

ECMWF Newsletter

Number 145 – Autumn 2015

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

Predictability: Living with the butterfly effect

An all-scale, finite-volume module for the IFS

Reducing temperature errors at coastlines

GEOLOW boosts access to Earth observation data



CONTENTS

EDITORIAL

The quiet revolution 1

NEWS

- Forty years of improving global forecast skill 2
- Predicting this year's European heat wave 4
- ECMWF meets its users to discuss forecast uncertainty 6
- Trans-polar transport of Alaskan wildfire smoke in July 2015 8
- Ensemble of Data Assimilations applied to atmospheric composition 10
- Using the OpenIFS model to describe weather events in the Carpathian Basin 11
- ECMWF helps ESO astronomers peer deep into space 12
- Surface verification in the Arctic 14
- ECMWF assimilates data from two new microwave imagers 14
- Improved spread and accuracy in higher-resolution Ensemble of Data Assimilations 15
- A first look at the new ecFlow user interface 16

VIEWPOINT

Living with the butterfly effect: a seamless view of predictability 18

METEOROLOGY

- An all-scale, finite-volume module for the IFS 24
- Reducing surface temperature errors at coastlines 30
- GEOWOW project boosts access to Earth observation data 35

GENERAL

- ECMWF Calendar 2016 42
- Contact information 42
- ECMWF publications 43
- Index of newsletter articles 44

PUBLICATION POLICY

The *ECMWF Newsletter* is published quarterly. Its purpose is to make users of ECMWF products, collaborators with ECMWF and the wider meteorological community aware of new developments at ECMWF and the use that can be made of ECMWF products. Most articles are prepared by staff at ECMWF, but articles are also welcome from people working elsewhere, especially those from Member States and Co-operating States. The *ECMWF Newsletter* is not peer-reviewed.

Editor: Georg Lentze

Typesetting and Graphics: Anabel Bowen with the assistance of Simon Witter

Cover photo: nicholashan/iStock/Thinkstock

Any queries about the content or distribution of the *ECMWF Newsletter* should be sent to Georg.Lentze@ecmwf.int

Guidance about submitting an article is available at www.ecmwf.int/en/about/news-centre/media-resources

CONTACTING ECMWF

Shinfield Park, Reading, Berkshire RG2 9AX, UK

Fax: +44 118 986 9450

Telephone: National 0118 949 9000

International +44 118 949 9000

ECMWF website: www.ecmwf.int

ECMWF NEWSLETTER GOING DIGITAL

As part of our aim to reduce our carbon footprint, we would like to encourage our readers to access the digital version of the ECMWF Newsletter.

Newsletters are already available as PDF documents here: www.ecmwf.int/en/about/news-centre/media-resources, and soon they will be available in fully digital form.

If you would like to be notified when the ECMWF Newsletter is available online, please contact us at newsletter@ecmwf.int and ask to be added to our email alerts list.

Alan Thorpe looks back at his five years as ECMWF's Director-General in his last Newsletter editorial before he steps down in December this year.



© Copyright 2015

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, England

Literary and scientific copyright belong to ECMWF and are reserved in all countries. This publication is not to be reprinted or translated in whole or in part without the written permission of the Director-General. Appropriate non-commercial use will normally be granted under condition that reference is made to ECMWF.

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.

The quiet revolution

As I look back over the past five years, when I have been privileged to be involved with ECMWF, the progress in the science and practice of numerical weather prediction (NWP) has been both significant and exciting. The THORPEX decade came to an end in 2014 and the follow-on research programmes have taken shape, pointing to the way forward. The first World Weather Open Science Conference in 2014, leading to the publication of *Seamless Prediction of the Earth System: From Minutes to Months*, provided a major assessment of how far the science has come and where the next challenges lie. Of course ECMWF itself is in reflective mood, this being its 40th anniversary year. And finally, the Centre is developing its next ten-year strategy, paving the way for the future contributions this organisation will make to the meteorological and environmental information enterprise.

As presented in the *Nature* review article 'The quiet revolution of numerical weather prediction', progress in meteorology and in particular in numerical weather prediction has been stunning over the past 40 years, but it has resulted from the steady accumulation of knowledge that often has gone relatively unnoticed by the rest of the scientific community – hence it's been a quiet revolution! ECMWF has played a major role in it, but this revolution has above all been a worldwide scientific collaborative achievement. In the review article we summarise four main ingredients that make up the quiet revolution. These are: observations and the definition of the initial state (the diagnosis, as envisaged by Vilhelm Bjerknes); understanding how the weather works (the physical processes); the predictions, including an assessment of uncertainty providing the degree of confidence in the forecast (the prognosis); and finally the computational science and technology of solving the equations and handling the data flows.

A feature of scientific progress in this field is that future improvements will also come from advances in these four ingredients. Clearly we can't exclude the possibility that a future discovery may overturn our understanding and take us in an entirely different direction, but a lesson of the quiet revolution is that it is only by a balanced development of all four ingredients that further progress can be made.

The time range over which NWP provides advance warning of the detailed evolution of weather events has grown from about two days ahead to nearly seven days ahead over the past 40 or so years. These predictions are fundamentally based on an ensemble approach because this is the best way to provide confidence in a forecast by enabling the range of possible scenarios of what might happen to be explored. Accuracy and reliability are the touchstones of what it is to be able to predict future weather. ECMWF is particularly well known for its contributions to advances in data assimilation and in ensemble prediction. However, this shouldn't obscure the advances made also in understanding and describing physical processes and in the way supercomputing can be exploited to enable the predictions to be made.

Over the past five years, ECMWF has made strides forward in all these areas. The pay-off from these advances can take time to become fully apparent. The atmosphere has variability in its potential predictability and this can temporarily mask the impact of the scientific advances (Thorpe et al., *ECMWF Newsletter*

No. 137, 2013). However, it can also throw up fantastic examples of the power and potential of NWP, and there is no better example of this than ECMWF's forecasts of Superstorm Sandy in 2012 (Magnusson et al., *Monthly Weather Review*, 2014), or more recently the forecast of Hurricane Joaquin. Glimpses of what might become routinely possible, as exemplified by those forecasts, motivate ECMWF's future goals.

Of course it is inevitable that the skill of weather forecasts tends to be judged by extreme events. Storms and hurricanes provide spectacular examples. But there are many other high-impact events such as local flash floods, heat waves, droughts, El Niño and fluvial flooding that can occur over an incredibly wide range of space and time scales. An understanding of what we can hope to predict regarding these events is also at the frontier of today's knowledge, but we can anticipate lots of progress to come (see articles in this Newsletter by Magnusson et al. and Buizza et al.). Figuring out where the forecast horizon lies is an exciting challenge. The World Weather Research Programme's new initiatives in sub-seasonal to seasonal prediction, polar prediction and high-impact weather will all contribute to an international collaborative effort to address the challenge.

During my time at ECMWF, I have also been privileged to be part of the dawning of a new era of analysing and predicting a wide range of environmental factors. This is fully consistent with and complementary to the Earth system approach that is developing in weather prediction. It is exemplified by the Copernicus services that are now part of the activities at the Centre. These are giving ECMWF a significant role in the critically important societal grand challenges of air pollution and climate change as well as associated issues such as flooding and greenhouse gas concentrations.

Few revolutions in the modern world have been as profound as those enabled by computers and computation, whether they be the internet, mobile technology, or big data. NWP has been and continues to be in the vanguard of this revolution. Modelling the world around us has become key, whether it be the weather, the human brain, galaxies, the economy, traffic flows, explosions etc. All these modelling tasks have in common the need for huge models, computational power and the ability to handle huge data volumes. The birth of ECMWF's Scalability Programme over the past two years is an exciting development and is going to lead to major advances in computation, including improved energy-efficiency. This will help NWP advance even more quickly than it could otherwise, given the kind of computational technology we expect to see over the coming decades. European collaboration in this area is taking off not least because of the leadership provided by ECMWF.

As I look ahead, the prospects for further major scientific and technological advances are huge. Whilst of course it is the ambition of all scientists to keep advancing, weather science has another critical feature. This is that people and society increasingly depend upon better predictions to get out of harm's way, protect infrastructures and stimulate the economy. I am confident that ECMWF will continue to lead the way and I look forward to following the Centre's progress as a member of its alumni.

Alan Thorpe

Forty years of improving global forecast skill

**PATRICK LALOYAUX,
LENNART BENGTSOON**
(Director of ECMWF 1981–1990)

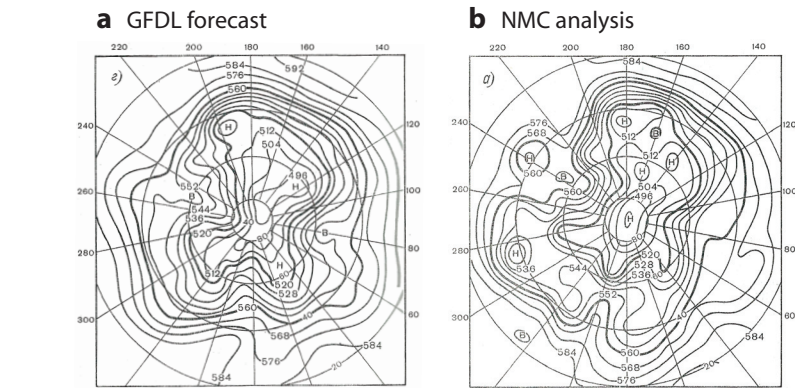
The Convention establishing ECMWF came into force on 1 November 1975. The Centre’s 40th anniversary provides an opportunity to look back at the first ECMWF weather forecast produced in 1975 for the period 1 March to 11 March 1965. To illustrate the impact of 40 years of research conducted at ECMWF, we have repeated this first forecast experiment using the CERA system, which is at the cutting edge of research in data assimilation and forecasting.

A series of medium-range forecasts between 20 February and 20 March 1965 were also produced using the CERA system. Their skill shows that ECMWF has fulfilled the initial promise that a future 6-day forecast would be as skilful as 2 to 3-day forecasts were in the 1970s.

ECMWF’s first forecast

The first global forecast experiments at ECMWF were carried out in the early autumn of 1975 using the general circulation models developed at GFDL (Geophysical Fluid Dynamics Laboratory) and UCLA (University of California, Los Angeles), kindly made available to ECMWF by the two institutes. The two models were used to produce a 10-day forecast starting on 1 March 1965 00 UTC.

The GFDL model was available with two different horizontal resolutions



500-hPa geopotential height charts showing (a) the GFDL forecast for day 5, valid for 6 March 00 UTC, and (b) the NMC analysis for 6 March 00 UTC, used for verification.

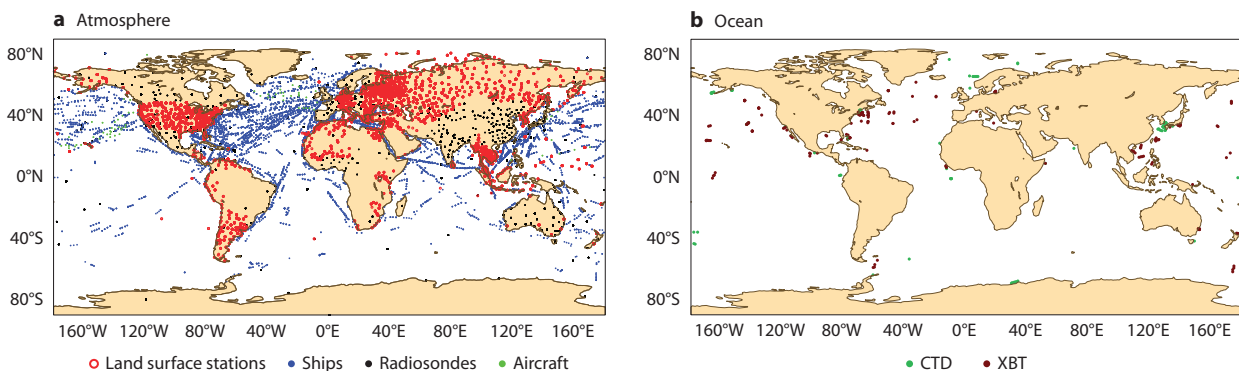
(250 and 500-km grids) and 9 vertical levels, 7 in the troposphere and 2 in the stratosphere. The UCLA model had a lower horizontal resolution corresponding to a grid spacing of 550 km with 6 vertical levels, 5 in the troposphere and 1 level in the stratosphere. Both models were mainly intended for climate simulation studies. They fully represented physical processes, including the radiation effects of H₂O, CO₂ and O₃. In this respect the models were more advanced than typical NWP models in the early 1970s.

At that time, atmospheric observations used in NWP consisted essentially of surface observations (weather stations and ship reports) and observations from radiosondes. The radiosonde network was very sparse in the tropics and the southern hemisphere, but coverage

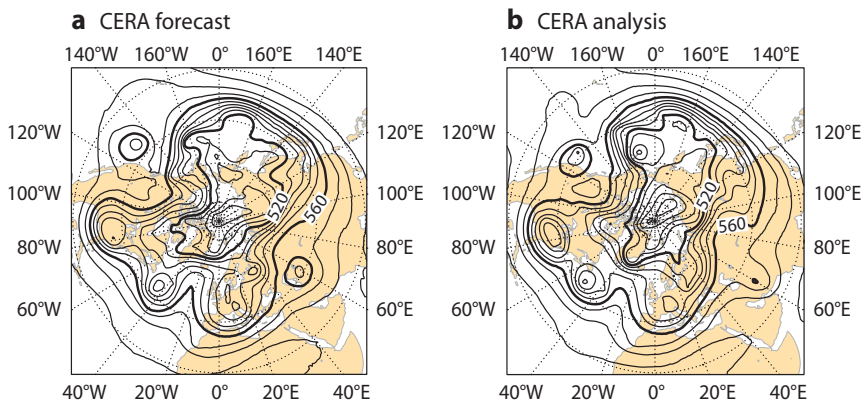
was rather good at middle and high latitudes in the northern hemisphere.

The available observations were processed by the NMC (National Meteorological Center) to compute the atmospheric analyses used to initialize and verify the forecast on 1 March 1965 00 UTC. It is likely that the NMC analysis was based on a Cressman scheme which used the operational forecast from NOAA (National Oceanic and Atmospheric Administration) and the Australian Met Service as the background, which was corrected using the observations.

Here we focus only on the results of the GFDL model-run at 250 km resolution, which was found to be the most accurate forecast. Although, as shown in the figures, the GFDL model indicates the large-scale weather patterns reasonably well in its 5-day



Observations assimilated by the CERA system for 1 March 1965 from (a) the atmosphere and (b) the ocean using Conductivity Temperature Depth (CTD) and Expendable BathyThermograph (XBT) measurements. A total of 126,703 atmospheric observations and 2,407 ocean observations were assimilated.



500-hPa geopotential height charts showing (a) the CERA forecast for day 5 (valid for 6 March 00 UTC) and (b) the CERA analysis for 6 March 00 UTC used for verification.

forecast, it fails to predict important synoptic features, such as cyclones. A more detailed assessment of the GFDL model shows that even the low-frequency part of the prediction had limited skill, suggesting that the forced Rossby waves had errors, presumably due to the incomplete representation of physical processes such as moist convection.

Today’s forecast

ECMWF has been developing a coupled atmosphere–ocean reanalysis system for the production of a 20th century reanalysis (*ECMWF Newsletter* No. 144, p. 15). To repeat the first ECMWF forecast experiment, we ran this system, called CERA, at the operational resolution of ECMWF’s Integrated Forecasting System (IFS), coupling a 16 km grid to the ocean NEMO model at a 1-degree resolution and with meridional refinement at the equator. We produced an analysis every six hours from 15 February to 1 April 1965, and 10-day forecasts for each day at 00 UTC.

The different types of atmospheric and ocean observations assimilated by the CERA system for 1 March 1965 come from the BUFR data used in the atmospheric reanalysis ERA-40 and the EN4 ocean dataset, respectively. It is thought that the same kind of atmospheric data was used by the NMC to provide the analysis for 1 March 1965 00 UTC. In addition to assimilating conventional observations, the CERA system constrains the air–sea interface with a nudging scheme towards the monthly sea-surface temperature analysis HADISST2. The figures show the CERA 500-hPa geopotential height forecast for day 5 and the CERA analysis for 6 March 1965 00 UTC used for verification. In the CERA forecast, the cyclone over the Atlantic is well predicted, as are other important synoptic features.

Improvement in forecast skill

In the 1970s, a series of 2-week forecasts were produced using the GFDL model for 12 winter cases


selected from the period 1964–69. Using the NMC analysis for verification and the anomaly correlation as a measure of skill, the 1.5-day forecast of 500 hPa geopotential height over the northern hemisphere has a skill score of 80%, and the 3.5-day one has a skill score of 60%. These two thresholds represent subjectively a good forecast and a useful prediction of the large-scale weather situation.

Twenty-nine medium-range forecasts produced by the CERA system for 00 UTC between 20 February and 20 March 1965 have been used to compute similar forecast skill scores. The anomaly correlation of the CERA 500-hPa geopotential height forecast has been calculated as a function of lead time for the northern hemisphere extratropics with respect to the CERA analysis. The 80% anomaly correlation skill level is reached at 6 days and the 60% level at 8.5 days. This means that the use of a modern ECMWF assimilation and forecasting system extends those skill levels by 4.5 and 5 days, respectively.

It is interesting to note that the general forecast skill obtained using the non-space observing system of the northern hemisphere extratropics in 1965 is not far removed from present forecast skill, which can draw on a number of space-based observing systems. However, the 1965 observing system is unlikely to have captured intense fine-scale weather systems, and outside the northern hemisphere extratropics it was not even close to the performance of present-day space-based observing systems.

When ECMWF was created, it was hoped that a future 6-day forecast would be as accurate as 2 to 3-day forecasts were in the early 1970s. The numerical experiments carried out with the CERA system show that the results of 40 years of research have allowed ECMWF to fulfil this commitment, even for the poorly observed period of 1965.

Despite the fact that here we have only examined a limited number of forecasts over a short period, the results obtained highlight the impressive progress in modelling and data assimilation at ECMWF from the beginning of operational predictions in 1979 to the present.



Anomaly correlation of the CERA 500-hPa geopotential height forecast for the northern hemisphere extratropics with respect to the CERA analysis, based on the 29 forecasts produced by the CERA system for 00 UTC from 20 February to 20 March 1965.

Predicting this year's European heat wave

**LINUS MAGNUSSON,
ALAN THORPE,
ROBERTO BUIZZA,
FLORENCE RABIER,
JEAN NICOLAU** (Météo-France)

A heat wave affected large parts of Europe during the summer of 2015. Records include an all-time high for Germany (5 July), an all-time high for Switzerland north of the Alps (7 July), a June record for Madrid (29 June) and a July record for the UK (1 July). The 2-metre temperature anomaly from ERA-Interim for the period 1 July to 15 August was more than 3 °C for large parts of central Europe as well as Spain. In this article we examine how successive ensemble forecasts for Paris picked up the onset of the heat wave increasingly well. On 1 July, the observed temperature at 12 UTC was between 35.5 °C and 36.8 °C among SYNOP stations in Paris, and later that day one station reached 39.7 °C, the second warmest temperature on record for the city. There are, of course, many other aspects of the heat wave worth discussing but not covered in this article.

Synoptic situation

The onset of the heat wave was associated with a ridge over western Europe, bringing hot tropical air from Africa. The amplification of the ridge started on the last days of June and coincided with the arrival of a Rossby wave train over western Europe. This Rossby wave had propagated from the west and could be traced back in the analysis to the Western Pacific, where it originated around 22 June.

Another synoptic-scale element was a cyclone that developed on 26 June over the eastern US and propagated eastward over the Atlantic. These dynamical ingredients probably contributed on different timescales to the establishment and predictability of the heat wave.

Ensemble visualisations

The ECMWF ensemble is designed so that its 51 members provide a set of scenarios consistent with our knowledge of the initial conditions and the laws of physics, bearing in mind the

associated uncertainties. The ensemble forecast is typically represented in the form of probability distributions, based on the relative frequency of ensemble members making particular predictions. This can take the form of a probability distribution function (PDF), shown for Paris in the top-left panel, or of a cumulative distribution function (CDF), shown in the top-right panel, or of box-and-whiskers diagrams, as shown in the mid-left and mid-right panels. We also show 2-metre temperature anomaly forecasts for Europe initialised on 18 June (bottom-left) and 22 June (bottom-right) for the week starting on 29 June to highlight the spatial pattern of the anomaly.

The different PDFs and CDFs for 2-metre temperature in Paris at 12 UTC on 1 July represent consecutive forecasts. Each PDF has been estimated by calculating a histogram of ensemble members using bins of 1 °C and applying a smoothing function with a weight of 0.5 for the central bin and 0.25 for the neighbouring bins. We have also included the distribution for the model climate, which is a useful reference as it includes the same systematic errors as the forecasts.

Results

The longest lead time illustrated here is the forecast issued one month before the event, on Monday 1 June. For this lead time – the so-called monthly forecast – it is expected that, to isolate a predictable signal, it is necessary to average in time over periods such as a week (as shown in the mid-right and lower panels). However, it is interesting to ask at which lead time the extended-range forecast for a specific time (such as 12 UTC on 1 July) acquires predictive skill. In answering this question, it is important to consider the range of scenarios predicted within the ensemble forecast.

At a lead time of a month (and certainly for point values) we expect to see only very weak evidence of predictive skill if any, and indeed there is no visible distinction between the climate and forecast PDFs and CDFs. In the forecast from Thursday 18 June, the ensemble PDF is shifted

towards warmer conditions, although the chance of near-record-breaking temperatures is still rather small. The shift in the PDF was due to a large-scale warm anomaly (covering most of Europe) for a longer period, and this anomaly put its stamp on the local PDF for Paris. The forecast from Monday 22 June has a peak in the PDF close to 34 °C, indicating that several members are picking up the risk of extreme temperatures, although a cold tail is still present in the distribution. At this stage forecasters in France started to pay attention to the risk of a heat wave. In the weekly temperature anomaly map for this forecast, the strongest anomalies are present over Western Europe. One could speculate that this signal is connected to the presence of the Rossby wave packet in the initial conditions.

Forecasts change substantially from around 26 and 27 June: the signal of extreme temperatures is dominating and the cold tail in the distribution has vanished, although the predicted temperatures are still a little cooler than the observed 36 °C. From 27 June, the forecast is sufficiently sharp for the coolest ensemble members to exceed the 90th percentile of the climatology. As early as 26 June, Météo-France was able to put warnings on its website for the extreme temperatures. Synoptically, at this time the cyclone had formed over the eastern US and its further development was probably less uncertain than in earlier forecasts. Finally, the 12-hour short-range forecast from 1 July 00 UTC has a very sharp distribution as little uncertainty remains for this short lead time.

To summarize the evolution, the forecast PDF peak progressively shifts towards higher temperatures and the distribution sharpens and loses its asymmetric cold tail. These are characteristics that one would hope to see in an ensemble forecast that has increasing levels of skill relative to climatology as the event approaches.

