PerezO (1965b) Open fires and transport of firebrands. Instituto Nacional de Tecnica Aeroespacial, Report no. Grants FG-SP-114. (Madrid, Spain)

Tarifa CS, del Notario PP, Moreno FG, Villa AR (1967) Transport and combustion of firebrands. Instituto Nacional de Tecnica Aeroespacial, Report no. Grants FG-SP-114, FG-SP-146. (Madrid, Spain)

Thurston, W., K. J. Tory, J. D. Kepert, and R. J. B. Fawcett, 2014: The effects of fire-plume dynamics on the lateral and longitudinal spread of long-range spotting. *Proceedings of the Research Forum at the Bushfire and Natural Hazards CRC & AFAC 2014 Conference*, M. Rumsewicz, Ed., Bushfire and Natural Hazards CRC, 85–94, ISBN: 978-0-9941696-3-15.

Thurston, W., K. J. Tory, R. J. B. Fawcett and J. D. Kepert, 2016: Large-eddy simulations of pyro-convection and its sensitivity to moisture, *5th International Fire Behaviour and Fuels Conference proceedings.* 6pp.

Thurston, W., J. D. Kepert, K. J. Tory, and R. J. B. Fawcett, 2017: The contribution of turbulent plume dynamics to long-range spotting. International Journal of Wildland Fire, 26, 317–330.

Toivanen, J., Engel, C. B., Reeder, M. J., Lane, T. P., Davies, L., Webster, S., et al. (2019). Coupled atmosphere-fire simulations of the Black Saturday Kilmore East wildfires with the Unified Model. *Journal of Advances in Modeling Earth Systems*, **11**. https://doi.org/10.1029/2017MS001245

Tolhurst K, Shields B, Chong D (2008) Phoenix: development and application of a bushfire risk management tool. *Australian Journal of Emergency Management* **23**, 47.

Tory, K. J., 2014: The turning winds with height thermal advection rainfall diagnostic: Why does it work in the tropics? Aust. Meteorol. Oceanogr. J., 64, 231-238.

Tory, K. J., and J. D. Kepert, 2020: Pyrocumulonimbus Firepower Threshold: A pyroCb prediction method. *Weather and Forecasting*, in press.

Tory, K. J., M. Peace and W. Thurston, 2016: Pyrocumulonimbus forecasting: Needs and issues. *Bushfire and Natural Hazards CRC Report, NO. 2016.239.*

Tory, K. J. and W. Thurston, 2015: Pyrocumulonimbus: A literature review. *Bushfire and Natural Hazards CRC Report*, NO. 2015.067.

Tory, K. J., W. Thurston, and J. D. Kepert, 2018: Thermodynamics of pyrocumulus: A conceptual study. *Mon. Wea. Rev*., **146**, 2579–2598, doi:10.1175/MWR-D-17-0377.1.

Trentmann, J., G. Luderer, T. Winterrath, M. D. Fromm, R. Servranckx, M. Herzog, H.-F. Graf and M. O. Andrea, 2006: Modeling of biomass smoke injection into the

lower stratosphere by a large forest fire (Part I): reference simulation. *Atmos. Chem. Phys.*, **6**, 5247—5260.

Vase, J., and Coauthors, 2013: The Australian Water Resource Assessment Modelling System (AWRA). 20th International Congress on Modelling and Simulation, Adelaide, Australia, The Modelling and Simulation Society of Australia and New Zealand Inc., 3015–3021.

Verdon-Kidd, D. C., A. S. Kiem, and G. R. Willgoose, 2016: East Coast Lows and the Pasha Bulker storm - lessons learned nine years on. J. South. Hemisph. Earth Syst. Sci., 66, 152-161, doi:10.22499/3.6002.005.

Vosper, S. B., and D. J. Parker, 2002: Some perspectives on wave clouds. Weather, 57, 3–8.

Wang, H.-H., 2011: Analysis on Downwind Distribution of Firebrands Sourced from a Wildland Fire. *Fire Technology*, **47**, 321–340, DOI: 10.1007/s10694-009-0134-4.

Wells, T., C. S. Yeo, and R. J. B. Fawcett, 2014: Meteorological and fire behavioural lessons learned from the Aberfeldy Fire, Victoria, 17 January 2013. 2014 Conference of the Australian Fire and Emergency Services Council and the Bushfire and Natural Hazards Cooperative Research Centre., Wellington, New Zealand.

Werth PA, Potter BE, Alexander ME, Clements CB, Cruz MG, Finney MA, Forthofer JM, Goodrick SL, Hoffman C, Jolly WM (2016) Synthesis of knowledge of extreme fire behavior. Volume 2 for fire behavior specialists, researchers, and meteorologists. US Department of Agriculture, Forest Service, Pacific Northwest Research Station. General Technical Report PNW-GTR-891. (Portland, OR)

White, P.V.M., 1973: *The Eye of the Storm.* Jonathan Cape, 608pp.

Whiteman, C. D., and J. G. Whiteman, 1974: A historical climatology of damaging downslope windstorms at Boulder, Colarado. NOAA Tech. Rep. ERL336-APCL35, NOAA, 62 pp.

Willoughby, H. E., J. A. Clos, and M. G. Shoreibah, 1982: Concentric eyewalls, secondary wind maxima, and the evolution of the hurricane vortex. *J. Atmos. Sci*., **39**, 395–411.

Woycheese JP, Pagni PJ, LiepmannD (1998) Brand propagation from largescale fires. In 'The Second International Conference of Fire Research and Engineering (ICFRE2)', 10–15 August 1997, Gaithersburg, MD. pp. 137–150. (Society of Fire Protection Engineers: Boston, MA)

Woycheese JP, Pagni PJ, Liepmann D (1999) Brand propagation from large-scale fires. Journal of Fire Protection Engineering 10, 32–44. doi:10.1177/104239159901000203