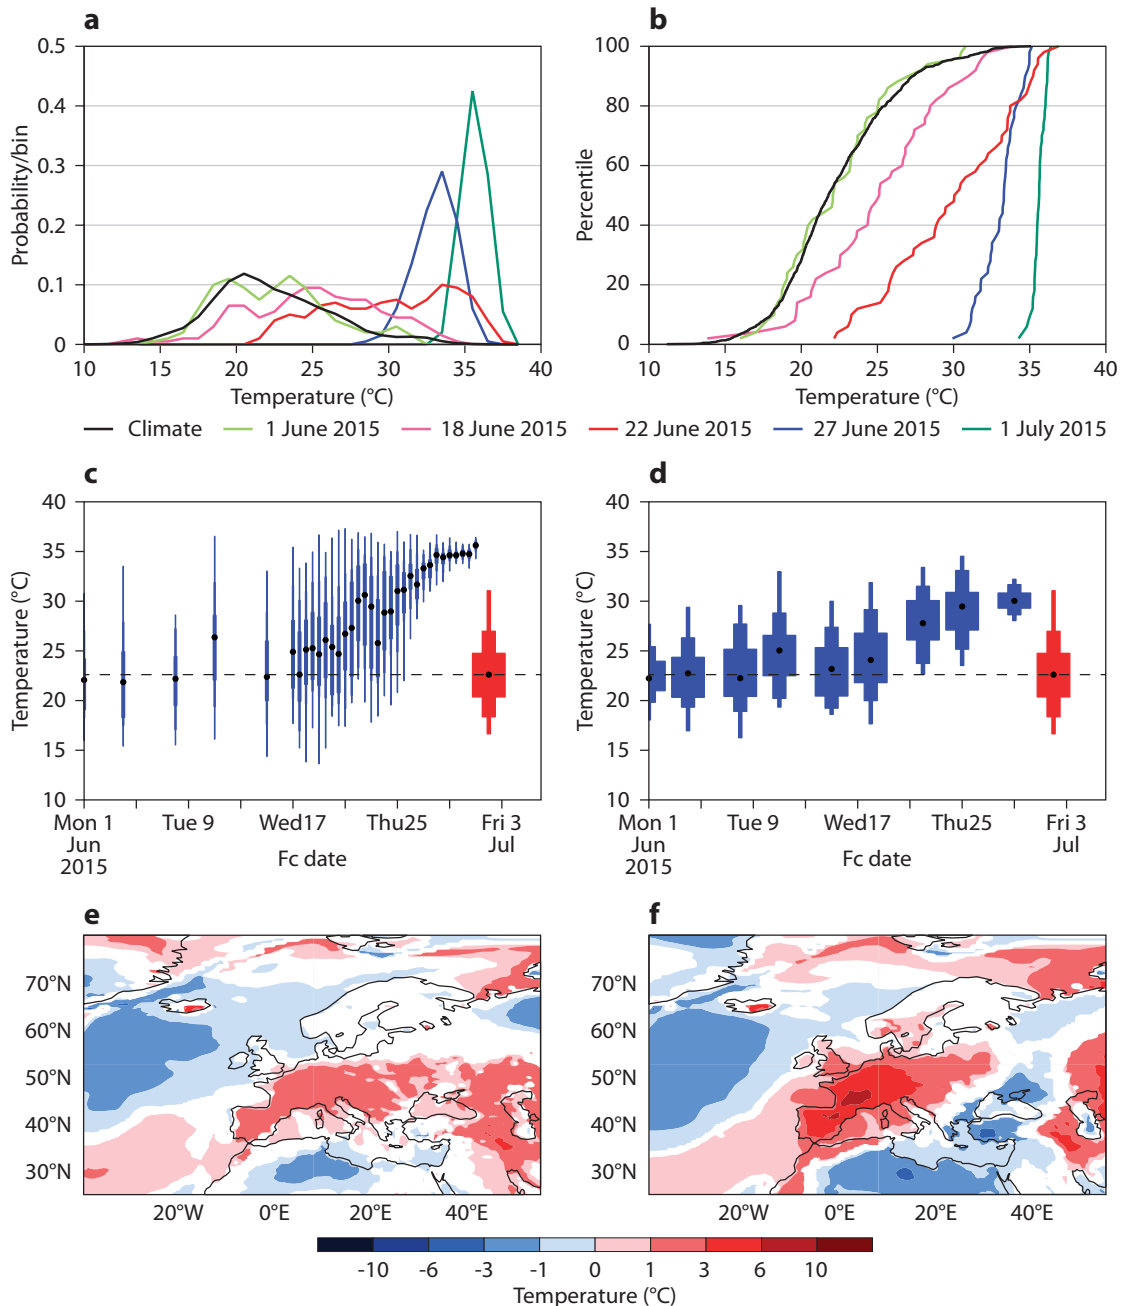
'Ready-set-go'

This case illustrates rather well how the concept of 'ready-set-go' could be used by forecasters to warn users. The ensemble distribution evolves from a

broad distribution indistinguishable from the climatology towards a very sharp distribution as the day of interest approaches. As valuable information is added to the initial conditions, the sharpness of the distribution increases, reducing uncertainty and increasing confidence. At the longer time range (2–3 weeks ahead), the

ensemble provides an indication that the distribution is shifted towards warmer values – the forecaster can tell users to be ready for the possibility of an extreme weather event. At the medium range (e.g. ten days ahead), the forecaster can tell users that they are now more confident about the forecast and can suggest that they

get set, i.e. start taking appropriate action that could help to manage the forthcoming weather conditions. Finally at short range (e.g. a few days ahead), forecasters can tell users that they are now very confident about the forecasts and can issue the warning: it is time to go and take action, since the extreme weather is coming.



Two-metre temperature forecasts. Ensemble forecasts valid 12 UTC on 1 July in Paris visualised by (a) probability density functions (PDFs) for forecasts initialised at 00 UTC on 1 June, 18 June, 22 June, 27 June and 1 July, (b) cumulative distribution functions (CDFs) for forecasts initialised at the same times and (c) a box-and-whisker plot for forecasts initialised at different times as shown, and for the model climate shown in red; also shown are (d) a box-and-whisker plot for ensemble forecasts of the average 2-metre temperature at 12 UTC for 29 June to 5 July in Paris, and for the model climate shown in red; and 2-metre temperature anomaly forecasts for 29 June to 5 July (shading in areas where the ensemble distribution is significantly different from climatology, significance level of 10%), initialised on (e) 18 June and (f) 22 June. The box-and-whisker symbols mark the 1st, 10th, 25th, 75th, 90th and 99th percentile and the median is marked with a dot. The dotted line represents the median of the model climate.

ECMWF meets its users to discuss forecast uncertainty

ANNA GHELLI,
FLORENCE RABIER



The annual meeting between ECMWF and its users, 'Using ECMWF's Forecasts' (UEF), took place in June 2015 and attracted 100 participants from Member and Co-operating States, national meteorological services, weather operators and academia. This year's theme, 'Quantifying and communicating uncertainty', explored the role of uncertainty in weather forecasting.

Information about forecast uncertainty has been around for a number of decades. Forecasters have long compared forecasts from different models, and they can now use ensemble forecasts to assess forecast uncertainty. More and more national meteorological services and weather forecast providers recognise that information on uncertainty is an integral part of forecasts and that it can be used to make better decisions based on weather forecasts. A discussion with users on quantifying, visualising and



UEF participants create a scatter plot. Participants were asked to indicate how confident they are in using and communicating information on uncertainty.

communicating uncertainty in weather forecasting thus seemed to be timely.

User engagement

While many users make use of uncertainty information, the extent to which they communicate it is still limited. This is illustrated by the scatter plot shown above, created by UEF participants, who tried to unravel the reasons why we are reluctant to communicate uncertainty information. *"Studies have shown that people understand that forecasts are uncertain and that they can make decisions taking into account this uncertainty,"* Dr Rebecca Morss from the US National Center for Atmospheric Research (NCAR) told the audience.

She also stressed that it is crucial to understand the different users' needs to be able to present the information in an accessible way. Therefore, engaging with end users to understand what kind of decisions they have to make and what time frames, risks and potential costs are associated with those decisions is paramount to providing appropriate support.

Ensembles of weather forecasts enable us to quantify uncertainty consistently with our knowledge of the initial conditions and model capability. It was generally agreed that this uncertainty can be translated and communicated to end users as the *confidence* they can have in a weather forecast.



Participants in the UEF 2015 meeting. About 100 people from around the world attended the event.

Uncertainty in practice

Sectors which have traditionally been reluctant to make use of uncertainty are now incorporating it into their risk maps. An example was presented at the meeting by the Swedish armed forces, who produce a risk map by combining the probability of a severe event with information about potential impacts.

In the UK, a number of changes were made after the 2007 floods. Warnings are now issued against lower probability thresholds to increase preparation lead times for emergency responders. A representative of the UK's Flood Forecasting Centre (a partnership between the Environment Agency and the Met Office) said four main challenges were the improvement of forecast accuracy; the need for a better and more widespread understanding of uncertainty; more training for emergency responders; and better integration with their action plans.

Visualising the large amount of data produced by ensemble forecasting systems is important to help users to translate quantified uncertainties into information on confidence to be provided to end users or impact models. Nicole Girardot from Météo-France showed a number of visualisations that are currently used or being tested at the French meteorological service.

Professor Tim Palmer from the University of Oxford asked whether we are wasting valuable resources by computing forecasts overly precisely and deterministically. According to him, it is very likely that we over-engineer our models by using high-precision, bit-reproducible computations throughout even though models use parametrizations which are partially stochastic. He believes that, *“taking account of the information content in NWP and climate model*

Key points

The meeting provided an opportunity to exchange ideas on uncertainty in forecasting and to define priorities. Here we present four key points made during the meeting that we think are crucial for the future:

- Ensemble forecasts enable us to quantify uncertainty consistently with our knowledge of the initial conditions and model capability. This uncertainty can be translated and communicated to end users as the confidence they can have in a weather forecast.
- Training is paramount to enabling forecasters and users in general to appreciate and exploit the information obtained from ensemble forecasts.
- The ensemble system needs to be reliable (the forecast probability needs to be equal to the observed frequency) and accurate (the distance between the observed probability density function (PDF) and the forecast PDF must be small). Knowing forecast model skill scores is essential to understand a model's strengths and weaknesses. Appropriate user-oriented verification methods need to be made available or developed.
- The outputs of ensemble models need to be presented in condensed form to avoid information overload, while ensuring that relevant weather event details are captured. The Extreme Forecast Index and the use of clusters are good examples, and regular exchanges between users and providers will result in the development of more products.

variables, we should be able to free up computing resources to develop higher-resolution models which will be better suited to future needs for increasingly localised weather and climate forecasts.”

Role of training

The programme not only included lectures and poster sessions but also group activities where participants could provide feedback and request additional products. One of the main concerns raised in group discussions was the large amount of data provided by current ensemble models and the lack of tools to summarise it. Moreover, training in understanding the strengths and weaknesses of models and in integrating the wealth of information they provide into current weather forecasting practices was highlighted as a priority, and this is an area where

ECMWF could provide support.

The meeting concluded with a panel discussion on the importance of quantifying uncertainty in NWP and how this information should be communicated to end users in terms of the confidence we can have in a weather forecast. The panellists highlighted the growing public appetite for weather to be put into context. *“We need a description of possible scenarios with associated probabilities of their occurrence so that we can understand what causes bifurcation in forecasts and how we can communicate it to the public,”* said Dan Suri from the UK Met Office. Training was high on the agenda of the participants, who suggested that more training for non-experts and forums for discussions with other users would be very helpful to improve the use of uncertainty information.

Feedback

We would like to thank ECMWF staff who helped during the meeting in various ways. We would also like to extend our thanks to all participants, who contributed actively to the activities and gave us constructive feedback. Here are some of their comments:

“I enjoyed all the breakout sessions and the conversation and dialogue that occurred in them.”

“Very useful meeting. Lots of ideas, I get the impression that many of us face the same challenges.”

“This was a fantastic workshop! We rely more and more on probabilistic information, it would be very useful to have more meetings like this one.”

“It was good to be able to discuss with ECMWF staff and people from the national meteorological offices.”


Trans-polar transport of Alaskan wildfire smoke in July 2015

MARK PARRINGTON,
ALESSIO BOZZO,
SAMUEL REMY (Laboratoire de
 Météorologie Dynamique),
IVAN TSONEVSKY

The Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF on behalf of the European Commission, provides a unique perspective on the transport and impact of atmospheric constituents associated with wildfire emissions, utilising near real-time information on wildfire locations and emissions of aerosols and trace gases. Five-day forecasts of smoke aerosol and trace gases, produced on a daily basis with the ECMWF Integrated Forecasting System (IFS), allow the track of plumes to be predicted to several days ahead.

A smoke plume from wildfires in Alaska was monitored in July 2015 as anomalous high pressure over the North Pole led to favourable circulation for trans-polar air mass transport. The plume was entrained into this flow on 6 July and passed over the eastern Arctic Ocean to Svalbard, where it was observed as haze at the Norwegian Polar Institute's Zeppelin Observatory on Spitsbergen on 10 July. Evidence of this plume was also observed in measurements of aerosol optical thickness (AOT) made by the Aerosol Robotic Network (AERONET) measurement site at Hornsund, also on Spitsbergen, and surface measurements of carbon monoxide made at the Zeppelin Observatory.

Wildfires are a significant component of the Earth system, emitting vast quantities of trace gases and aerosols and perturbing the chemical composition of the atmosphere. Boreal forests experience extensive wildfires every spring and summer. The 2015 wildfire season in boreal North America has been particularly strong, exacerbated by persistent dry conditions in the west. Approximately 5.1 million acres (an area larger than Slovenia) have burned in Alaska this year, making 2015 the second most

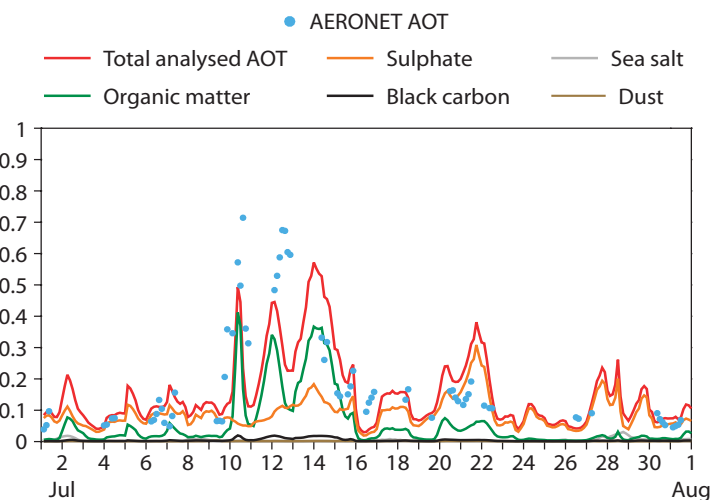


View from the Zeppelin Observatory, Spitsbergen. Photos taken (a) on 9 July 2015 and (b) on 10 July 2015 (courtesy of Per Erik Hanevold, Norwegian Polar Institute).

devastating Alaskan fire year on record after 2004, when 6.5 million acres burned. High-latitude wildfires are of particular concern for climate change in the Arctic, where temperatures are increasing at twice the rate of the rest of the planet. However, the impact of wildfire emissions on the Arctic climate is quite complex. For example, increased greenhouse gas emissions and deposition of black carbon aerosol on snow and ice can lead to increased temperatures and snow melt, while enhanced aerosol loading in the atmosphere can absorb

and reflect light and lead to surface cooling during the day.

In addition to forecasts and analyses, the application of the IFS to the Earth system provides an ideal test-bed for quantifying the impact of this plume on the radiation field over the Arctic region. Letting prognostic aerosols interact with radiation computations allows us to assess their radiative impact. The change in long-wave radiation is significant only very close to the source, where the aerosol load is large. As shown in the figures, the extinction in the short-

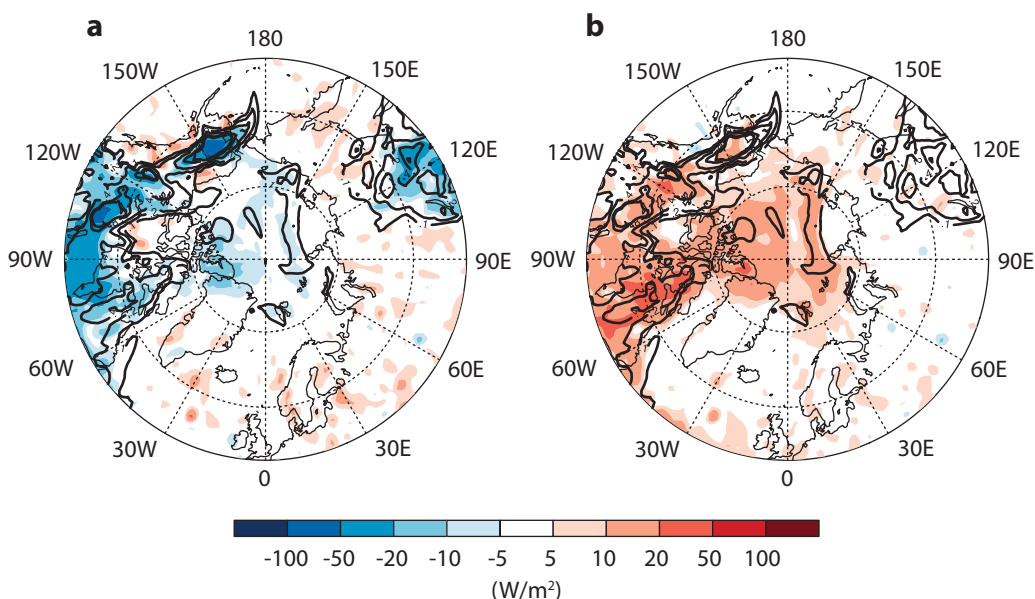


Observations and analysis of aerosol optical thickness (AOT). July 2015 AERONET measurements of AOT at 550 nm made at Hornsund, Spitsbergen, and analysed AOT from the IFS for different aerosol sources. The peaks in organic matter AOT between 10 and 16 July show the effects of a smoke plume from Alaskan wildfires, while sulphate AOT values are indicative of fossil fuel emissions (chart courtesy of CAMS website <http://atmosphere.copernicus.eu>).

wave is significant, with an average decrease in the surface solar radiation of 5–10 W/m² along the plume as it was transported across the Arctic Ocean and up to 50 W/m² in the vicinity of the fires over northern Alaska. The plume also reduces the overall albedo of the region, with about 10–20 W/m² less radiation reflected at

the top of the atmosphere and increased absorption within the aerosol layer. These radiative perturbations can lead to a cooling of the surface and a heating of the aerosol layer with potential impacts on weather patterns. This example of anomalous Alaskan wildfires and polar meteorology provided an excellent opportunity

to use the IFS in investigating the interaction between Arctic atmospheric composition, transport and radiation. Further investigation is under way to assess the overall radiative impact and possible meteorological feedbacks of events such as this, which will be monitored by CAMS in future boreal wildfire seasons.



The impact of using prognostic aerosol on radiation forecasts. The shading in the charts shows the difference between a 24-hour forecast with prognostic aerosol in the radiation scheme and another using a monthly mean aerosol distribution for (a) surface downward short-wave radiation and (b) top-of-atmosphere net short-wave radiation, averaged from 8 to 13 July 2015. Contours show organic matter AOT at 550 nm for values of 0.1, 0.2, 0.4, 0.8 and 1.0. Net radiation is defined as down- minus up-welling fluxes.

Forecast verification metric challenge

The public, industry, emergency managers and other decision-makers can use weather, climate and impact forecasts more effectively in their decision-making when the quality of the forecasts is measured in terms that are meaningful to them. Yet very few metrics exist to measure forecast quality in user-relevant terms.

To encourage the development of user-oriented verification approaches, the World Meteorological Organization’s Joint Working Group on Forecast Verification Research (JWGFVR) is conducting a challenge to develop and demonstrate new user-oriented forecast verification metrics. The new metrics will support the World Weather Research Programme (WWRP) High Impact Weather, Sub-seasonal to Seasonal, and Polar Prediction Projects.

The JWGFVR warmly encourages all interested researchers and practitioners to participate.

The deadline for entries is 31 October 2016.

More information and an entry form is available at: <http://www.wmo.int/pages/prog/arep/wwrp/new/FcstVerChallenge.html>



Ensemble of Data Assimilations applied to atmospheric composition

SÉBASTIEN MASSART

The EU-funded Copernicus Atmosphere Monitoring Service (CAMS) produces daily analyses of atmospheric species: aerosols, reactive gases and greenhouse gases. These analyses result from the assimilation of data from various satellites. The assimilation process balances the weight given to the satellite data and that given to the background (short-term forecast). The background weight is determined by the background error statistics. The background error statistics are currently static for atmospheric species. It has been demonstrated that using flow-dependent background errors is beneficial for ECMWF's operational meteorological analysis. A preliminary study is currently being carried out to investigate whether using flow-dependent background errors for atmospheric species would also be beneficial for the CAMS analyses. As a first step in this investigation, the temporal and spatial variability of the background error standard deviation

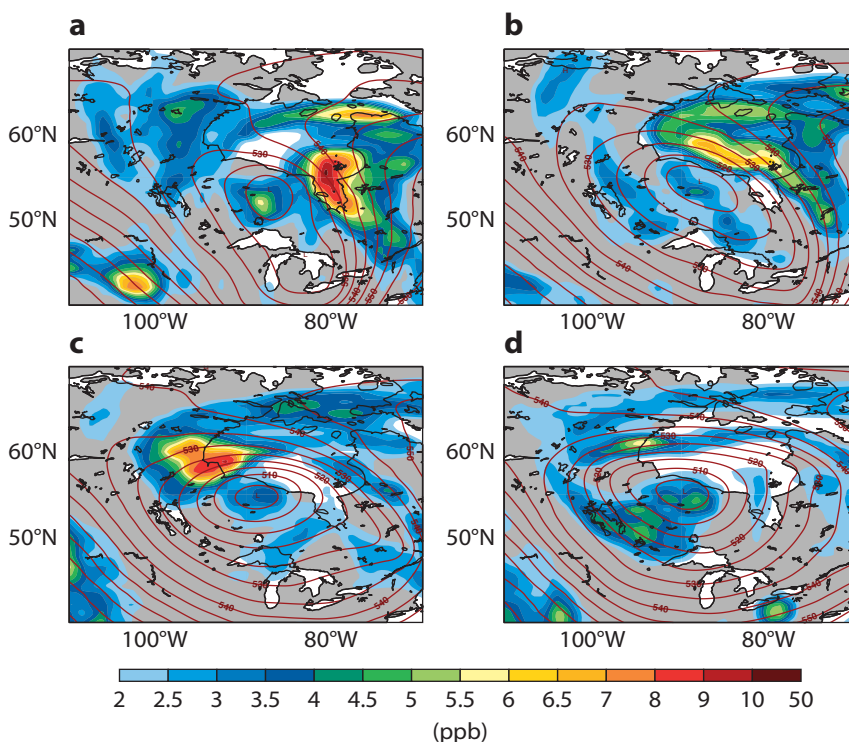
has been studied for two species of the greenhouse gas family: carbon dioxide (CO₂) and methane (CH₄).

To estimate the CO₂ and CH₄ background error statistics, we followed the approach used for the operational meteorological analysis based on an Ensemble of Data Assimilations (EDA). We ran a 25-member ensemble where each member used perturbed observations, perturbed sea-surface temperature fields and perturbed model physical tendencies. Moreover, we introduced perturbed surface fluxes to better account for the specific error sources of the studied species.

From the 25 members we computed the CO₂ and CH₄ ensemble spread, which is an estimate of the background error standard deviation. We found that the spread has large spatial variations and a strong diurnal cycle close to the surface for both CO₂ and CH₄. In the mid-troposphere, the spread is influenced by the atmospheric flow. The figure shows an example of changes in the

methane ensemble spread at 500 hPa for four consecutive analyses. On 4 October 2014 at 18 UTC a low pressure system is present north of the Great Lakes of North America. Meanwhile, there is a structure of high values of the methane ensemble spread to the east of the low pressure system (panel a). This comes from a large ensemble spread in the methane concentration at the surface that is uplifted in the low pressure system. As the low pressure system deepens over the following hours, the methane ensemble-spread structure is advected by the flow, moving around the low pressure system (panels b to d).

This particular situation highlights the variability in the ensemble spread for atmospheric species and therefore the variability in the background error standard deviation. It also illustrates the ensemble spread's dependence on the meteorological situation. The next step will be to assess the impact of using EDA-based flow-dependent background error statistics for the analysis of atmospheric species.



Methane ensemble spread for four consecutive analyses. Evolution of the methane 25-member EDA spread (shading) in parts per billion (ppb) and geopotential height (contours) at 500 hPa in m²/s² from (a) 4 October 2014 at 18 UTC to (b) 5 October 2014 06 UTC, (c) 5 October 2014 18 UTC and (d) 6 October 2014 at 06 UTC.

Using the OpenIFS model to describe weather events in the Carpathian Basin

**KAROLINA SZABÓNÉ ANDRÉ (ELTE),
GABRIELLA SZÉPSZÓ (OMSZ)**

The Hungarian Meteorological Service (OMSZ) started to use ECMWF's OpenIFS global model in 2013 for educational purposes. In cooperation with Eötvös Loránd University (ELTE) we examined a 3-week high-pressure weather situation with a persistent inversion, which is usually difficult to predict in the Carpathian Basin. The results show that the OpenIFS correctly reproduces pressure, temperature and humidity conditions of the period, but that it has some phase delay in the calculated inversion tops.

Until 2013, OMSZ and ELTE used limited-area models (CHAPEAU, WRF) to teach numerical modelling. The adoption of the OpenIFS global model brought challenges as well as opportunities. It was first applied to study a persistent inversion situation in the framework of a Master's thesis. Model runs and the supervision of numerical modelling aspects were provided by OMSZ, while synoptic evaluation was overseen by the University.

Cold air pools

The weather situation we considered is called a persistent cold air pool. It consists of a cold air layer trapped by the topography and a warmer air mass above. It can be a diurnal feature or persistent. The cold air is typically more humid than the air aloft and cools down to its dew point temperature, leading to the formation of fog.

Persistent cold air pools may cause

hazards for the population. Air pollutants accumulate below the inversion leading to health issues. Fog, smog and cold weather can also disrupt transport.

However, numerical weather prediction models can struggle to resolve and accurately predict the formation and decay of this important weather feature.

Case study

We examined a 3-week period from 13 November to 4 December 2011. On 13 November, the weather over Eastern Europe was marked by an anticyclone and a cold air pool formed in the Carpathian Basin. Ten days later, a warm front partly broke up the cold air pool. Following the influence of a cold front (on 28 November), the cold air pool formed again and it finally broke up on 3 December.

We performed a series of 24-hour forecasts covering the 3-week period on OMSZ's SGI supercomputer using 64 CPUs and 128 GB RAM. Due to capacity constraints we used a relatively low spatial resolution (T639, corresponding to a grid spacing of about 32 km, and 91 vertical levels). A 24-hour run took about 2.5 hours, so all the integrations took 2 days and 7 hours.

Investigated parameters

Cold air pools can be characterized by different parameters. We examined the mean sea level pressure (MSLP), the planetary boundary layer's convective potential energy (referred to here as shallow convective potential energy, SCP), and the height of the top and the

bottom of the inversion layer and the cloud layer.

SCP indicates the specific internal energy (J/kg) needed for a neutral lapse rate. SCP is calculated by modifying the CAPE (Convective Available Potential Energy) equation. We chose 850 hPa as the upper level because cold air pools are usually below this level.


Results

To validate the OpenIFS forecasts, we compared its hourly outputs to daily upper air and hourly surface observations at Szeged (46.25°N, 20.10°E, altitude: 83 m; nearest grid point in the OpenIFS: 46.23°N, 19.97°E, altitude: 100 m).

The figure shows the time series of MSLP and SCP below 850 hPa. Disruptive weather events (such as the warm front on 23 November and the cold fronts on 28 November and 3 December) can be identified. The forecasts gave relatively accurate MSLP and SCP values. Based on observations, SCP was negative during the period we considered, so the lower atmosphere was stable. Its three minimums (under 400 J/kg) indicate the three cold air pool events. Compared to observations, the OpenIFS forecasts show a less stable atmosphere (a higher SCP value) between 20 and 24 November, during the strong cold air pool period. The same is true at the beginning of December. However, the model values were below 400 J/kg, similar to the observations.

Comparing the calculated inversion tops, a phase delay was noticed between the observations and the model data. The model built up and decayed the inversion associated with the cold air pool earlier than the observed times. Investigating the calculated cloud top, it is apparent that the OpenIFS forecasts were accurate when clouds (relative humidity exceeding 90%) and the cold air pool were in place. Before the first cold air pool period, the model did not show any cloud. However, we could not identify observed low-level clouds (below 1,000 hPa) and fog in the model output.

For more information on the OpenIFS, visit: <http://software.ecmwf.int/oifs/>



Time series of observations and OpenIFS forecasts at Szeged, Hungary, between 13 November and 4 December 2011, showing (a) mean sea level pressure and (b) SCP below 850 hPa.

ECMWF helps ESO astronomers peer deep into space

MARC SARAZIN (ESO)

ECMWF weather forecasts help one of the world's top astronomical organizations, the European Southern Observatory (ESO), to fine-tune its telescopes in the quest to explore the outer reaches of the universe.

ESO's Paranal Observatory

These are exciting times for astronomy as modern technology enables us to study objects at the edge of the observable universe and to detect planets orbiting other stars. To do this, ESO operates three world-class observing sites in the Atacama Desert in northern Chile: La Silla, Paranal and Chajnantor. In this remote region, the sky is clear more than 300 nights a year and light pollution is minimal. And now ESO is building the 39-metre European Extremely Large Telescope (E-ELT), which will become 'the world's biggest eye on the sky'.

The Very Large Telescope (VLT) at Mount Paranal is ESO's current premier observatory for observations in visible and infrared light. It is made up of four telescopes, each 8.2 metres in diameter, which operate individually using a diverse range of instruments. Together with four smaller Auxiliary Telescopes, they also function as the light-collecting elements of the site's Very Large Telescope Interferometer (VLTI).

Weather conditions are generally favourable for the operation of the VLT. The cold Humboldt Current in the ocean 14 km to the west maintains a temperature inversion well below the observatory, which ensures cloudless skies most of the year.



The Paranal Observatory. Weather conditions tend to be favourable for the operation of the observatory's telescopes, which are located at 2,635 metres above sea level. *Photo: ESO/G. Hüddepohl (atacamaphoto.com)*

But the observatory's optimum performance depends on more than just clear skies. Temperature and wind also have a bearing on the correct functioning of the facility.

Temperature forecasts

Astronomical objects can appear blurred as a result of variations in the refractive index of air, which are in turn caused by variations in temperature. It is therefore important to minimize differences between air temperature inside the telescope enclosure and outside it.

To achieve this, the VLT is actively thermally controlled during daytime. Air conditioning brings the inner volume of the enclosure to the target temperature estimated from ECMWF

forecasts for the time at which the enclosure is due to open.

The forecasts are extracted from ECMWF data by applying a bilinear interpolation from the four nearest grid points and a linear interpolation from the two nearest pressure levels. A Kalman filter is then applied, trained on the past two weeks of local temperature measurements taken at 30 metres above ground (the height of the telescope enclosure). A Kalman filter is an algorithm designed to estimate a quantity based on predictions of the quantity on the one hand and observations on the other.

Wind profiles


Turbulence in the Earth's atmosphere distorts images obtained even at




The Very Large Telescope system shortly after sunset. During daytime, the huge telescope structures are cooled to the air temperature forecast by ECMWF. The people who can be seen on the left provide a sense of scale. *Photo: ESO/B. Tafreshi (twanight.org)*

Improvements brought by adaptive optics.

The left-hand image of two interacting galaxies is blurred as a result of atmospheric seeing. The right-hand image illustrates the improvement obtained by using adaptive optics. Image: ESO



Vertical wind profile. ECMWF's forecast of wind speed at different pressure levels above the Paranal observatory (743 hPa) for three days starting on 24 September 2015. The information is used to predict times of slow-moving turbulence, which is easier to correct for by adaptive optics techniques.



the world's best sites for astronomy, including at Paranal. To minimize this 'atmospheric seeing' effect, astronomers have turned to a method called adaptive optics.

Sophisticated, deformable mirrors controlled by computers can make adjustments hundreds or even thousands of times a second to correct for the distortion caused by turbulence in the Earth's atmosphere. This makes the images obtained almost as sharp as those taken from space, and it allows the optical system to observe finer details of much fainter astronomical objects than would otherwise be possible from the ground.

However, this technique performs well only if the turbulence does not move faster than the adaptive optics control loop. As the turbulence is carried by the wind, a forecast of

the vertical wind profile above the observatory gives precious advance information about the expected efficiency of the adaptive optics instruments. The forecast can be used to arrange the observing schedule accordingly.

Cooperation since 1998


ECMWF has supported astronomical research at ESO by providing meteorological products for its observatories in northern Chile since 1998. Thanks to a recent decision by its Council, ECMWF is now able to provide products from the Boundary Conditions Programme to research projects and international organisations. ESO is the first international organisation to benefit from this decision and is now receiving hourly forecasts four times a day.

About the author: Marc Sarazin is an Applied Physicist at the European Southern Observatory. For more information on ESO, visit the European Southern Observatory's website: www.eso.org/

Surface verification in the Arctic

THOMAS HAIDEN

To improve weather and climate prediction in polar regions, in June 2012 the World Weather Research Programme (WWRP) established the Polar Prediction Project. One of its key elements is the Year of Polar Prediction (YOPP), which is scheduled to take place from mid-2017 to mid-2019 and is now in its preparation phase. The main goal of YOPP is to improve environmental prediction capabilities for the polar regions by coordinating a period of intensive observing,



Assessing the 'footprint' of Arctic radiation observations. Correlation in ERA-Interim of monthly anomalies of downward longwave radiation between the grid point nearest to the location of Barrow, Alaska, and other areas in the Arctic. The analysis period is 2010–2014. Positive correlations are shown in red, negative correlations in blue.

modelling, prediction, verification, user engagement and education.

As part of these activities, ECMWF has begun to look more into surface verification in the Arctic. One of the challenges in routine verification in this region is the relative sparseness of ground-based observations north of 70°N. There is, however, a set of comprehensively equipped stations coordinated by the International Arctic Systems for Observing the Atmosphere (IASOA) project, which was initiated as part of the International Polar Year (IPY) in 2007 to 2009. In addition to standard meteorological parameters, IASOA provides observations of radiation fluxes, cloud properties, aerosol measurements, air chemistry, surface albedo and stratospheric properties. Most stations are situated along the coastline of the Arctic basin, such as Barrow, Alaska (71°N); Eureka, Canada (80°N); Ny-Ålesund, Svalbard (79°N); and Tiksi, Russia (72°N). The coastal nature of these locations, combined with the effects of seasonally and inter-annually changing sea-ice conditions, makes the forecasting of radiation fluxes at these locations challenging.

Verification results for Barrow over the period 2012–2014 suggest that high-latitude surface radiation is well represented in the IFS, with mean errors less than 10 W/m² in the shortwave and less than 5 W/m² in the longwave part of the spectrum. These values are comparable to, and slightly

smaller than, those obtained from verification against Baseline Surface Radiation Network (BSRN) data from mid-latitude sites. However, day-to-day variability in cloud conditions introduces substantial random errors into the radiation forecast and currently limits the range of positive daily-mean cloud forecast skill in mid- and high latitudes to about 5 days. Since part of the apparently random error may turn out to be systematic if stratified appropriately, further work will focus on the regime-dependence of cloud forecast skill. The goal is to identify which processes contribute most to cloud and radiation errors at different timescales and forecast ranges.

The use of point observations in verification raises the question of the spatial representativeness of the observational data. ECMWF model analyses are being used to assess the degree of correlation at different timescales between anomalies at a given IASOA location and those in other Arctic regions. The results may inform the planning process for existing and future observation sites. For the location of Barrow, positive correlations in monthly downward longwave radiation anomalies are found to extend well into the Arctic basin and southward through the Bering Strait, while negative values appear over European and western Siberian seas. They reflect an anti-correlation of the meridional flow component between these two Arctic sectors resulting from the dominant large-scale modes of northern hemispheric variability, the North Atlantic Oscillation and the Pacific-North American pattern.

ECMWF assimilates data from two new microwave imagers

KATRIN LONITZ,
PETER LEAN, ALAN GEER

Since 12 August 2015, data from two new microwave imagers, the Advanced Microwave Scanning Radiometer (AMSR-2) on board GCOM-W1 and the GPM Microwave Imager (GMI), have been assimilated into operations

using the all-sky approach. Both are research instruments on board satellites launched by the Japan Aerospace Exploration Agency (JAXA) and, in the case of GPM, jointly by JAXA and NASA. The activation of those two instruments has been motivated by the end of the Tropical Rain Measurement Mission (TRMM), including the TRMM

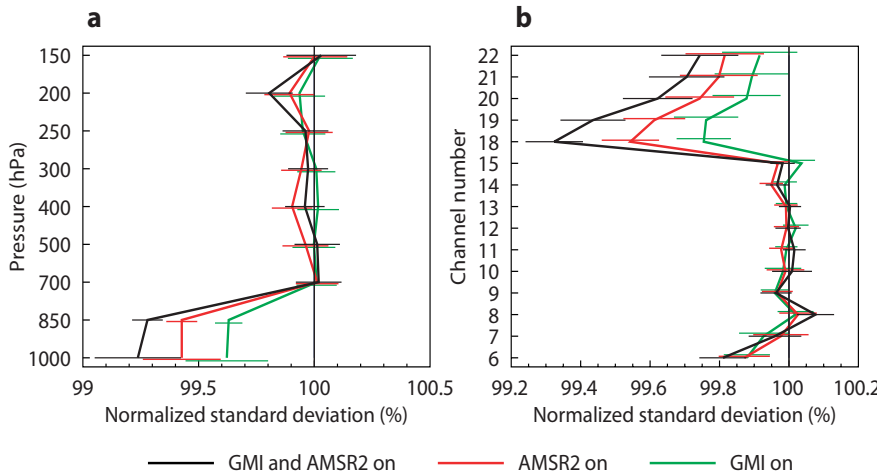
Microwave Imager (TMI), in April 2015 after 17 years of service. This left only one microwave imager, the Special Sensor Microwave Imager/Sounder (SSMIS) on board DMSP-F17, in active assimilation.

The assimilation of data from these two additional microwave imagers has led to an improvement in short-range low-

level wind, temperature and moisture fields, as shown in the figure. Results obtained by Masahiro Kazumori of the Japan Meteorological Agency, who visited ECMWF in 2013 and prepared

our system for AMSR-2, also showed that adding microwave imagers can benefit medium-range forecast scores. It is hoped that, for many years to come, SSMIS, GMI and AMSR-2 will

continue to provide high-quality data improving water vapour, cloud, precipitation and wind fields through the all-sky assimilation system.



Impacts of using additional microwave imager data. The charts show the normalized standard deviation of first-guess departures from (a) atmospheric motion vectors and (b) NASA's Advanced Technology Microwave Sounder, which is sensitive to humidity (channels 18–22) and temperature (channels 6–8). The normalization is done using the results of an experiment where SSMIS-F17 is the only microwave imager assimilated. Values less than 100% indicate beneficial impacts from the use of additional microwave imagery. The horizontal bars indicate 90% confidence intervals. Results cover the period from 26 February 2015 00 UTC to 20 July 2015 12 UTC.

Improved spread and accuracy in higher-resolution Ensemble of Data Assimilations

**ELIAS VALUR HOLM,
MASSIMO BONAVITA,
LINUS MAGNUSSON**

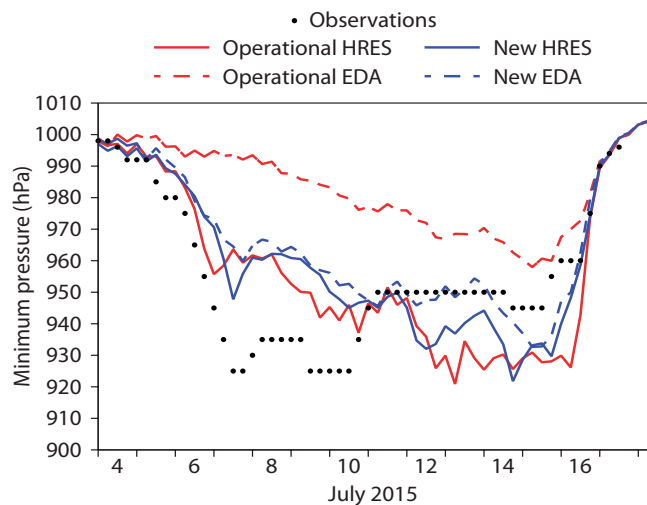
Tests of the next operational resolution of the IFS show that the planned increase in the Ensemble of Data Assimilations (EDA) resolution from 50 km to 18 km brings the EDA into closer agreement with high-resolution forecasts (HRES), which are also being upgraded from 16 km to 9 km. The effect can be seen in the analysis of the central pressure of tropical cyclones.

The chart shows the evolution of the analysed minimum pressure in tropical cyclone Nangka, which made landfall in Japan in July 2015, over two weeks. The estimated minimum pressure from observations is also shown. The results show that in this case the new EDA is in close agreement with the operational and new HRES analyses, whereas the old EDA lacked the resolution to capture the cyclone.

The EDA is used to estimate background errors for HRES from the spread of the EDA ensemble members as well as providing initial conditions for ensemble forecasts

(ENS). The spread has become more realistic with the new 18-km EDA, and has increased, especially around tropical cyclones. The analysis uses the information that errors are increased in the vicinity of a cyclone to allow available observations to have a larger impact there. In addition to increased resolution, the new EDA also uses its

own dynamical errors, whereas before it used a more static estimate of errors. This makes the error estimation in EDA similar to what is done in HRES. It further improves the realism of the error estimates for the HRES analysis and the initial conditions for ENS obtained from the EDA.



Evolution of minimum pressure in analyses for tropical cyclone Nangka. The chart shows observation-based estimates; the current operational HRES; the current operational 50-km EDA control analysis; the new higher-resolution system; and the new higher-resolution EDA control analysis (perturbed member analyses are similar); all from 4 to 17 July 2015.

A first look at the new ecFlow user interface

IAIN RUSSELL,
SÁNDOR KERTÉSZ

ecFlow is a sophisticated workflow package developed at ECMWF for the management of jobs submitted for operational and research tasks. It enables the design, submission and monitoring of jobs in a controlled environment, providing tolerance for software and hardware failures and good restart capabilities.

The existing graphical user interface for ecFlow, ecFlowview, is a modified version of the user interface for ecFlow's predecessor, SMS, which was called xcdp. Originally developed more than 20 years ago using the Motif

toolkit, its user interface is dated, hard to maintain and difficult to add new features to.

Qt is an open source application framework already used at ECMWF for the development of Metview's user interface. With good experiences of using it, we assessed Qt to be the best option for ecFlowUI, the new user interface for ecFlow.

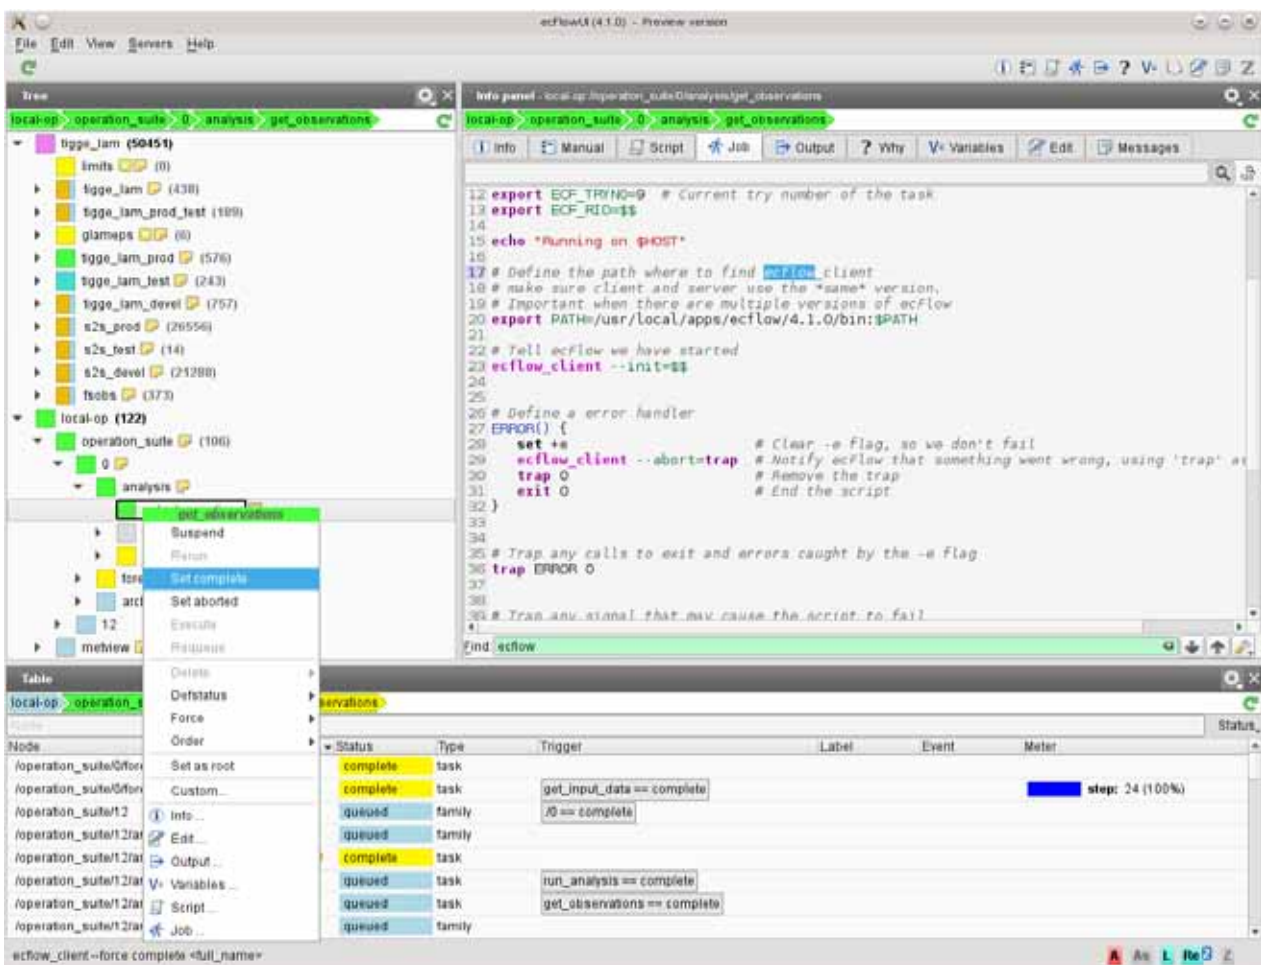
The new user interface

ecFlowUI was designed to look familiar to existing users of ecFlowview while providing new features and new ways of working with ecFlow. The first thing to note is that the user interface is based on a *dashboard* approach, where different panels can be arranged to form the user's

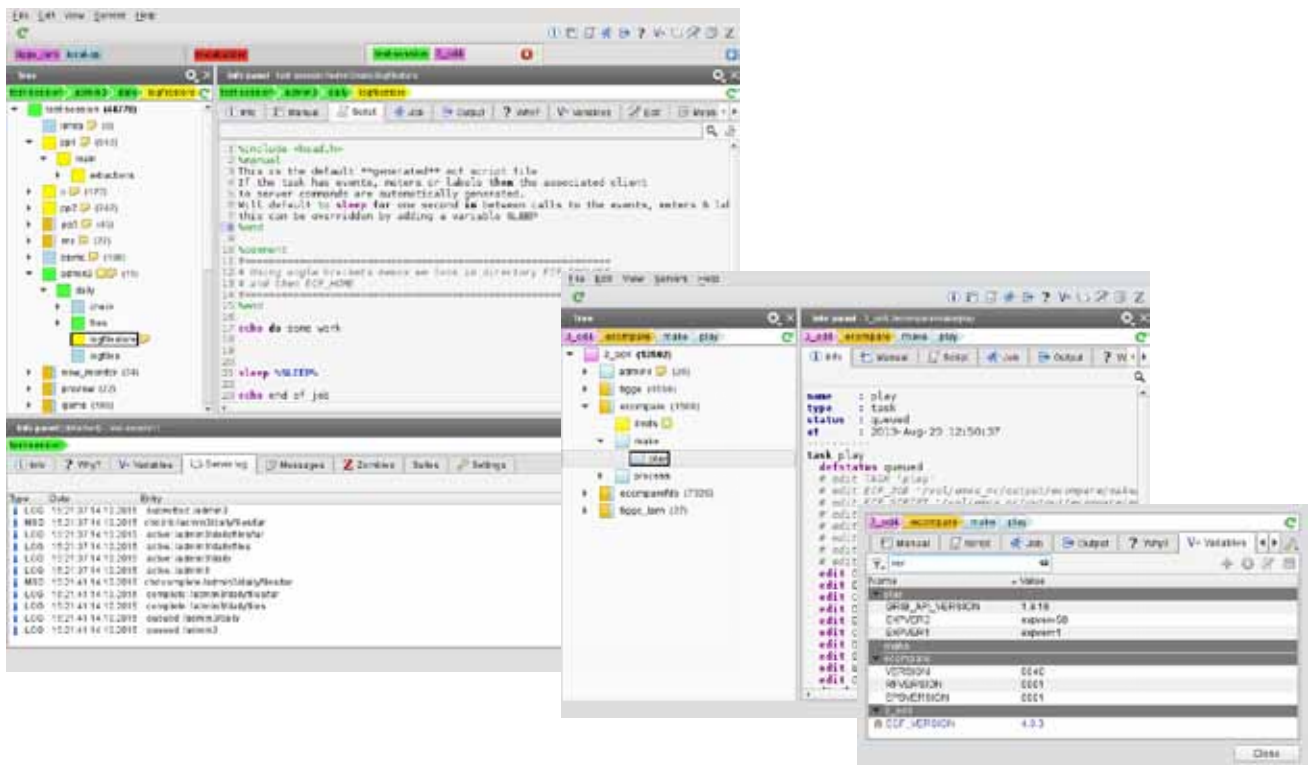
preferred working environment.

There are currently three types of panel which can be used. The **Tree View** panel shows the hierarchical structure of the tasks on each server. A configurable colour scheme shows the status of each node, and options are available to display attributes and icons, and to filter the nodes based on their status. Interaction with the nodes is performed from a node-sensitive context menu.

The **Table View** panel provides a flat view of the tasks on each server. The table can be sorted according to the values in a specified column and provides the same filtering facilities as the Tree View. It is planned to expand



The main ecFlowUI interface. There is a Tree View panel on the left, an Info Panel on the right, and a Table View panel at the bottom. The breadcrumb toolbars at the top of each panel provide further navigational facilities, and a context menu enables interaction with the servers.



ecFlowUI supports multiple windows and tabs. Each window or tab communicates with its own set of servers.

the filtering capabilities to enable more complex queries, effectively turning it into a powerful search panel.

The **Info View** panel provides a set of tabs showing information relevant to the currently selected node. Different tabs show details such as a description of the task, the script behind it, and its latest output. Variables defined for the task can also be viewed, filtered and modified here.

One feature common to all the panels is the navigational *breadcrumbs* at the top, providing a convenient way to change the currently selected node and to see its status and that of its parents.

ecFlowUI supports multiple windows and multiple tabs. Each tab contains its own list of active servers, and the tab heading itself shows the status of each server in that tab.

Performance improvements

Since ecFlowUI potentially has to communicate with many busy ecFlow servers, communication time can become significant. To allow the application to continue working during these transactions, all network communication is performed in separate threads, allowing the user interface to remain responsive. With a revised way of locally storing node information, ecFlowUI also has a smaller memory footprint than ecFlowview, enabling more servers to be viewed at once.

Current status

ecFlowUI is ready to be tested, but some features, such as the time line view, have yet to be implemented. There are also many small improvements that need to be made before it can be considered

a full replacement for ecFlowview. Once this work has been done, and its functionality tested and approved by users, ecFlowUI can become the default user interface for ecFlow.

At the time of writing, ecFlowUI is available to internal users at ECMWF and will become available to external users with the release of ecFlow version 4.1.0. Initial feedback from users is being used to prioritise developments.

More information about ecFlow can be found at: <https://software.ecmwf.int/ecflow>

Living with the butterfly effect: a seamless view of predictability

ROBERTO BUIZZA,
MARTIN LEUTBECHER, ALAN THORPE

Errors in numerical weather prediction grow from the smaller to the larger scales, and even if the initial-time errors are very small, they can eventually grow to affect the larger scales. The less faithful our model representation of the atmosphere is, the faster the error propagates and the sooner predictable signals coming from the initial conditions are lost. This destructive interaction between initial-time and model errors is a consequence of chaos in such predictions, as first noted by French mathematician Henri Poincaré in his book *Science and Hypothesis* (first published in 1902) and then developed by Lorenz in a series of landmark scientific papers (for example 1969a). It is often referred to as ‘the butterfly effect’ and quoted as the main reason why predicting the weather on particular days beyond two weeks is an extremely challenging, if not impossible, venture.

However, ECMWF has been issuing operational monthly forecasts since 2004 and seasonal forecasts since 1997. They show that two weeks can be exceeded provided care is taken in defining the scales one wants to predict. Indeed, some parameters can be predicted with extremely high skill, and average accuracy measures indicate skill definitely beyond two weeks.

Results published in a recent paper (Buizza & Leutbecher, 2015; hereafter *BL15*) and summarized in this article confirm earlier indications (see, e.g., Shukla 1998) that the ‘forecast skill horizon’, defined as the lead time when ensemble forecasts cease to be more skilful than a climatological distribution, is longer than two weeks. This conclusion can be reached if we follow a seamless approach to measure it for forecast fields with increasingly coarse spatial and temporal scales. Thanks to major advances in numerical weather prediction, for some weather parameters forecast skill horizons longer than two weeks are now achievable even on relatively fine scales.

In other words, we have learned to live with the butterfly effect.

How did we manage to live with the butterfly effect?

One key aspect that has provided encouragement is the realisation that there are some large-scale/low-frequency phenomena that are more predictable than smaller-scale/higher-frequency ones. These include Rossby waves, organized tropical convection (e.g. the Madden-Julian Oscillation, MJO), long-range teleconnections, El Niño, and boundary forcing due to long-lived anomalies (e.g. in sea-surface temperatures, soil moisture and sea ice).

Shukla (1998) talked about “predictability in the midst of chaos” to explain how skilful long-range predictions of phenomena such as El Niño were possible despite fast error growth rates from small to large scales. Hoskins (2013) talked about “discriminating between the music and the noise” and introduced the concept of a predictability chain, whereby, for example, “a large anomaly in the winter stratospheric vortex gives some predictive power for the troposphere in the following months”.

The second key aspect has been the adoption of an ensemble-based, probabilistic approach in operational weather forecasting. The shift from a purely deterministic to a probabilistic approach, with ensembles of numerical integrations used to estimate the probability density function of forecast states, has made it easier to extract predictable signals in the extended forecast range. Reliable ensembles have given us a scientifically sound way to assess when a long-range forecast can be trusted.

Other factors that have helped to extract predictable signals in the extended time range include:


- model improvements, such as better convection schemes (see, e.g., Bechtold *et al.*, 2012)
- the inclusion of more relevant processes, for example relating to the ocean
- higher model grid resolutions thanks to much faster supercomputers
- more accurate estimates of the initial conditions thanks to advances in the global observing system and data assimilation methods.

If we consider the ECMWF monthly ensemble, which is based on a coupled ocean–land–atmosphere model, forecast skill has improved significantly since operational production started in 2004, in particular for the prediction of large-scale, low-frequency events. Looking at the MJO, Vitart *et al.* (2014) showed that the 2013 version of the ECMWF monthly ensemble predicted it skilfully up to about 27 days, indicating a 1-day gain in skill per year since 2004. Considering the North Atlantic Oscillation (NAO), they reported that the ECMWF ensemble showed skill in 2013 up to about forecast day 13, compared to about day 9 ten years earlier.

On the longer, seasonal time scale, large-scale phenomena such as monthly-average sea-surface temperature anomalies linked to El Niño, or monthly-average ocean-basin tropical storm activities, can be skilfully predicted months ahead using the ECMWF seasonal System 4, operational since November 2011. It is time to move forward and shift our thinking from a two-week predictability limit to the idea that we can live with the butterfly effect.

The concept of a skilful probabilistic forecast

A



The concept of a skilful probabilistic forecast has been used as the basis to estimate the forecast skill horizon of ECMWF ensemble forecasts. The climatological (reference) probability density function (red line) has only a small overlap with the observation, represented by a probability density function which becomes arbitrarily narrow in the limit of zero observation error (black line). By contrast, the forecast probability density functions (violet, green and blue) match the observation probability density function increasingly well as the forecast time shortens. Consistently, the forecast cumulative distribution functions (defined by integrating the probability density functions; not shown) also approach the observation cumulative distribution function. To measure how close the forecast and observed distributions are, the Continuous Ranked Probability Score (CRPS), which is equal to the mean squared distance between the forecast cumulative distribution function and the observed cumulative distribution function, can be used.

The forecast skill horizon

BL15 assessed the predictive skill of the ECMWF operational medium-range/monthly ensemble (ENS) from July 2012 to July 2013. At that time, ENS included 51 members and was run with a spectral triangular truncation T_{L639} (corresponding to a horizontal resolution of about 32 km) up to forecast day 10, and with a T_{L319} spectral truncation (corresponding to a horizontal resolution of about 64 km) from day 10 to day 32. It had 91 vertical layers up to 0.01 hPa and was coupled to a dynamical ocean model from forecast day 10.

As is normal practice when issuing long-range forecasts, ENS forecasts were corrected by removing the model bias computed using the ENS re-forecast suite, which at that time included five-member ensembles covering the preceding 20 years. As a reference forecast from which to compute a skill score, they used a climatological ensemble (ENSCLI), defined by 100 sequences of reanalyses covering the re-forecast period (see *BL15* for more details).

They looked at predictability in a seamless way by considering ensemble-based, probabilistic forecasts averaged over different spatial and temporal scales, and applying the same approach (i.e. the same metric to the same variables) to investigate whether the forecast skill was sensitive to the scale characteristics. Spatial- and time-averaging procedures were applied to extract the most predictable signal from the grid point fields. Spatial smoothing was obtained by applying spectral filters to the original fields, from T_{120} (spectral truncation with total wave number 120, corresponding to about 170 km horizontal grid spacing) to T_7 (about 3,000 km horizontal grid spacing). Time-averaging was obtained by considering instantaneous fields, or to be more precise fields defined for a 40-minute period (i.e. twice the time step used in the numerical integration), and fields averaged over 1, 2, 4, 8 and 16 days.

BL15 defined the 'forecast skill horizon' as the lead time when the bias-corrected ensemble forecast ceased to be

more skilful than the climatological distribution. More precisely, the forecast skill horizon was computed as the forecast time when the average CRPS of the bias-corrected ensemble stopped being statistically significantly lower, at the 99th-percentile level, than the CRPS of the climatological ensemble.

Sensitivity to spatial and temporal scales

Figure 1 shows the annual-average CRPS of the bias-corrected ensemble (ENSBC) and of the climatological forecasts (ENSCLI), and the skill score of ENSBC defined simply as the difference between the two, as follows:

$$CRPSS(ENS_{BC}) = CRPS(ENS_{CLI}) - CRPS(ENS_{BC})$$

Forecasts are for the 2-day time averages of 850 hPa temperature at T_{120} spectral truncation over the northern hemisphere (NH, points with latitude north of 30°N), the southern hemisphere (SH, points with latitude south of 30°S) and the tropics (TR, points with latitude between 20°S and 20°N), valid from 12 hours to 32 days, every 12 hours. The CRPSS confidence intervals hit the zero line at forecast day 25 for the northern hemisphere, at day 18 for the southern hemisphere and at day 26 for the tropics, indicating that the forecast skill horizon can be longer than two weeks.

Figure 2 shows the sensitivity of the forecast skill horizon to the temporal scale under consideration, for T_{120} spatial fields. It shows that time-averaging reduces not only the difference between the climatology and observations (measured by $CRPS(ENSCLI)$) but also the error growth rate (the rate at which the difference between the forecast and observations grows as lead times increase, measured by the slope of $CRPS(ENSBC)$). The net effect is that, as the time-averaging period is progressively extended from 40 minutes to 16 days, the forecast time when the two average CRPS curves intersect moves to larger values. If we consider, for example, the northern hemisphere (Figure 2a), the forecast skill horizon increases from 23 days for 40-minute average fields to 24.5 days for 1-day average fields, 25 days for 2-day

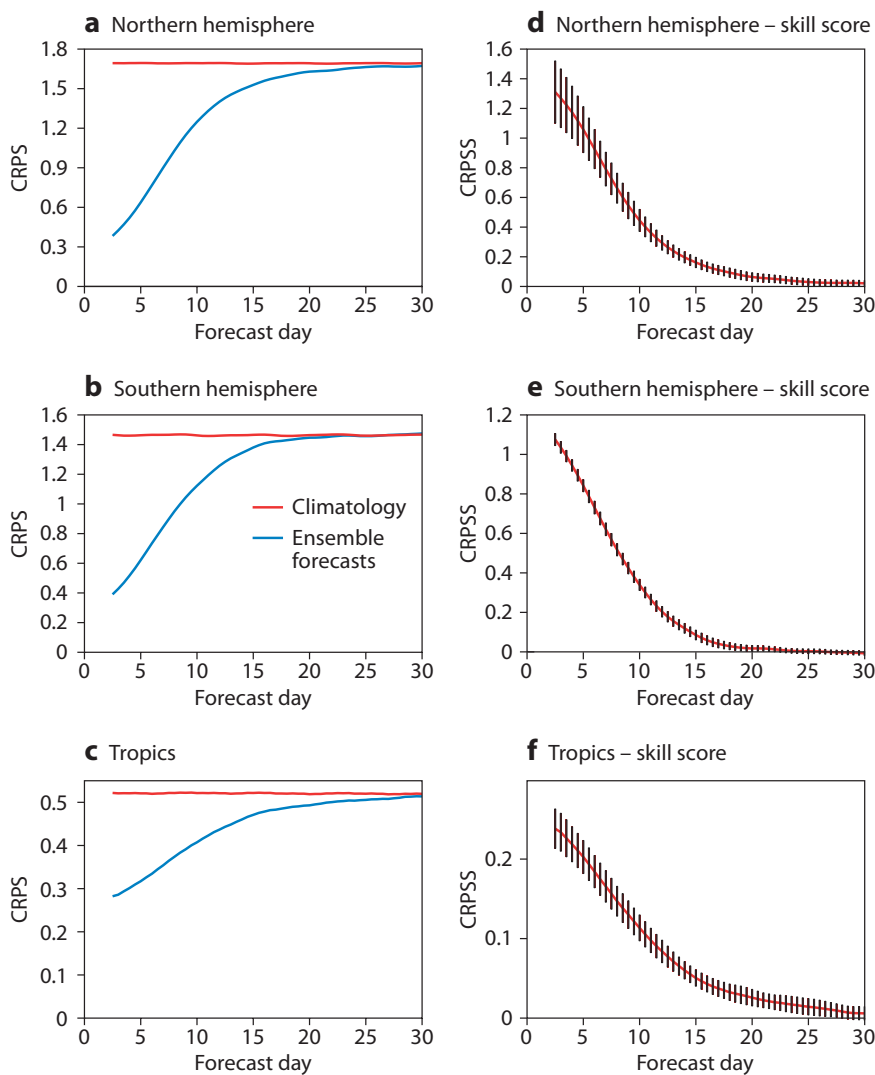


Figure 1 CRPS of the bias-corrected ensemble (ENSBC) and the reference climatological ensemble (ENSCLI) for 850 hPa temperature fields, for a 2-day time average and a T120 spectral truncation (about 170 km horizontal spacing) for (a) the northern hemisphere, (b) the southern hemisphere and (c) the tropics; and the CRPSS(ENSBC) (see text for definition) with 98th percentile confidence intervals for (d) the northern hemisphere, (e) the southern hemisphere and (f) the tropics.

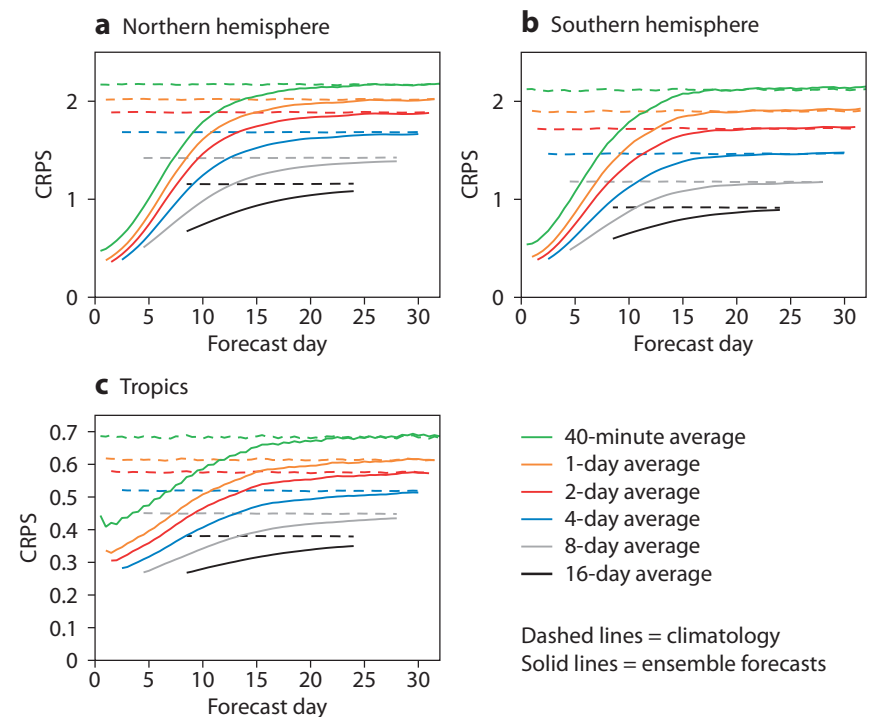


Figure 2 Annual-average (107 cases) CRPS of the bias-corrected ensemble (ENSBC, solid lines) and the reference climatological ensemble (ENSCLI, dashed lines), for 850 hPa temperature fields with a T120 spectral truncation (about 170 km horizontal spacing) and with different degrees of time-averaging for (a) the northern hemisphere, (b) the southern hemisphere and (c) the tropics. Confidence intervals, which are essential to determine the forecast skill horizon, are not shown here for simplicity.

average fields and more than 30 days for time-averaging periods of 4 days.

Results shown in Table 1 for 850 hPa temperature forecasts over the three regions (NH, SH and TR) indicate that the sensitivity to time-averaging is stronger than the sensitivity to spatial filtering. It is worth pointing out that the forecast skill horizon depends also on the geographic area, the season and the forecast field.

Generally speaking, these results indicate that larger spatial and temporal scales are more predictable by between 5 and 12 days than finer scales. They show that we should be more specific when we talk about the forecast skill horizon. The horizon is finite, but it is not the same for all scales, fields, areas and seasons.

Clearly, close to the forecast skill horizon the skill level is very small in absolute terms, and the number of users who can exploit this level of skill may be very limited. Nevertheless, our definition of the forecast skill horizon is objective, and, in line with general practice, defined by comparing the skill of a forecast with that of a well-defined unskilled reference forecast.

Implications for forecasters

The results presented here show that ECMWF bias-corrected ensemble forecasts over a wide range of spatial and temporal scales can have forecast skill horizons longer than the two weeks estimated by Lorenz (1969a, b). This result is not new: it confirms results published over the last 20 years (e.g. Shukla, 1998) that were obtained by applying different methodologies and metrics and looking at some specific phenomena, such as the MJO, the NAO and El Niño. What is new in this work is that the forecast skill horizon has been measured in a seamless way for forecasts with different spatial and temporal scales, by applying the same metric from day 0 to forecast day 32.

These results have clear implications for forecasters: large-scale features, such as the presence of blocking conditions over the Euro-Atlantic sector, can be predicted with some skill three to four weeks ahead. This means that forecasters can look at long-range forecasts of large-scale, time-average features to detect whether, for example, the weekly-average temperature distribution at the surface is shifted towards warmer or colder conditions. But at this forecast range they cannot expect to extract finer details, such as what the temperature at a specific location will be. They will have to wait until the forecast range is, for example, only a few days to be able to skilfully predict these finer details, such as the temperature on a specific day and at a precise location. In other words, as the forecast lead time shortens, the forecaster can zoom in on details at increasingly finer spatial and temporal scales.

Visualising the forecast skill horizon

Figure 3 introduces the forecast skill horizon diagram, which can be used to visualise seamlessly the fact that predictive skill depends on the spatial and temporal scales of forecast phenomena. In the diagram, the x-axis is the horizontal spatial scale of predicted aspects of the weather (in km), and the y-axis is the forecast skill range, in days. The grey rectangle to the left of $x = 32$ km identifies the scales that are definitely unresolved in the ECMWF ENS, which today has a grid spacing of about 32 km. It should be noted that the effective resolution, in the sense of the ability to represent a weather feature adequately, can be lower than the grid spacing by a factor of five or more.

In Figure 3, the forecast horizon skill diagram is applied to the ECMWF medium-range/monthly ensemble (ENS) forecasts studied in BL 15. Each coloured line illustrates how the forecast skill horizon varies with the spatial scale after time averaging has been applied (from the most detailed, 40-minute to the 8-day time-average). Each line represents an average forecast skill horizon, computed considering

Temperature 850 hPa	40-minute average			2-day average			8-day average		
	NH	SH	TR	NH	SH	TR	NH	SH	TR
T120 (170 km)	23.0	16.5	22.0	25.0	18.0	26.0	> 28.0	25.0	> 28.0
T30 (680 km)	24.0	17.0	23.0	25.0	18.0	27.0	> 28.0	25.5	> 28.0
T7 (3,000 km)	> 32.0	23.0	26.5	> 31.0	23.5	28.0	> 28.0	> 28.0	> 28.0

Table 1 Forecast skill horizons for the probabilistic prediction of 850 hPa temperature over the northern hemisphere (NH), the southern hemisphere (SH) and the tropics (TR), for fields with increasingly larger spatial scales (T120, T30 and T7 spectral triangular truncation) and longer time averages (40-minute, 2-day and 8-day averages). The greater than symbol (>) indicates that the forecast skill horizon is larger than the last time step that could be verified (i.e. 32 days for 40-minute average forecasts, 31 days for 2-day average forecasts and 28 days for 8-day average forecasts).




Figure 3 Forecast skill horizon diagram with the forecast skill horizon values for upper-level forecasts with different temporal characteristics, computed by *BL15* for the ECMWF medium-range/monthly ensemble (ENS).




Figure 4 Forecast skill horizon diagram highlighting the fact that the forecast skill horizons of the ECMWF medium-range/monthly ensemble (ENS), shown here for the northern hemisphere, have a seasonal dependence.

seven upper-air variables (geopotential height at 500 hPa, temperature and wind components at 850 and 200 hPa) and three areas (NH, SH and TR). The 8-day average line is positioned higher in the vertical than the 40-minute line, reflecting the fact that time-averaged fields are more predictable than instantaneous fields. The lines for the other time-average periods (1, 2 and 4 days) lie between these two.

As shown in Figure 4, the diagram can be used to visualise the fact that the forecast skill horizon is seasonally dependent: over the northern hemisphere, for example, the forecast skill horizon is longer in winter than in summer. In other words, everything else being equal, ENS probabilistic forecasts for the period from December 2012 to February 2013 had longer forecast skill horizons than forecasts for

the period from June 2012 to August 2012. The fact that over the northern extra-tropics predictability is higher in winter than in summer could be linked to the fact that winter is dominated more by large-scale and slow-evolving phenomena such as blocking, while in summer fast and smaller-scale convection is more active.

To give a more complete picture of where the forecast skill horizon of ECMWF ensemble forecasts lies, Figure 5 shows skill estimates also for shorter-range and extended-range forecasts, for example by including ECMWF seasonal forecasts. We have also included skill estimates for surface fields, such as 2-metre temperature, and for precipitation, as well as for indices of large-scale/low-frequency patterns of variability, such as the North Atlantic Oscillation (NAO), the Madden-Julian Oscillation and El Niño.

Figure 5 visualizes the forecast skill horizon considering all these fields. The red lines relating to the instantaneous and finer-scale surface variables are closer to the x-axis, illustrating the fact that surface variables are less predictable. By contrast, the blue lines relating to large-scale patterns identified by teleconnection indices (e.g. NAO and MJO) and to the average sea-surface temperature (SST) in the Pacific regions affected by El Niño are further away from the x-axis and closer to the top-right part of the diagram. This illustrates the fact that these large-scale patterns can be skilfully predicted months ahead.

Two further features have been added to the diagram: a blue-line envelope, drawn schematically to include all the individual lines, and a red 'no-skill' region. The blue line shows where the forecast skill horizon for the ECMWF ensemble forecast is today: relatively short, less than 10 days, for very detailed forecasts, but transitioning to very long horizons of up to a year for monthly-average SST forecasts for regions in the tropical Pacific. The blue line is not straight, parallel to the x-axis and with a value of about two weeks, as it would be if there were a fixed limit to predictability, but it is curved, reflecting the fact that the forecast skill horizon is scale-dependent, variable-dependent, area-dependent, and season-dependent.




Figure 5 The forecast skill horizon of ECMWF operational forecasts, constructed using published skill measures of medium-range/monthly (ENS) and seasonal (S4) forecasts.

FURTHER READING

Bechtold, P., P. Bauer, P. Berrisford, J. Bidlot, C. Cardinali, T. Haiden, M. Janousek, D. Klocke, L. Magnusson, A. McNally, F. Prates, M. Rodwell, N. Semane, F. Vitart, 2012: Progress in predicting tropical systems: The role of convection. *ECMWF Research Department Technical Memorandum no. 686*.

Buizza, R., & Leutbecher, M., 2015: The Forecast Skill Horizon. *ECMWF Research Department Technical Memorandum no. 754*. Also *Q. J. Roy. Meteorol. Soc.*, in press.

Hoskins, B. J., 2013: Review article: the potential for skill across the range of the seamless weather-climate prediction problem: a stimulus for our science. *Q. J. R. Meteorol. Soc.*, **139**, 573–584.

The forecast skill horizon challenge

The forecast skill diagram, built with a seamless approach to fields with different spatial and temporal characteristics, visualises our understanding of the predictability of different scales of atmospheric variability. It illustrates the fact that, thanks to scientific and model advances, a better estimation of initial conditions using more and better observations and more accurate assimilation methods, and the use of ensemble methods, we can extract predictable signals from the larger scales, notwithstanding the upscale error propagation from the small scales.

The interplay between the downscale propagation of predictability and the upscale error propagation – the battle between the sources and sinks of predictive skill – determines where the forecast skill horizon lies.

By further improving our ensembles, increasing the resolution of the models, reducing the initial error, and adding a better and more complete description of all relevant phenomena, we can aim to further push this line towards the top-left corner. In other words, we can live with the butterfly effect and further extend the forecast skill horizon.

Lorenz, E. N., 1969a: The predictability of a flow which possesses many scales of motion. *Tellus*, **XXI**, 3, 289–307.

Lorenz, E. N., 1969b: How much better can weather prediction become? *Technology Review*, July/August, 39–49.

Shukla, J., 1998: Predictability in the midst of chaos: a scientific basis for climate forecasting. *J. Atmos. Sci.*, **38**, 2547–2572.

Vitart, F., G. Balsamo, R. Buizza, L. Ferranti, S. Keeley, L. Magnusson, F. Molteni, & A. Weisheimer, 2014: Sub-seasonal predictions. *ECMWF Research Department Technical Memorandum no. 734*.

An all-scale, finite-volume module for the IFS

PIOTR SMOLARKIEWICZ, WILLEM DECONINCK,
MATS HAMRUD, CHRISTIAN KÜHNLEIN,
GEORGE MOZDZYNSKI, JOANNA SZMELTER, NILS WEDI

ECMWF hosts the European Research Council-funded project *PantaRhei*, which explores novel numerical methods to complement existing, highly optimised numerical weather prediction (NWP) models. The need for such innovation stems from the fact that state-of-the-art global NWP models using the spectral transform method may become computationally inefficient at very fine resolutions due to the communication overhead associated with global spectral transformations.

As a first step, we have developed an autonomous, all-scale numerical module which uses the finite-volume method (Box A) to supplement ECMWF's Integrated Forecasting System (IFS). This module is compatible with emerging energy-efficient, heterogeneous hardware for high-performance computing (HPC), and it is able to represent elements of real weather on a large range of scales, including cloud-resolving scales.

Motivation

The advance of massively parallel computing in the 1990s and beyond has encouraged finer grid intervals in NWP models. This has improved the spatial resolution of weather systems and enhanced the accuracy of forecasts, while stimulating the development of global non-hydrostatic models. Today many operational NWP models include non-hydrostatic options either for regional predictions or research. However, to date no NWP model runs globally in operations at resolutions where non-hydrostatic effects are important (*Wedi & Malardel, 2010; Wedi et al., 2012*). Such high resolutions are still computationally unaffordable and too inefficient to meet the demands of the limited time window for distributing global forecasts to regional NWP recipients and, ultimately, the public.

Efforts to ensure the computational affordability of global non-hydrostatic forecasts face a twofold difficulty. On fine grids the spectral transform method becomes computationally inefficient because of the required global data-rich inter-processor communications (*Wedi et al., 2013*). Therefore, simply scaling up the number of processors would be unaffordable, not least due to the huge increase in electric power consumption this would entail.

At the same time, replacing hydrostatic primitive equations (HPE) that have been central to the success of weather and climate prediction exacerbates the efficiency problem. In particular, with the simulated vertical extent of the atmosphere thin compared to its horizontal extent, the vertically propagating sound waves supported by the non-hydrostatic Euler equations, from which HPE derive, impose severe restrictions on the numerical algorithms. The hydrostatic balance assumption

underlying HPE conveniently filters out vertically propagating sound waves, therefore permitting large time steps in the numerical integration. Moreover, HPE imply the separability of horizontal and vertical discretisation, thus facilitating the design of effective flow solvers, such as the semi-implicit semi-Lagrangian (SISL) time stepping combined with the spectral-transform spatial discretisation that is used today. Such separability does not apply in non-hydrostatic models.

While NWP strives to extend its skill towards finer scales, non-hydrostatic research models endeavour to extend their realm towards the global domain. The two routes of development must meet, but the way to merge the different areas of expertise is far from obvious. Altogether, NWP is at a crossroads. Although massively parallel computer technology promises continued advances in forecast quality, the latter cannot be achieved by simply applying the existing apparatus of NWP models to ever finer grids.

A new way forward

Recognising the predictive skill of the IFS, we seek to address the challenges outlined above by supplying a complementary non-hydrostatic dynamical module with the capabilities of a cloud-resolving model, concurrently driven by large-scale IFS predictions based on the HPE. The first step towards this paradigm is the development of an autonomous, global, finite-volume, non-hydrostatic dynamical module capable of working on the IFS's reduced Gaussian grid and, in principle, on any horizontal mesh.

Partial differential equations (PDEs) require the calculation of differential operators. In the IFS, spatial differentiation is conducted in spectral space, and there are no practical means of calculating derivatives locally in the physical

Finite-volume method

The finite-volume method is an approach to the approximate integration of partial differential equations (PDEs) describing natural conservation laws. Similar to the finite difference method or finite element method, solutions are calculated at discrete places on a meshed geometry.

'Finite volume' refers to the small volume surrounding each node point on a mesh. In PDEs, integrals of the divergence terms over these finite volumes are converted into surface integrals using the Gauss divergence theorem. On a discrete mesh, these surface integrals are then evaluated as a sum of all fluxes through individual surfaces bounding each finite volume (Figure 1).

Because the flux entering a given volume is identical to that leaving the adjacent volume, these methods are conservative. An important advantage of the finite-volume method is that it can easily be formulated for unstructured meshes.

A

space. Moreover, the increased complexity of the currently available non-hydrostatic IFS option significantly increases the number of spectral transforms required, nearly doubling the cost compared to the hydrostatic forecast, which is already projected to be too slow on existing supercomputers at non-hydrostatic resolutions within the operational time window of 1 hour.

Supplementing the IFS with a complementary non-hydrostatic module makes it possible to use numerical procedures unavailable in the SISL spectral dynamical model with minimal disruption to its highly optimised code. In particular, such a module can replace global communication and computation with local equivalents, providing an effective testing ground for assessing the utility of innovations from outside the realm of spectral methods in the context of real weather. Realisations of this concept can be as simple as advecting selected critical fields in a conservative finite-volume fashion. Or they can be as complex as replacing the whole dynamical core with a cloud-resolving model, while driving the latter with a coarser solution of the spectral-transform hydrostatic model. Regardless of the envisaged realisation, having in place an alternative non-hydrostatic module capable of efficiently exchanging dependent variables with the IFS model, but also of providing complete solutions by itself, is a prerequisite for progress.

Discretisation

The new, autonomous, all-scale finite-volume module (hereafter FVM) is a hybrid descendant of the interdisciplinary all-scale Eulerian/semi-Lagrangian research model EULAG (Smolarkiewicz *et al.*, 2014) and the generalisation of its non-oscillatory forward-in-time numerics to unstructured meshes (Szmelter & Smolarkiewicz, 2010; Szmelter *et al.*, 2015). The default governing PDEs of the FVM are based on the fully compressible Euler equations, and simplified soundproof PDEs are included as an option. The standard latitude-longitude (lat-lon) spherical framework forms the basis of the analytical




Figure 1 An unstructured median-dual mesh in 2D. The edge connecting the nodes i and j of the primary polygonal mesh intersects, at its mid-point, the face S_j shared by computational dual cells surrounding the nodes i and j ; the white circles represent the centres of mass of the primary mesh, while the blue and black lines belong to the primary and the dual mesh, respectively.

formulation, as it does in the IFS. The notorious stiffness of the global lat-lon system is circumvented by the flexibility of the finite-volume discretisation on an unstructured mesh. Figure 1 illustrates such a discretisation. Its key geometric objects are the *nodes* in which data reside and the solution is sought, the *edges* that connect nodes forming the primary mesh, and the *dual cells* over which differential operators are evaluated using vector calculus. In the FVM, unstructured meshes can be built around arbitrarily specified nodes. Of particular interest are meshes built around the grids admissible in the IFS, shown in Figure 2, as they provide support for spectral transforms.

In the FVM, all dependent variables are co-located in the nodes, accommodating both spectral transforms and grid-point solutions within the same code while facilitating time-stepping schemes with a large degree of implicitness. In three spatial dimensions, the mesh is prismatic, that is, the horizontal discretisation is common to all vertical levels

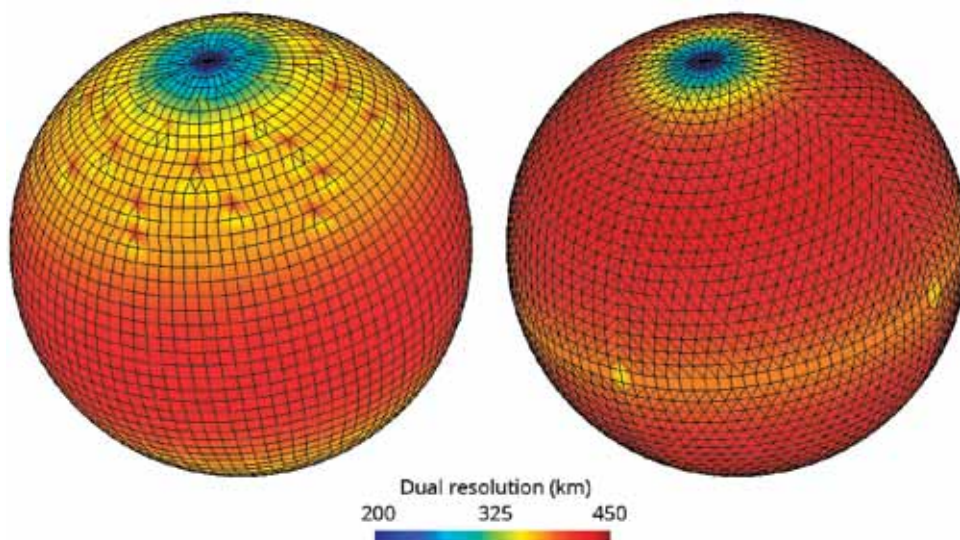


Figure 2 Two primary meshes generated on N24 reduced Gaussian grid points with an approximate resolution of 3.75° (415 km). N24 indicates 24 latitudes between a pole and the equator. Gaussian grids are latitude-longitude grids. In reduced Gaussian grids, the number of grid points on latitudes near the poles is reduced to achieve a more even spacing of grid points. The shading represents the dual resolution, computed as the square root of the local dual volume.

of the computational space. A uniform finite-difference discretisation in the vertical is chosen to facilitate the solution of intricate elliptic problems in thin spherical shells. To accommodate planetary orography and a pliant height coordinate in physical space (such as hydrostatic isobars), the governing equations are cast in a generalised time-dependent curvilinear reference frame. To numerically solve the governing PDEs, the FVM employs proven semi-implicit non-oscillatory forward-in-time integrators of the governing PDE systems (Smolarkiewicz *et al.*, 2014) conceptually similar to, but more general than, those used in the IFS.

Data structure and parallelisation

Flexible object-oriented data structure

The conceptual building blocks of the FVM were established in EULAG, and they are well documented in the literature. Nonetheless, many aspects of the discrete apparatus had to be customised for the parallel implementation on the unstructured mesh constrained by the IFS grid. In particular, the FVM module is based on a newly developed framework called Atlas.

The Atlas framework provides parallel, distributed, flexible, object-oriented data structures for both structured and unstructured meshes on a sphere. It separates concerns of mathematical model formulation and numerical solutions from the cumbersome management of unstructured meshes, distributed memory parallelism, and input/output of data. While the FVM uses FORTRAN 2003, it is recognised that handling flexible structures and carefully controlled memory-management are not easily achievable using the FORTRAN language. Hence the language of choice for Atlas is C++, which provides excellent object-orientated support and builds upon C's strengths in memory management. A FORTRAN 2003 interface exports all of Atlas's functionality to the FVM.

In principle, the FVM and Atlas can accommodate nearly any form of horizontal meshing. However, in the envisaged context of being able to use spectral transforms on the finite-volume mesh, the nodes have to satisfy certain constraints. They must be distributed in a structured way in a lat-lon domain, where latitudes are distributed as the roots of the Legendre polynomials, and longitudes are distributed uniformly on each latitude while reducing in number going from the equator towards the poles. Given these fixed node locations, Atlas is capable of generating a hybrid triangular/quadrilateral mesh. The IFS's set of reduced Gaussian grids satisfy these requirements, and the meshing approaches shown in Figure 2 illustrate Atlas's capabilities. These meshes are very coarse, for illustrative purposes only. For the target applications, the required meshes are N1280 and N2000, corresponding to grid increments of about 10 and 5 km. For such meshes the coordinates and connectivity tables forming the *Mesh* object in Atlas have an enormous memory footprint. This is why *Mesh* is distributed over parallel tasks.

Parallelisation scheme

Because the FVM operates at the nodes of the IFS grid, it seamlessly inherits the load-balanced equal-regions domain decomposition parallelisation scheme of the IFS (Mozdzyński *et al.*, 2015). Like the IFS, the FVM is designed to run on HPC facilities which group computational cores into tasks. Computational cores in each task have access to shared memory among themselves, whereas tasks do not share memory and are connected by message passing.

The FVM hybridises two standards – the Message Passing Interface (MPI) and Open Multi-Processing (OpenMP) – for parallelisation, each with a different purpose. MPI is a standard handling the communication of data between tasks. It involves distributing data between tasks and requires information to be sent and received at regular intervals to advance a simulation. OpenMP is a standard for easily managing multiple computational threads working with shared memory. For example, several iterations in one *do* loop can be executed simultaneously within an algorithm. Experience gained with the IFS shows that the hybrid MPI/OpenMP parallelisation delivers about 20% better performance than MPI alone. This is the result of an improved load balance and memory scalability for large grids.

The parallelisation of the FVM distributes the computational mesh with its unstructured horizontal index in such a way that structured vertical columns are always preserved in memory as a contiguous entity. The Atlas framework, upon which the FVM is built, is responsible for generating a distributed mesh, and it provides communication patterns using MPI to exchange information between the different partitions of the mesh. To minimise the cost of sending and receiving data, the distribution of the mesh is based on an



Figure 3 Equal region domain decomposition with 1,600 partitions for the N1280 octahedral reduced Gaussian grid (see right-hand panel in Figure 2).

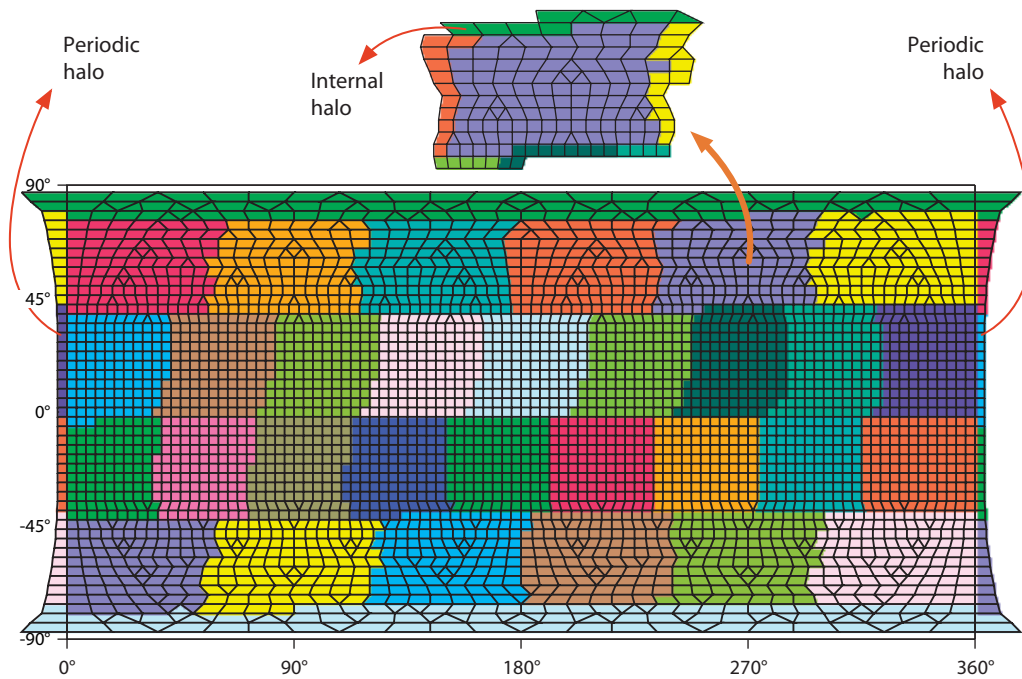


Figure 4 The computational mesh shown in the left-hand panel of Figure 2 is here mapped onto tasks (32 partitions). Also shown is the internal halo of one partition, and the periodic halo responsible for the periodic boundary condition.

equal-regions domain decomposition algorithm for a quasi-uniform node distribution on the globe.

As shown in Figure 3, such a decomposition divides the globe into bands oriented in zonal direction, and it subdivides each band into a number of regions so that globally each region has the same number of nodes. The bands covering the poles are not subdivided.

To evaluate finite-volume differential operators, the FVM requires only nearest-neighbour communications for halos of adjacent partitions. As shown in Figure 4, the halo thickness is limited to one element of the primary mesh. The nearest-neighbour communication available in the FVM contrasts with the global communication required for spectral transforms. It is an essential part of the computational arrangements in the FVM since global communication becomes a severe bottleneck as more and more cores are added to the HPC cluster. The underlying Atlas framework is used to create the internal halos and to manage communication patterns to exchange data in the halos, thus separating this concern from the FVM.

Examples

In the following, we discuss solutions generated with the compressible option of the FVM. In the course of the module’s development, the acoustic and two soundproof options were also extensively used to benchmark and verify the module. These solutions are not shown here, as they corroborate the conclusions of *Smolarkiewicz et al.* (2014) regarding the relative comparability of compressible and soundproof results. Two physical problems are considered: the sheared orographic flow on a small planet with a 20.4 km radius, to highlight the non-hydrostatic quality

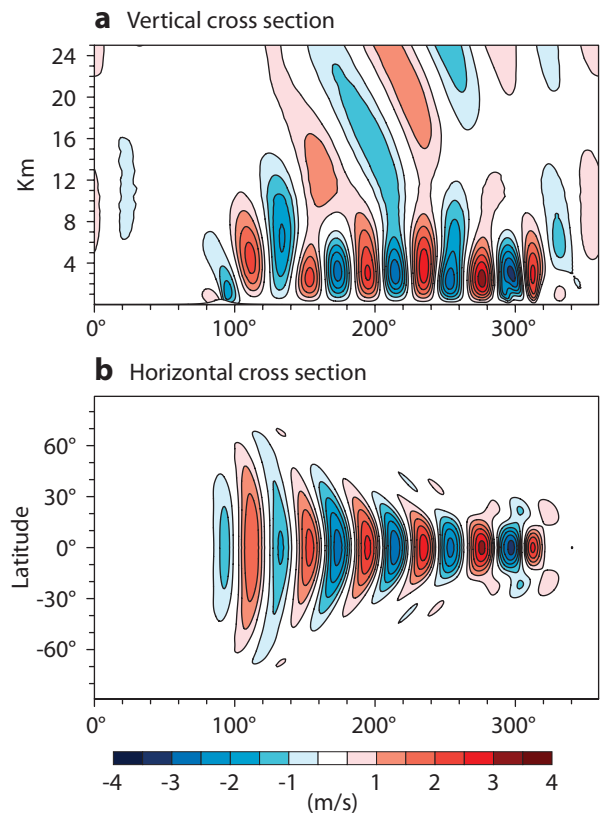


Figure 5 Gravity wave developing in the lee of a hill placed on a tiny planet. The hill can be seen in panel (a), near the lower left corner between 80° and 100°. Red and blue patterns show updrafts and downdrafts, respectively, with contours of vertical velocity in m/s, in (a) vertical cross section along the equator, and (b) horizontal cross section at 3 km above the surface.

of the solver; and baroclinic instability on an Earth-like planet, which concerns essentially hydrostatic global weather. The two problems serve to illustrate diverse aspects of the FVM.

Quasi-two-dimensional orographic flow with linear vertical shear

This classical problem constitutes a particularly discriminating test of a non-hydrostatic model's capabilities, because in the presence of shear the non-hydrostatic and hydrostatic equations predict a fundamentally different propagation of orographically-forced gravity waves. While hydrostatic models produce a vertically propagating gravity wave, the correct solution is that of a trapped, horizontally propagating gravity wave.

The setup of the numerical experiment assumes a smooth mountain range centred at the equator of a tiny planet with a radius of about 20 km. The hill is only 500 m tall and has gentle slopes with a height-to-width aspect ratio of 20%. The assumed atmosphere has a simple structure with a constant Brunt-Väisälä frequency of 0.01 s^{-1} (corresponding to a typical tropospheric value), and wind increasing linearly from 10 m/s at the surface at the equator to about 36 m/s at tropopause altitude (10.5 km), and remaining constant aloft.

The FVM domain is discretised in the horizontal using a finite-volume dual mesh generated around an N128 octahedral reduced Gaussian grid (see Figure 2), with a horizontal mesh spacing of about 250 m at the equator. In the vertical the model domain is resolved with 137 levels stretched smoothly, such that the vertical spacing increases from 70 metres near the ground to 1,400 metres near the top, set at 85 km. The integration time is 2 hours with a time step of 3 seconds.

Figure 5 shows isolines of the vertical velocity after two hours of wave evolution simulated with the FVM. Figure 5a shows the solution in the vertical cross section at the equator, and Figure 5b shows a horizontal cross section over the spherical surface at a height of 3 km. The key result is that there is a trapped lee wave behind the mountain, reminiscent of lee waves of clouds often seen downstream from mountains. Such clouds are a sign of wave updrafts, in which ascending moist air cools adiabatically so the water vapour condenses, before evaporating in the downdrafts. The importance of this result lies in the fact that it cannot be produced with hydrostatic models. If the same problem were solved with HPE, the solution would be a single vertically propagating mountain wave with consecutive lenticular clouds eventually forming above the mountain.

It should be noted that our solution agrees with the numerical results and linear analysis discussed in *Wedi & Smolarkiewicz (2009)*. The lee wave has a dominant horizontal wavelength of about 14 km, corresponding to 40 degrees of longitude on the small planet. The wave energy leakage through the tropopause excites a weaker stratospheric wave with a dominant wavelength, supported by the ambient conditions, about twice that of the trapped wave below. Due to the domain periodicity, the upstream influence of the mountain can

already be seen after 2 hours at the downstream end of the wave packet.

Baroclinic instability

While the preceding section addresses the FVM's non-hydrostatic performance, here we illustrate its skill in simulating essentially hydrostatic motions. To this end, we have repeated the inviscid baroclinic instability simulations used in *Smolarkiewicz et al. (2014)*, where the authors discuss a series of experiments with soundproof and compressible equations using a coarse horizontal resolution of 128 by 64 nodes on a regular lat-lon grid. Here we show the corresponding compressible result generated with the FVM using an N800 octahedral reduced Gaussian grid with a horizontal resolution of 12 km (roughly corresponding to the current operational resolution of the IFS), and 61 stretched vertical levels with the resolution changing from 50 m near the ground, through 150 m at 2 km, to about 850 m at 24 km (the model top).

Figure 6 shows the baroclinic wave train at day 8 after evolving from a weakly perturbed unstable-equilibrium initial state consisting of the two planetary jets in the mid-latitudes. Figure 6a shows isolines of potential temperature

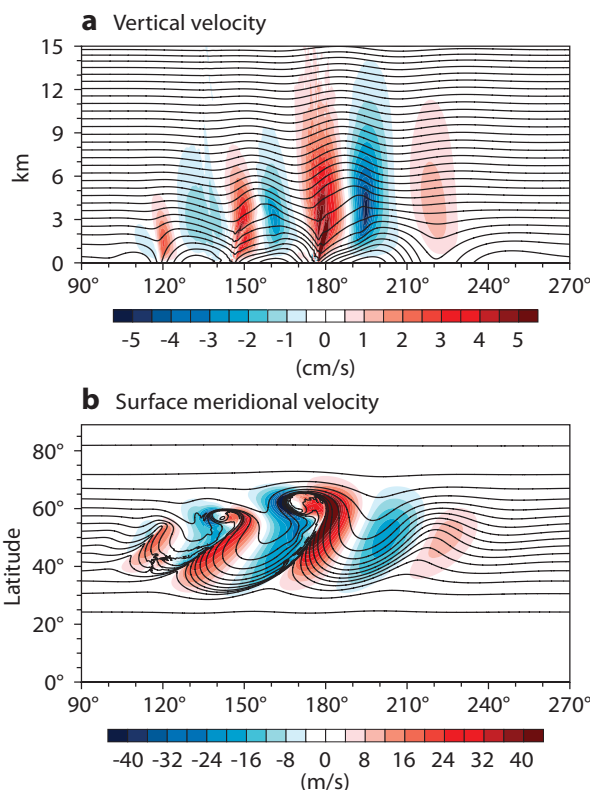


Figure 6 Baroclinic instability, day 8, shown by (a) contours of vertical velocity (shading, in cm/s) overlaid with isentropes (shown with contour intervals of 5 °C) in the vertical cross section at 53°N, and (b) contours of surface meridional velocity (shading, in m/s) and the isentropes. The steepening of the isentropes to the vertical together with the radiation of mesoscale gravity waves and the generation of grid-scale features in (a) are indicative of the frontal collapse in (b).

(isentropes) overlaid with contours of the vertical velocity, in the vertical cross section through the centre of the northerly jet. Figure 6b shows isentropes and the meridional velocity at the sphere's surface. Together these figures illustrate the 3D structure of idealised frontogenesis and the formation of weather systems in the mid-latitudes. These results are consistent with those familiar from the literature. Notably, at the relatively high resolution used, the simulation begins to capture sharp frontal discontinuities and the associated radiation of mesoscale gravity waves (Plougonven & Zhang, 2013). The grid-scale features seen above the collapsing front are inevitable and are controlled by the model numerics.

Summary and outlook

The PantaRhei project pursues an entirely flexible and hierarchical approach capable of satisfying the needs of current and future NWP. The first step is the development of the autonomous FVM supplementing the IFS. The FVM on its own can provide solutions representative of elements of real weather over a large range of scales. It does so using computation and communication patterns which are very different from spectral transform-based global NWP models. The idealised results shown in Figures 5 and

6 relate to key elements of natural weather at very different scales. The same elements can be identified in Figure 7, which shows a snapshot of the FVM simulation of global circulation on an N256 reduced Gaussian grid using the IFS orography together with an idealised frictional/diabatic forcing. The figure illustrates the importance of the FVM's capabilities in the broader NWP context.

The proposed concept offers exciting new opportunities for combining the strengths of spectral-transform models and cloud-resolving models. The FVM and Atlas provide a toolbox of methods previously inaccessible to spectral transform-based models, while the IFS provides the required advanced environment to explore alternative discretisation techniques in the context of operational weather forecasting.

The key to success are the advancement of numerical methods that can be applied to all-scale global flows, and the associated software development for large-scale high-performance computing. The FVM represents an essential algorithmic building block for future NWP and climate services and thus forms part of the future co-design process between ECMWF and leading HPC manufacturers.

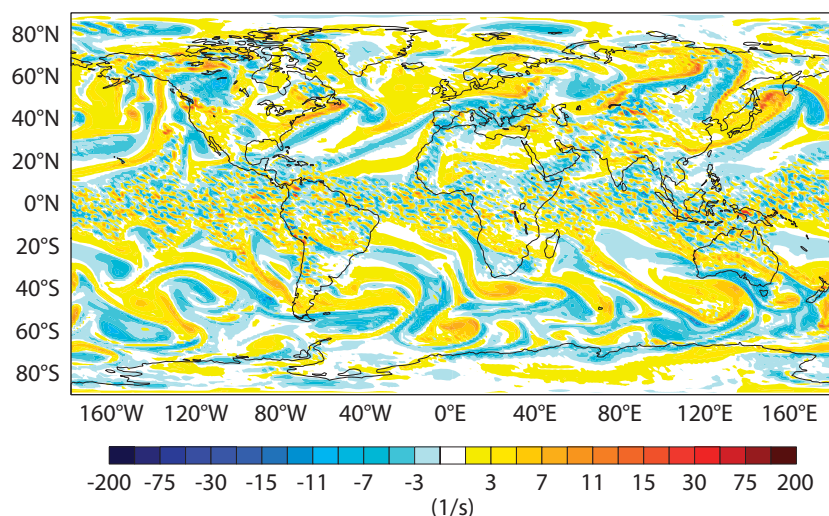


Figure 7 FVM-produced simulation of a global circulation using realistic orography. The shading shows the vertical component of instantaneous relative vorticity ($\times 10^6$) at about 4 km above the surface, revealing baroclinic eddies in the mid-latitudes of both hemispheres and fine-scale features in the equatorial area and mountainous regions indicative of convection and gravity waves.

FURTHER READING

Mozdzynski G., M. Hamrud, & N.P. Wedi, 2015: A PGAS implementation of the ECMWF Integrated Forecasting System (IFS). *Int. J. High Perform. C.*, **29**, 261–273.

Plougonven R., & F. Zhang, 2013: Internal gravity waves from atmospheric jets and fronts. *Rev. Geophys.*, **52**, 33–76.

Smolarkiewicz P.K., C. Kühnlein & N.P. Wedi, 2014: A consistent framework for discrete integrations of soundproof and compressible PDEs of atmospheric dynamics. *J. Comput. Phys.*, **263**, 185–205.

Szmelter, J. & P.K. Smolarkiewicz, 2010: An edge-based unstructured mesh discretisation in geospherical framework. *J. Comput. Phys.*, **229**, 4980–4995.

Szmelter, J., Z. Zhang & P.K. Smolarkiewicz, 2015: An unstructured-mesh atmospheric model for nonhydrostatic

dynamics: Towards optimal mesh resolution. *J. Comput. Phys.*, **294**, 363–381.

Wedi, N.P. & P.K. Smolarkiewicz, 2009: A framework for testing global non-hydrostatic models. *Q. J. R. Meteorol. Soc.*, **135**, 469–484.

Wedi, N.P., M. Hamrud, G. Mozdzynski, 2013: A fast spherical harmonics transform for global NWP and climate models. *Mon. Weather Rev.*, **141**, 3450–3461.

Wedi N.P. & S. Malardel, 2010: Non-hydrostatic modelling at ECMWF. *ECMWF Newsletter No. 125*, 17–21.

Wedi N.P., M. Hamrud, G. Mozdzynski, G. Austad, S. Curic, J. Bidlot, 2012: Global, non-hydrostatic, convection-permitting, medium-range forecasts: progress and challenges. *ECMWF Newsletter No. 133*, 17–22.

Reducing surface temperature errors at coastlines

ROBIN HOGAN, ALESSIO BOZZO

To limit its high computational cost, the current radiation scheme in ECMWF's Integrated Forecasting System (IFS) is run on a coarser spatial grid than the rest of the model. As a result, surface radiative fluxes computed over the ocean may be used over land, where the surface temperature ('skin temperature') and surface albedo are very different. This can lead to large near-surface temperature errors at coastal land points. For example, in Norway night-time 2-metre temperature forecasts at coastlines can sometimes be too low by more than 10 °C. The radiation scheme is also run infrequently, which leads to errors in the diurnal cycle of skin temperature.

A computationally efficient solution to these problems has been developed, in which the surface longwave and shortwave fluxes are updated at every model time step and grid point according to the local skin temperature and albedo. The computational cost of performing approximate radiation updates is only 2% of the cost of the full radiation scheme, and the overall run time of the model is only increased by about 0.2%. Testing the new scheme by running daily five-day forecasts over an eight-month period has revealed a significant improvement in 2-metre temperature forecasts at coastal stations compared to observations.

The ECMWF radiation scheme: accurate but slow

Since 2007, the ECMWF radiation package has been based on the Rapid Radiative Transfer Model for General Circulation Models (RRTMG). This model is very accurate and well validated against both observations and reference radiation calculations that represent individual spectral lines. It also captures cloud structure via the Monte Carlo Independent Column Approximation. However, the model's accuracy comes at a cost: 252 individual radiative transfer calculations are required to represent the full shortwave (solar) and longwave (thermal infrared) spectral regions. The associated computational cost means that we cannot afford to run the radiation scheme at every model time step (typically 10–20 minutes) and grid point (a configuration hereafter referred to as the Reference Scheme). The current operational practice is to run it every hour in the high-resolution (HRES) forecast and every three hours in all other model configurations. To reduce computational cost further, the model fields are first interpolated to a reduced spatial grid with 6.25 times fewer grid points globally, and the resulting radiative fluxes are interpolated back to the original grid (a configuration hereafter referred to as the Operational Scheme).

Figure 1 compares HRES skin temperature forecasts in the region of Long Island in the US state of New York using the Operational Scheme (Figure 1a) with forecasts using the Reference Scheme (Figure 1b). In this case the reduced grid in the Operational Scheme led to an underestimate by as

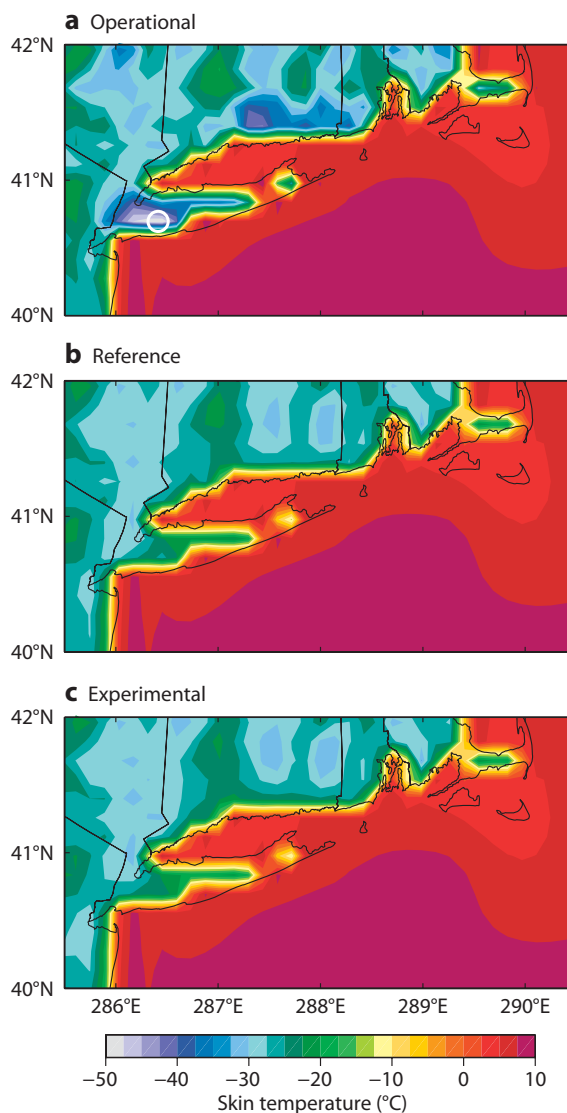


Figure 1 Skin temperature at 0700 local time on 4 January 2014 for a region around the coast of Long Island from forecasts initialized 24 hours previously with the radiation scheme run on (a) a reduced spatial grid every hour; (b) the full model grid at every time step; and (c) a reduced spatial grid every hour but with approximate updates to surface radiative fluxes on the full model grid and at every time step. The white circle highlights an area of particularly big differences between (a) on the one hand and (b) and (c) on the other.

much as 25 °C in one particular area (circled in Figure 1a), and forecasters at La Guardia Airport alerted ECMWF to the poor forecast at the time. Since this error occurred just before sunrise, it is purely a longwave effect. It can be seen that the ocean is much warmer than the land, leading to much larger longwave emission, but this was applied over some coastal land points, leading to them cooling too fast at night. This effect is compounded by the land having a much lower heat capacity than the ocean.

Figures 2a and 2b show early afternoon skin temperature forecasts along a desert coastline. This time the 10 °C overestimate is predominantly a shortwave phenomenon: the lower albedo of the ocean means that much more sunlight is absorbed than by the more reflective desert, but this extra absorption is applied to some coastal land points, leading to them warming too much.

Radiation updates: approximate but fast

We seek a computationally efficient way to update the radiative fluxes at every model grid point and time step to account for the local values of skin temperature and albedo. If we can work solely with broadband fluxes, which bundle different parts of the spectrum together, then this will be much faster than the radiation scheme and its 252 spectral intervals. Note that only differences in surface properties are considered; the atmosphere and in particular clouds are assumed to be constant.

In the longwave part of the spectrum, we can use the fact that the broadband flux emitted by a blackbody is equal to $\sigma T_{surface}^4$ (where σ is the Stefan-Boltzmann constant) to update the surface upwelling flux. This gives a much improved surface energy balance, but we have to decide where to deposit this excess energy in the atmosphere. It turns out that, due to strong water vapour and carbon dioxide absorption, typically around half of the longwave radiation emitted from the surface is absorbed in the first 500 m of the atmosphere. We have therefore configured the radiation scheme to compute the fraction of surface emission that is absorbed by each model level, and used this at every grid point to update also the longwave heating profile. This improves the response of the temperature profile in the lowest part of the atmosphere (boundary layer).

In the shortwave part of the spectrum, the IFS already scales the shortwave flux profile at each time step to account approximately for the change in sunlight incident at the top of the atmosphere as the sun moves through the sky (Morcrette, 2000). We have implemented an additional scheme proposed by Manners *et al.* (2009) that makes an approximate update to surface radiative fluxes to account also for the change in the path length of the solar beam through the atmosphere.

While these two schemes are important, they do not fix the problem highlighted in Figure 2. We therefore perform an update of the reflected surface radiation to ensure that it is consistent with the local value of broadband albedo. However, simply letting the extra reflected radiation escape to space would still result in errors of up to 70 W/m² at desert coastlines, since in reality some of it will be reflected back down to the surface, with the possibility of further upward reflection and so on. To capture this feedback, we use the top-of-atmosphere and surface fluxes from the radiation scheme to compute the effective broadband transmittance and reflectance of the entire atmosphere. This enables the response of both upwelling and downwelling surface fluxes to a change in surface albedo to be computed while accounting for multiple reflections.

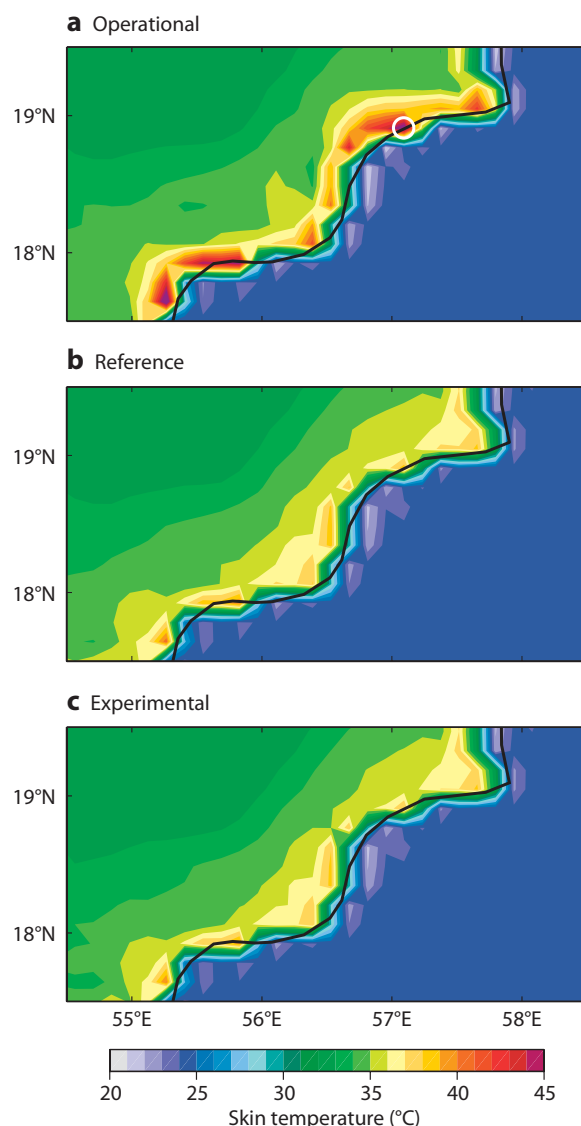


Figure 2 Skin temperature at 1400 local time on 4 January 2014 for a region around the southern coast of Oman from forecasts initialized 24 hours previously with the radiation scheme run on (a) a reduced spatial grid every hour; (b) the full model grid at every time step; and (c) a reduced spatial grid every hour but with approximate updates to surface radiative fluxes on the full model grid and at every time step. The white circle highlights an area of particularly big differences between (a) on the one hand and (b) and (c) on the other.

Figures 1c and 2c show that enhancing the Operational Scheme by the use of approximate radiation updates on the full model grid and at every time step (hereafter referred to as the Experimental Scheme) leads to skin temperature forecasts very similar to those obtained using the Reference Scheme (Figures 1b and 2b), but at a far lower computational cost.

Figure 3 shows the time series of net radiative fluxes (downwelling minus upwelling) and temperature for the point in Long Island with the largest errors in Figure 1. The forecast using the Operational Scheme can be seen to

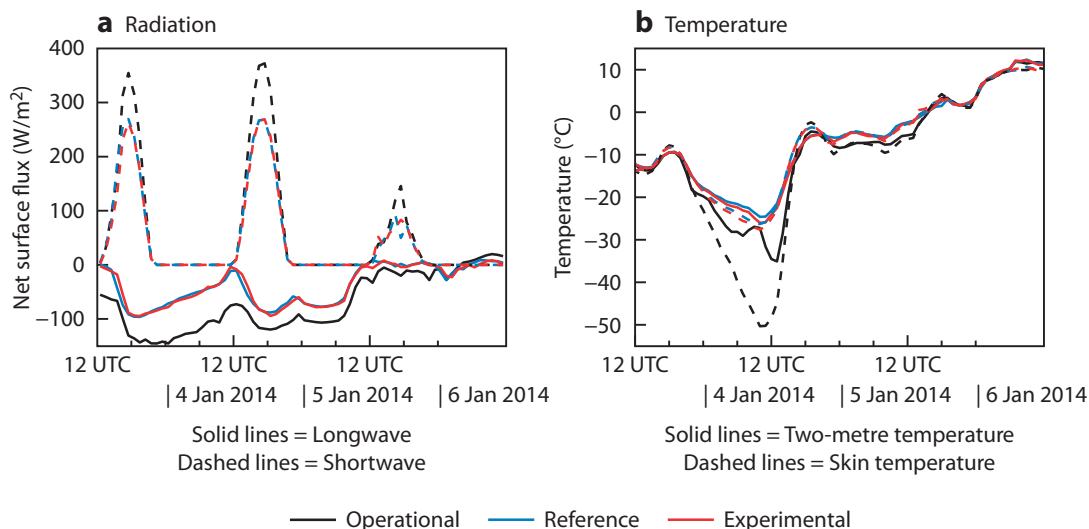


Figure 3 Time series of (a) surface net shortwave and longwave fluxes and (b) 2-metre temperature and skin temperature at 40.7°N 286.4°E (the point on Long Island indicated by the white circle in Figure 1a) for full 72-hour forecasts using the radiation scheme run on a reduced spatial grid every hour (Operational); the full model grid at every time step (Reference); and a reduced spatial grid every hour but with approximate updates to surface radiative fluxes at every time step and grid point (Experimental).

underestimate net longwave flux by around $50 W/m^2$ for the first day of the forecast, compared to a forecast using the Reference Scheme. This error is almost completely corrected by using the Experimental Scheme. Using the radiation updates also corrects a daytime error in shortwave fluxes, which was exacerbated by the snow cover at the time, leading to a large albedo contrast at the coast.

Comparison with observations

Eight months of daily forecasts have been produced using the HRES configuration (T1279, corresponding to a horizontal resolution of about 16 km, with radiation calls every hour) and the ensemble configuration (T639, corresponding to a horizontal resolution of about 32 km, with radiation calls every three hours). Two-metre temperatures have been compared with observations at European coastal stations. Figure 4 shows a time series at a site in northern Norway for December 2012, where it is clear that large temperature underestimates in the T639 model configuration are much improved by the use of the Experimental Scheme.




Figure 4 Comparison of T639 2-metre temperature forecasts at 00 UTC and 12 UTC (lead times of 24 hours and 36 hours, respectively) against observations at Sortland, Norway (68.7°N, 15.42°E) for December 2012.

Figure 5 depicts the full eight months of data in the form of probability distributions of 2-metre temperature error. The distributions are significantly skewed for the Operational Scheme, with a tail of large positive temperature errors in summer at 12 UTC and a tail of large negative errors in winter at 00 UTC. These errors are reduced when using the Experimental Scheme, and indeed the error distributions are quite similar to running the Reference Scheme.

Improvement of the diurnal cycle and forecast skill

The new approach of performing approximate radiation updates on the full model grid and at every time step clearly results in significant improvements in coastal areas at very small computational cost, but the fraction of the

globe affected is very small, so we might not expect a measurable improvement in global forecast skill. However, in model configurations that call the radiation scheme only every three hours, all land area forecasts currently suffer from the problem that longwave emission from the surface is held at a fixed value for three hours even though the skin temperature evolves significantly in this time. This lag in the longwave response to skin temperature changes means that land areas tend to warm up too rapidly in the morning and cool too rapidly at night.

Figure 6 demonstrates this effect: Figure 6a shows the mean June–September 2012 skin temperature error at 12 UTC for forecasts using the Operational Scheme, evaluated against forecasts using the Reference Scheme.

We see a warm bias in land areas where 12 UTC corresponds to early afternoon, and a cold bias for night-time areas. Figure 6b shows a significant improvement as a result of applying the Experimental Scheme.

The impact on forecast skill for other weather parameters has also been assessed. Running the Reference Scheme rather than the Operational Scheme, we find no significant improvement to pressure forecasts or to temperature forecasts above the mid-troposphere. There is, however, an improvement to 1,000 hPa temperature forecasts that is measurable out to at least five days. As shown by Hogan & Bozzo (2015), if we use the Experimental Scheme, then the skill of 1,000 hPa temperature forecasts is improved, but only by half as much as when running the Reference Scheme. This suggests that only half the improved skill from more frequent radiation calculations is associated with the better response to surface conditions, while the rest is associated with the response to changes in the atmosphere, which are not captured by approximate radiation updates.

Outlook

The method described here, for updating radiative fluxes and heating rates between calls to the full radiation scheme, successfully corrects coastal errors caused by the radiation scheme being run on a coarser grid than the rest of the model and less frequently than the model time step. It also improves the diurnal cycle of skin temperature over land. The approximate update method has been incorporated into IFS model cycle 41r2 and is expected to become operational in spring 2016.

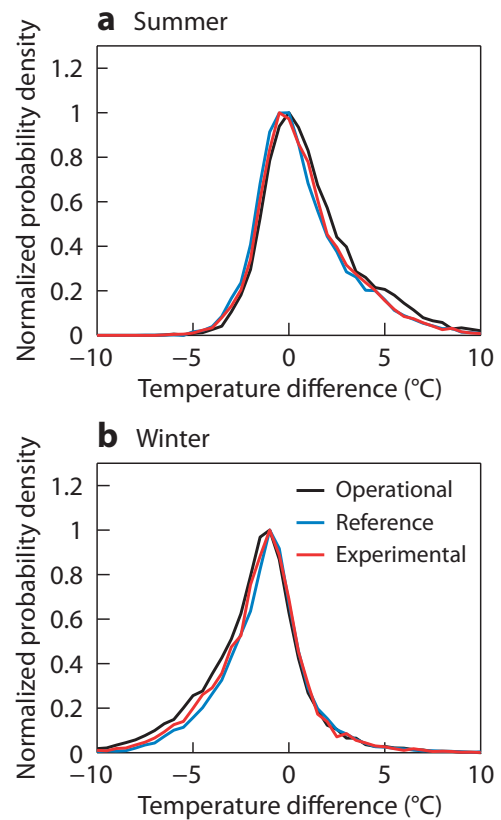


Figure 5 Probability density functions (normalized to the maximum value) of the difference between forecast and observed 2-metre temperature at European coastal stations for (a) June to September 2012 forecasts at 12 UTC, 36 hours into each daily forecast, and (b) December 2012 to March 2013 forecasts at 00 UTC, 24 hours into each daily forecast.

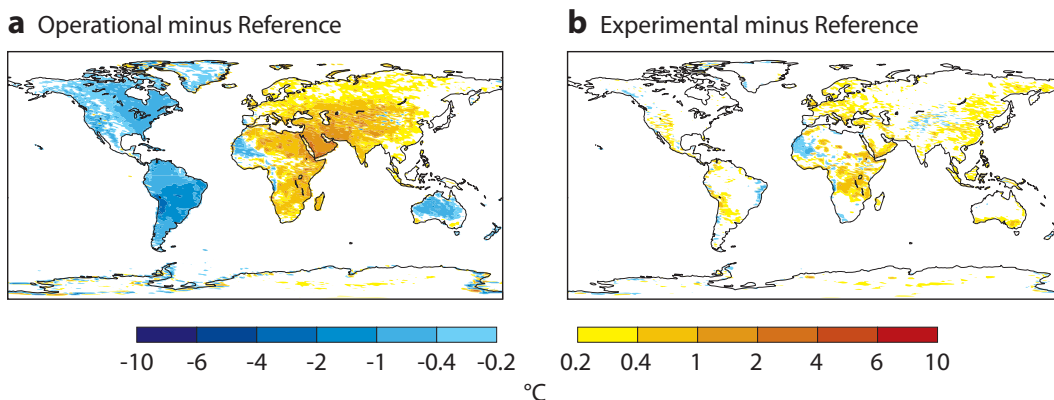
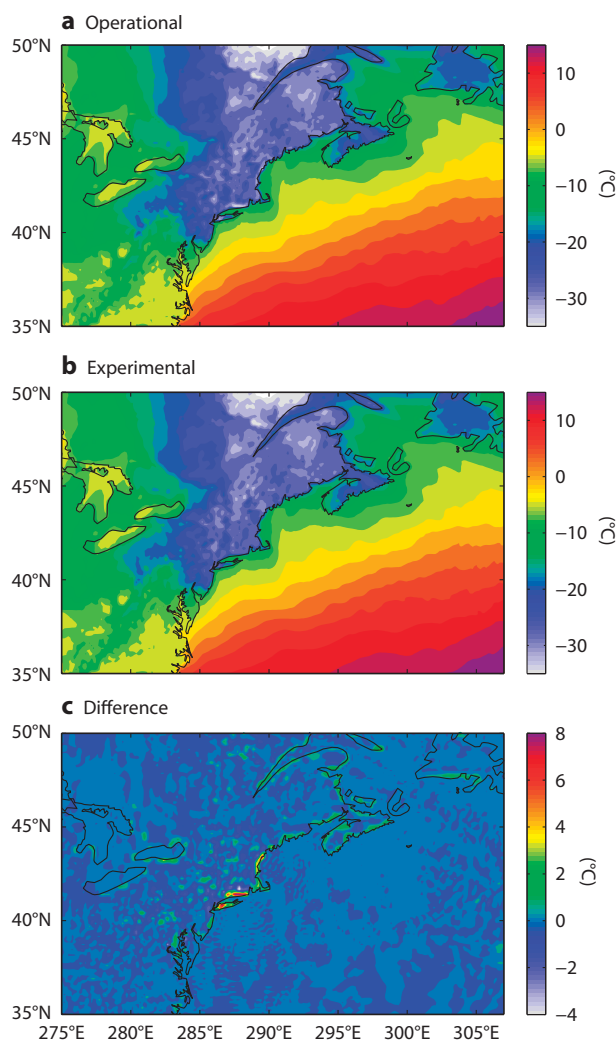


Figure 6 Difference in mean skin temperature (°C) at 12 UTC between reference forecasts with the radiation scheme run on the full model grid and at every time step (Reference) and (a) T639 forecasts with the radiation scheme run every 3 hours and on a reduced grid (Operational); and (b) the same as (a) but with approximate updates to the radiation fields on the full model grid and at every time step (Experimental). Forecasts were initialized from the analyses at 00 UTC every day for four months between 1 June and 30 September 2012, and values extracted 36 hours into each forecast.

This work highlights the importance of feedback from forecast users. ECMWF acted on it by advancing the science and designing, testing and implementing a solution. As illustrated in Figure 7, the results in this case indicate that there should be significant improvements to standard ECMWF forecast products, such as 2-metre temperature.

In terms of future developments, it should be recognised that the approximate updates described in this article are only needed because of the high computational cost of the current radiation scheme, and in particular the large number of spectral intervals it uses to represent gaseous absorption spectra. Reducing the accuracy of the spectral calculations has shown the potential to significantly improve the computational efficiency of the radiation scheme (Bozzo *et al.*, 2014). We plan to explore faster alternatives that retain as much of RRTMG's accuracy as possible, with the aim of developing a full scheme that is fast enough to be run much more frequently in space and time. Such a scheme would therefore be able to respond to the space–time variation not only in surface conditions but also in atmospheric state.

Figure 7 Two-metre temperature from (a) the Operational Scheme and (b) the Experimental Scheme, along with (c) the difference between the two. The data are for the same time and model resolution as shown in Figure 1, but for a larger area of the east coast of North America to put the large Long Island differences into context. Note that the exact locations of the largest errors in the Operational Scheme can change with both model resolution and model cycle (the latter if there is a change to the land-sea mask between cycles).



FURTHER READING

Bozzo, A., R. Pincus, I. Sandu & J.-J. Morcrette, 2014: Impact of a spectral sampling technique for radiation on ECMWF weather forecasts. *J. Adv. Modeling Earth Sys.*, **6**, 1288–1300. doi:10.1002/2014MS000386.

Hogan, R. J., and A. Bozzo, 2015: Mitigating errors in surface temperature forecasts using approximate radiation updates. *J. Adv. Modeling Earth Sys.*, **7**, 836–853. doi:10.1002/2015MS000455.

Manners, J., J.-C. Thelen, J. Petch, P. Hill & J.M. Edwards, 2009: Two fast radiative transfer methods to improve the temporal sampling of clouds in numerical weather prediction and climate models. *Q. J. R. Meteorol. Soc.*, **135**, 457–468.

Morcrette, J.-J., 2000: On the effects of the temporal and spatial sampling of radiation fields on the ECMWF forecasts and analyses. *Mon. Weath. Rev.*, **128**, 876–887.

GEOWOW project boosts access to Earth observation data

**ERVIN ZSOTER, RICHARD MLADEK,
FLORIAN PAPPENBERGER, DAVID RICHARDSON,
SUSANNA HOPSCH, JULIA KELLER,
RICHARD SWINBANK, PIERS BUCHANAN,
PHILIPPE ARBOGAST, BRUNO JOLY**

GEOWOW (GEOSS interoperability for Weather, Ocean and Water) was an EU-funded project that improved Earth Observation data discovery, accessibility and exploitability. It was coordinated by the European Space Agency (ESA). The project evolved the Global Earth Observation System of Systems (GEOSS) in the areas of Weather, Ocean Ecosystems and Water. Particular emphasis was placed on supporting multi-disciplinary approaches and interoperability across different domains, including the use of weather forecast ensembles in hydrological applications.

The weather community in GEOWOW, including the UK Met Office, Météo-France, the Karlsruhe Institute of Technology (KIT) and ECMWF (as work package leader), made major contributions to the success of GEOWOW, in particular by enhancing the content and use of the TIGGE forecast archive (Box A). Key achievements include:

- The creation of the TIGGE-LAM (limited-area model) archive for Europe to improve research relating to regional ensemble forecasts of high-impact weather
- The development of real-time products for high-impact weather and tropical cyclone track forecasts combining predictions from multiple TIGGE ensembles for forecasters as part of the WMO Severe Weather Forecasting Demonstration Project (SWFDP – Box B)
- A multidisciplinary showcase project to demonstrate the interoperable use of weather and water data in GEOWOW by generating discharge forecasts from TIGGE model inputs and providing them together with observations through a web service

In this article, we summarise the main outcomes of the weather-related work carried out during the three-year project, which ended in August 2014.

Improving data access

GEOWOW aimed to develop the existing infrastructure to improve data discovery, access and exploitation and to promote new documented datasets. To address these

AFFILIATIONS

Ervin Zsoter, Richard Mladek, Florian Pappenberger, David Richardson: ECMWF, UK

Susanna Hopsch, Julia Keller: Karlsruhe Institute of Technology, Germany

Richard Swinbank, Piers Buchanan: Met Office, UK

Philippe Arbogast, Bruno Joly: Météo France

THORPEX and TIGGE

A

THORPEX (The Observing System Research and Predictability Experiment) was a ten-year international research programme that was established in 2005 to accelerate improvements in the accuracy and utility of high-impact weather forecasts up to two weeks ahead. THORPEX was part of the World Weather Research Programme and is a key research component of the WMO's disaster risk reduction programme.

In the TIGGE (The International Grand Global Ensemble) project, which was started as part of THORPEX, ten of the leading global weather forecast centres provide regular global ensemble predictions to support research. The first global ensemble forecasts were archived in TIGGE at the end of 2006. Since then, TIGGE has become a focal point for a range of research projects, including research on ensemble forecasting, predictability and the development of products to improve the prediction of severe weather.

Although the THORPEX programme finished at the end of 2014, TIGGE will continue for a further 5 years, when its future will be reviewed. The database will continue to provide an invaluable resource for the science community, including by supporting WWRP THORPEX legacy projects such as the Polar Prediction Project (PPP), the High Impact Weather project (HIWeather) and the Sub-seasonal to Seasonal prediction project (S2S).

SWFDP

B

The World Meteorological Organization's Severe Weather Forecasting Demonstration Project (SWFDP) strengthens the capacity of National Meteorological and Hydrological Services in developing and least developed countries to deliver improved forecasts and warnings of severe weather to save lives, livelihoods and property.

The SWFDP is implemented as a number of Regional Projects, each using a 'Cascading Forecasting Process' (global to regional to national). The SWFDP was started in Southern Africa in 2006 and has now been rolled out in Eastern Africa, Southwest Pacific Islands and Southeast Asia. ECMWF and the Met Office both participate in the project as a global NWP centre, providing graphical products for the domains of each regional project as required for early warnings of high-impact and extreme weather events.

ECMWF contributes to three of the SWFDP Regional Projects. Participants in these projects can access the ECMWF forecast charts for their project using dedicated links (where a login is required).

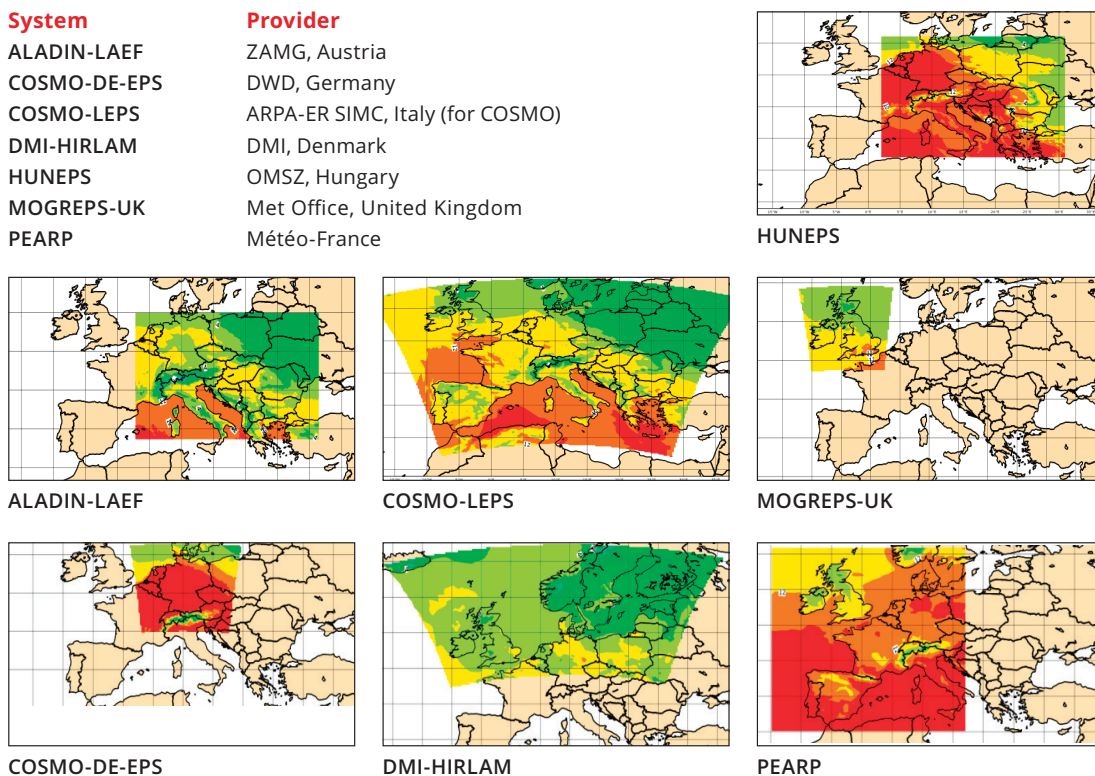


Figure 1 Examples of different 2-metre temperature forecasts produced by the seven LAM Ensemble Prediction Systems that contribute to the TIGGE-LAM archive.

requirements the weather work package made several improvements to the TIGGE archive. These include adding new datasets and experimenting with different ways of efficiently accessing the data.

TIGGE-LAM archive

During GEOWOW, the pre-existing global TIGGE archive was extended by adding a limited-area model (LAM) counterpart, the TIGGE-LAM archive for Europe. The new database is updated regularly with datasets from seven European LAM Ensemble Prediction Systems (EPS): ALADIN-LAEF, COSMO-DE-EPS, COSMO-LEPS, HUNEPS, MOGREPS-UK, PEARP, and DMI-HIRLAM (Figure 1).

The forecasts in the archive are produced with grid spacings of 2 to 12 km and provide detailed information for the short range, up to a few days ahead. Such forecasts complement the larger-scale information provided by the global data in the established TIGGE archive. Only high-priority surface parameters are archived (mean sea level pressure, 10-metre wind components, wind gust, 2-metre temperature and dew point temperature, total precipitation, CAPE and CIN convective indices, orography and land-sea mask). The TIGGE-LAM archive, which is hosted by ECMWF, was officially launched in March 2014.

TIGGE-LAM datasets are freely accessible using standard ECMWF tools (data portal, WEB-API). The data are available for scientific research (with a 48-hour delay) after a simple electronic registration process in the same way as the global data. Users have already started experimenting with

TIGGE-LAM data by comparing models (for example global vs regional) and improving the methodologies for the generation and application of regional ensemble forecasts.

Time-series archive prototype

In order to improve the accessibility of the TIGGE and TIGGE-LAM databases, a prototype of a time-series archive has been developed in GEOWOW. The aim is to provide an efficient way of accessing long time series (over several years) of forecast data at specific geographical locations. Currently the retrieval of forecast time series for any point requires accessing the series of whole fields (global or limited area), which is very inefficient and makes data access very slow.

Although the production version of the time-series archive has not been achieved, the experience gathered in this area constitutes a good basis for further development after GEOWOW.

The infrastructure developed in GEOWOW will now also be used by other research projects, such as UERRA (Uncertainties in Ensembles of Regional Re-analysis; project funded under EU Framework Programme 7) and S2S (Sub-seasonal to Seasonal prediction project; WWRP-WCRP joint research project).

Forecast product development

An important area where improvements in data exploitability could be demonstrated was the development of new TIGGE-based products for high-impact weather.

These products were developed in collaboration with users in developing and least developed countries and their value was demonstrated through the SWFDP framework.

Severe weather warning products

A survey of forecasters participating in the SWFDP confirmed that there is considerable interest in the use of ensemble-based products to support the forecasting of tropical cyclones, heavy precipitation and strong wind. The development and demonstration of ensemble forecast products in GEOWOW therefore focussed on these areas.

An important feature of the TIGGE archive is the two-day delay before the forecast data can be accessed. This comes from the status of TIGGE as a non-operational, non-commercial use research platform. The delay means that SWFDP forecasters are unable to use products calculated from the TIGGE data archive for real-time forecasting. To meet the needs of the SWFDP, a set of real-time products

has been developed in GEOWOW which will supply forecasters with plots showing the risks of strong winds and heavy rain up to 5 days ahead. Figure 2 shows an example of such a product.

The system is based on one originally developed, using TIGGE data, by Mio Matsueda of Tsukuba University and Oxford University (*Matsueda & Nakazawa, 2015*). The products use ensemble predictions from four TIGGE partners: the UK Met Office, ECMWF, NCEP (USA) and JMA (Japan). Three of the ensembles were already available to the Met Office, and agreement has been reached with JMA to supply predictions in real time for this project. This has enabled the warning products to be supplied to SWFDP regional projects, starting with the South Pacific Islands (Severe Weather Forecasting and Disaster Risk Reduction Demonstration Project, SWFDDP). It is planned to roll out the products to the Southern Africa and South-East Asia regions of the SWFDP in the near future.

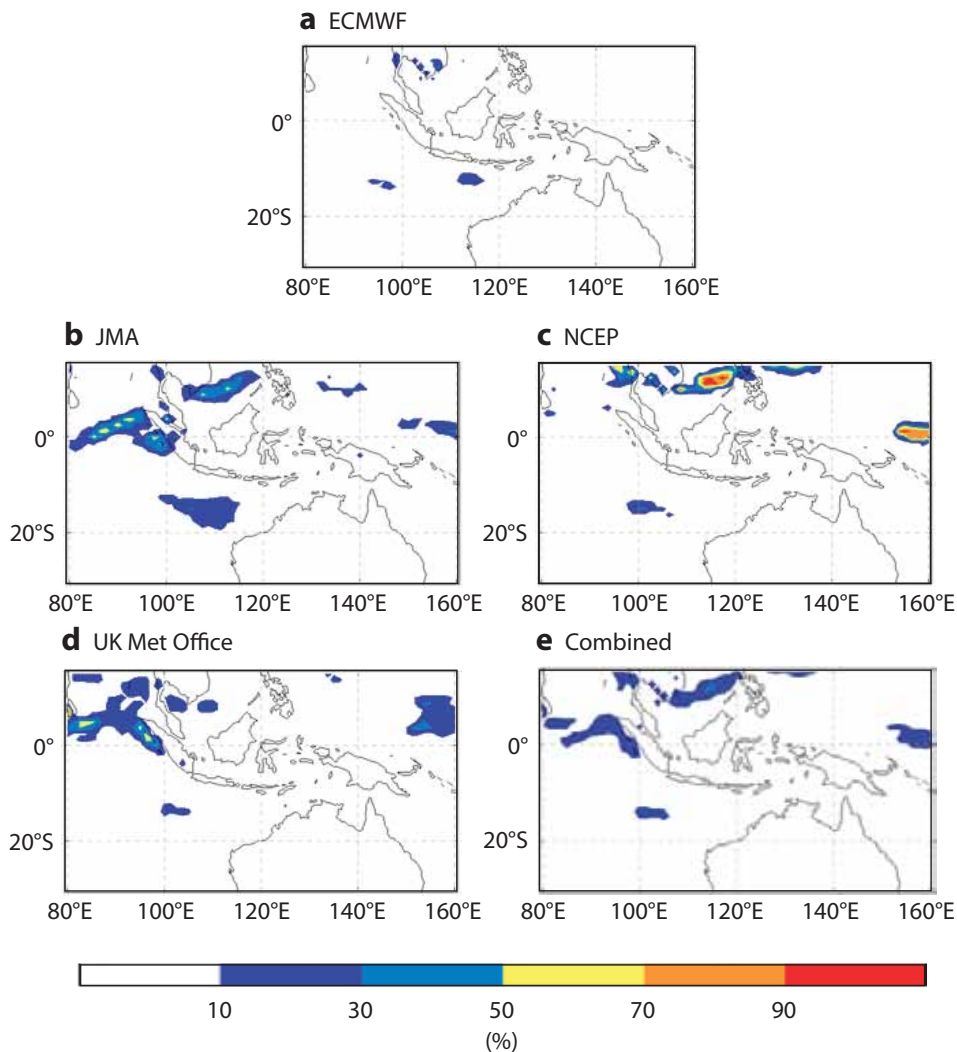


Figure 2 Example of product for warnings of heavy rainfall, showing ensemble forecasts produced by (a) ECMWF, (b) the Japan Meteorological Agency (JMA), (c) the US National Centers for Environmental Protection (NCEP), (d) the UK Met Office, and (e) a combination of those four forecasts. This example was produced for the SWFDDP showing the risk of the predicted rainfall exceeding the 99th percentile of the corresponding model climate rainfall in the 24 hours ending at 12 UTC on 10 August 2015, based on forecasts starting at 12 UTC on 5 August.

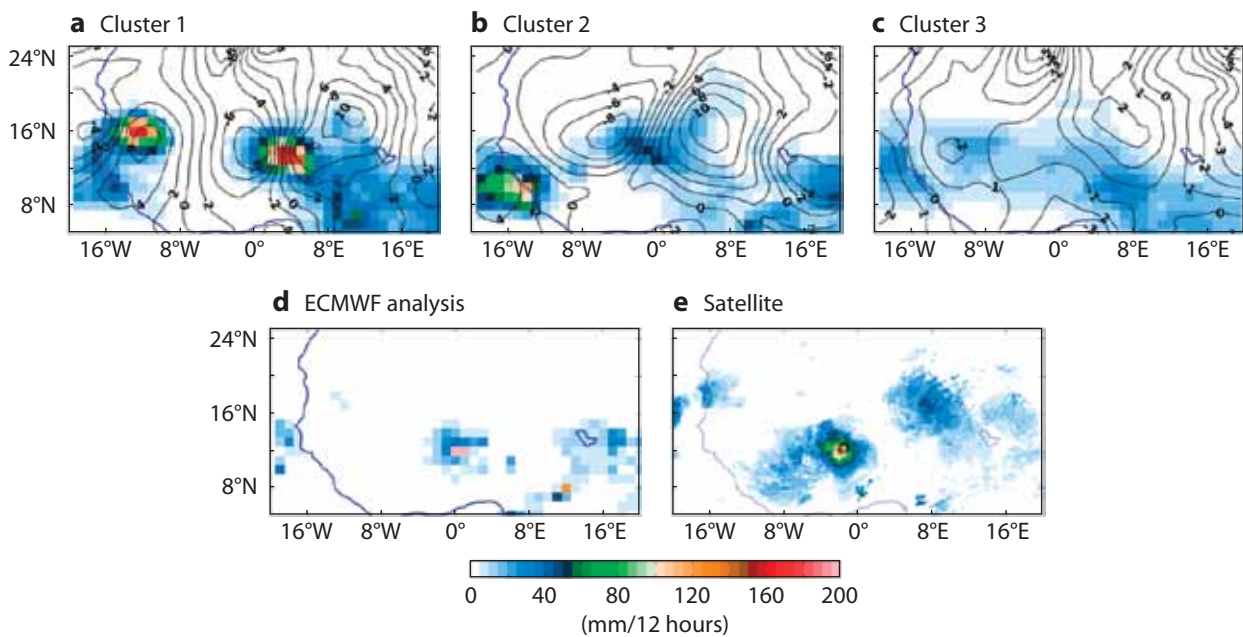


Figure 3 Example of weather scenarios obtained by EOF and fuzzy cluster analysis from the 96-hour TIGGE multi-ensemble forecast for a severe flooding event in Ouagadougou in Africa on 1 September 2009. Panels (a) to (c) show the cluster mean 12-hour precipitation (shading) and 700 hPa vertical velocity (contours) in the three derived clusters of ensemble members from a combination of TIGGE models, (d) shows the ECMWF 12-hour precipitation analysis (shortest-range forecast) and (e) the satellite-based 24-hour estimated precipitation.

Tropical cyclones

Additional products were also developed showing tropical cyclone track forecasts and strike probabilities combining predictions from multiple TIGGE ensembles. Tropical cyclone track forecast products have been available to the SWFDP forecasters for some time, and now the multi-centre versions are also being distributed. The three ensemble systems included are those of the Met Office, NCEP and ECMWF. It was found that there is significant benefit in terms of skill when the results of more than one ensemble are combined. In particular, the multi-model approach has the benefit of reducing the risk of observed storm tracks being outside the ensemble spread. The severe weather warning and tropical cyclone track products will be evaluated using regular SWFDP protocols.

High-impact weather forecasting

Within GEOWOW an exploratory study was carried out to demonstrate how TIGGE data can be used to aid forecasting of high-impact weather events from a multi-model perspective. The goal was to assess how much information can be extracted from the low-resolution ensembles to identify convective events, and to provide suggestions on how TIGGE data could be used for downscaling purposes. An analysis based on empirical orthogonal functions (EOF) followed by fuzzy clustering was used to identify distinct scenarios that are contained in the ensemble data.

In particular, case studies of high precipitation events in association with the passage of easterly waves over West Africa, the tropical cyclone Haruna (South Indian Ocean, impacting Madagascar), as well as several weather events affecting southern Africa were used to show how

the variability within ensembles can be used to gain information on possible ranges of intensity or evolution of weather features. An example is given in Figure 3.

There has also been a special collaboration with the South Africa Weather Service (SAWS) to explore the operational application of the clustering technique. A working product based on the EOF and fuzzy clustering technique was developed. Operational implementation at SAWS has started to provide guidance to forecasters in high-impact weather situations on the website for the southern Africa SWFDP domain (more details can be found in the June 2014 issue of the South African Society for Atmospheric Sciences (SASAS) newsletter).

Heavy precipitation events in Europe

In southern France, heavy precipitation events (HPEs) occur typically in the autumn, when highly local mechanisms can lead to dramatic rainfall. Due to the localised nature of these events it is often very difficult to achieve early detection and a correct magnitude estimate based on direct model outputs alone. During GEOWOW an experimental framework to predict HPEs was extended to TIGGE. An index of large-scale environments favouring HPEs was designed to estimate the probability of the most intense HPEs. The prediction of HPEs with sufficient lead time and sufficient certainty to mitigate the potential economic impact was implemented for two numerical weather prediction systems, the Météo France ensemble system PEARP and ECMWF's ensemble forecasts. An example of such a prediction is shown in Figure 4. Real-time analyses of HPE occurrence probability are provided to Météo France forecasters focussing on short lead times

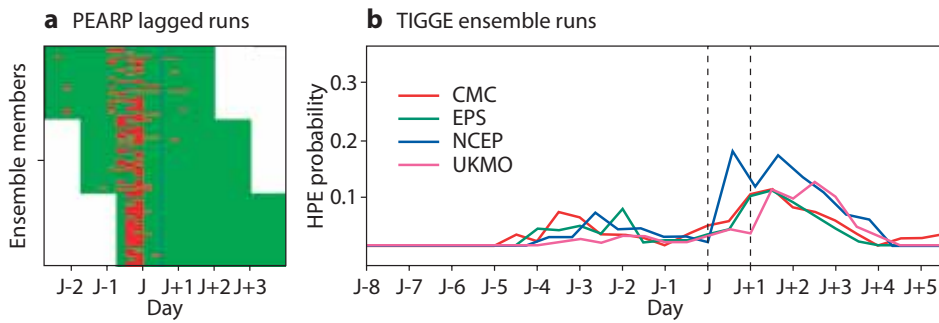


Figure 4 Example of forecast products for a heavy precipitation event (HPE) that occurred in southern France on 26 September 2012 showing (a) the HPE index computed from PEARP lagged daily runs from 18 UTC on 23 September 2012 to 18 UTC on 25 September 2012 (three runs), with each run represented by a third in the plot and each ensemble member shown as a line, with 6-hour periods that are favourable for an HPE highlighted in red, and (b) the probability of an HPE computed from four global medium-range ensemble systems (Canadian Meteorological Centre (CMC), ECMWF (EPS), US National Centers for Environmental Prediction (NCEP) and UK Met Office (UKMO)) from TIGGE (00 UTC runs on 18 September 2012). J denotes the HPE event and the days on the x-axis are referenced to this day.

(1 to 3 days). In addition to this, in GEOWOW extended experiments were performed to demonstrate the potential added value in providing early HPE warnings in the medium range by combining multiple ensemble systems. It could be shown that significant HPE occurrence probabilities could be extracted from TIGGE ensembles with lead times of up to 10 days.

Ensemble calibration and combination

The TIGGE archive contains forecast data from different models which have very different biases and error characteristics. The work in this area contributed to potential forecast error reduction by both calibrating forecast data to correct systematic errors and combining forecasts from multiple NWP models. The developments in this area were demonstrated mainly for heavy rain and strong winds.

New ensemble calibration scheme

A new ensemble calibration scheme has been developed in GEOWOW (Flowerdew, 2014). The scheme, which is

illustrated in Figure 5, aims to improve the local reliability of ensemble predictions – to minimise any systematic errors in the probabilities of weather outcomes derived from ensemble forecasts. This calibration scheme preserves the spatial, temporal and inter-parameter structure from the raw forecasts, which should be beneficial to downstream applications, such as hydrological models. The method was applied to recent, high-resolution TIGGE ensemble predictions for Europe for a range of parameters (precipitation, surface air temperature and dewpoint, surface wind speed and mean sea level pressure) over a two-year period. The results demonstrate that the calibration improved the skill of the forecasts as measured by probabilistic skill scores, including the Brier Skill Score. The results also confirmed that combining predictions from three skilful ensembles to form a multi-model grand ensemble gave superior results to the best single ensemble.

Calibration using reforecasts

A new reforecast dataset, covering autumn and winter periods of the last 32 years, has recently been developed

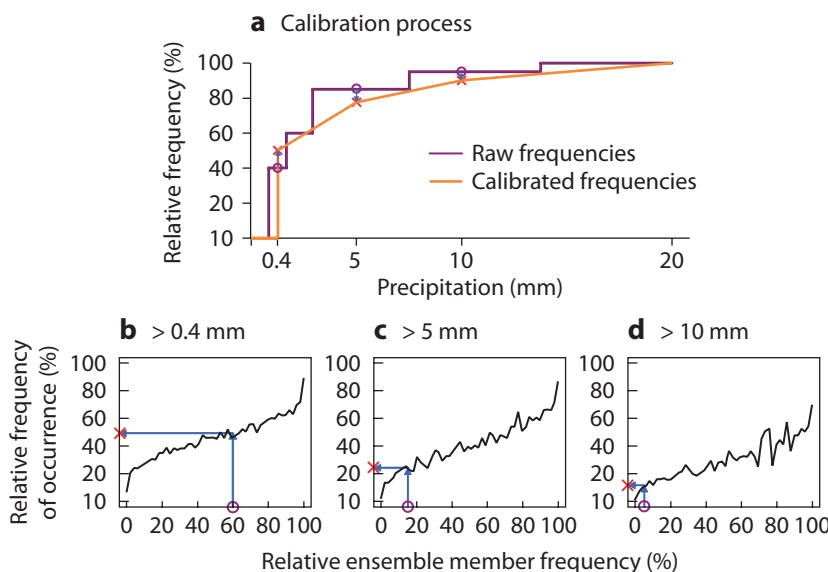


Figure 5 Schematic representation, adapted from Flowerdew (2014), of the new ensemble calibration method for a single grid point and a single forecast source, showing (a) the relative frequency of raw ensemble members (in %) which predict 12-hour precipitation amounts below different thresholds (including 0.4 mm, 5 mm and 10 mm) (purple line), and the relative frequency of cases in which past observed precipitation amounts really were below those thresholds for each of the raw ensemble member frequencies (orange line) determined from reliability diagrams (b) to (d) for the different thresholds. New, calibrated ensemble members can be obtained from the calibrated frequency distribution.

to document the model climate of the currently operational global ensemble forecast system of Météo France (called PEARP). This reforecast dataset will be used to calibrate PEARP forecasts. Preliminary results have shown that this calibration significantly improves the reliability and resolution of ensemble forecasts of 24-hour precipitation amount.

The reforecasts were also used in an analysis of major windstorms that have hit France over the last 30 years. Forecast indices such as the Extreme Forecast Index (EFI) and the Shift of Tails (SOT) index were computed for the most severe storms, using the reforecast-based model climate and historical ensemble forecasts produced by an ensemble forecasting system with a similar configuration to the current PEARP predictions (Figure 6).

It could be shown that the ensemble system together with its climatology has satisfactory predictive skill even for the most extreme events. It was also possible to determine optimal EFI and SOT thresholds which provide warnings for these extreme events while limiting the number of false alarms. This work related to the use of reforecast data is described in *Boisserie et al.* (2015).

Multidisciplinary use

One outcome of the GEOWOW project was to build bridges between different scientific disciplines. A key objective was the ability for users to exchange data, browse available resources and obtain relevant data through a common interface. To support these goals, different use cases have been initiated in GEOWOW where multidisciplinary and interoperable use of various data could be demonstrated.

Discharge modelling

The weather package has contributed to the demonstration of multidisciplinary use across different Societal Benefit Areas (SBAs) by developing one of the showcases in GEOWOW. This weather/water cross-domain activity was the ‘Modelling of discharge forecasts based on TIGGE ensemble inputs and validation with observations from GRDC’ (Global Runoff Data Centre). Weather partners worked on the hydrological modelling while water-related partners worked on the interoperable use of hydrological data standards and observations. Other project participants contributed to the technical parts of the scenario.

The discharge modelling work was based on the HTESSSEL land-surface model used operationally at ECMWF. The offline version of HTESSSEL was extended to accommodate ensemble forecast runs from models in the TIGGE archive, also using ECMWF climate and initial conditions. The hydrological model output runoff was coupled to the CaMa-Flood river routing scheme to provide river discharge data for about 400 catchments globally.

The production of the TIGGE-based hydrological discharge forecasts covered the period 2009–2011. The analysis of this discharge dataset was performed using GRDC observations and the skill properties of the hydrological predictions from the different TIGGE models were compared with post-processed and multi-model combinations.

To support the data provision, a dedicated data server with PostgreSQL database support and a Sensor Observation Service (SOS) server to service the data were installed at

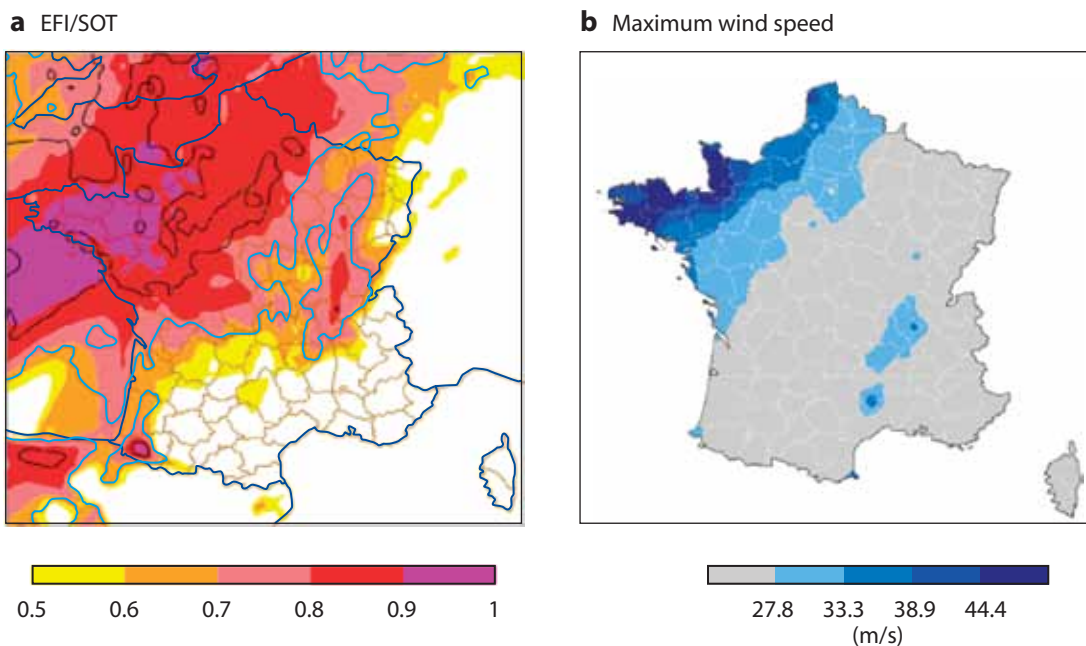


Figure 6 Maps of France showing (a) EFI and SOT indices based on a 102-hour lead time forecast of 24-hour maximum wind gust and (b) observed 24-hour maximum wind speed, both valid for 16 October 1987. The EFI is shown by shading (from 0.5 to 1.0) while the SOT index (with $p = 0.9$) is plotted using contours (light blue line indicates a value of -0.5, grey line indicates a value of 0). The observed 24-hour maximum wind speed values are shown in shades of blue for values greater than 27.8 m/s.

ECMWF. The modelled TIGGE discharge forecasts were uploaded onto the data server for a set of GRDC stations covering mainly the GEOSS DataCORE stations. In close collaboration with other project partners, the data server was connected to GEOSS through the GEOSS Discovery and Access Broker (DAB) and the visualisation of the modelled data (in parallel to the GRDC observations) was performed by the 52°North SOS client application (Figure 7).

This system using the modelled discharge data and the GRDC observations was a demonstration of the aim of the project to improve the discovery and accessibility of TIGGE data through the GEOSS framework. The system was presented at relevant events, such as the GEO-X Plenary and Ministerial Summit, and also in the media.

Conclusions to GEOWOW

The weather community has contributed to the success of GEOWOW through numerous improvements in ensemble data access and forecasting methodology, and also through the development of forecast products for high-impact weather and multidisciplinary applications across different GEO SBAs. The main achievements include:

- The creation of the TIGGE-LAM archive for Europe to improve research relating to regional ensemble forecasts of high-impact weather
- TIGGE time-series prototype development to improve access to long ensemble forecast time series
- The development of real-time products for high-impact weather and tropical cyclone track forecasts combining predictions from multiple TIGGE ensembles for SWFDP forecasters
- An exploratory study to demonstrate how TIGGE data can be used to support the forecasting of high-impact



Figure 7 Example of the web service run by 52° North showing some randomly selected ensemble members from TIGGE discharge forecasts for a point on the Susquehanna River in Maryland, US.

weather events from multi-models based on EOF analysis and subsequent fuzzy clustering

- The extension of the heavy precipitation event prediction products for southern France to TIGGE models
- A new ensemble calibration scheme that preserves the spatial, temporal and inter-variable structure from the raw forecasts
- A new 32-year reforecast dataset to document the model climate of the ensemble system of PEARP. This reforecast dataset is used to calibrate the PEARP forecasts and also to study weather extremes in France
- A multidisciplinary showcase project to demonstrate the interoperable use of weather and water data in GEOWOW by generating discharge forecasts from TIGGE model inputs and providing them together with observations through a web service
- The registration of TIGGE and TIGGE-LAM archives in GEOSS and their discoverability through the GEOSS portal

GEOWOW's contribution to the development of GEOSS was substantial. This project demonstrated how users can be supported in discovering, accessing and exploiting data from diverse sources through GEOSS. In this project, weather, water and ocean ecosystems were included, but in principle the same approach should work across a larger range of GEO Societal Benefit Areas and Earth Observations.

Useful links

TIGGE-LAM data portal

<http://apps.ecmwf.int/datasets/data/tigge-lam>

More information about the TIGGE/TIGGE-LAM project

www.ecmwf.int/en/research/projects/tigge

GEOWOW

www.geowow.eu

GEOSS portal

www.geoportal.org

SWFDP

www.wmo.int/pages/prog/www/swfdp/

S2S (Sub-seasonal to Seasonal prediction project)

<https://software.ecmwf.int/wiki/display/S2S>

UERRA (Uncertainties in Ensembles of Regional ReAnalysis)

<https://software.ecmwf.int/wiki/display/UER>

SASAS newsletter

<http://web.csag.uct.ac.za/sasas/index.php/newsletter>

FURTHER READING

Boisserie, M., L. Descamps, & P. Arbogast, 2015: Calibrated forecasts of extreme windstorms using Extreme Forecast Index (EFI) and Shift Of Tails (SOT). *Wea. Forecasting*. doi:10.1175/WAF-D-15-0027.1, in press.

Flowerdew, J., 2014: Calibrating ensemble reliability whilst preserving spatial structure. *Tellus A*, **66**, 22662, doi:10.3402/tellusa.v66.22662.

Matsueda, M. & T. Nakazawa, 2015: Early warning products for severe weather events derived from operational medium-range ensemble forecasts. *Met. Apps.*, **22**, 213–222, doi: 10.1002/met.1444.

ECMWF Calendar 2016

Jan 18–21	Computer User Training Course: ecFlow
Jan 25–29	Computer User Training Course: HPC Facility Cray XC30
Feb 1–5	Training Course for Trainers, Training Champions: Use and Interpretation of ECMWF Products
Feb 8–12	Training Course: Use and Interpretation of ECMWF Products
Feb 22–26	Computer User Training Course: Introduction for New Users/MARS
Mar 7–11	NWP Training Course: Data Assimilation
Mar 14–18	EUMETSAT/ECMWF NWP SAF Training Course: Assimilation of Satellite Data
Apr 11–14	Workshop on Model Errors
Apr 11–15	NWP Training Course: Advanced Numerical Methods for Earth System Modelling
Apr 18	EUMETSAT Data Policy Group – Switzerland
Apr 19	Advisory Committee for Data Policy – Switzerland
Apr 20	ECOMET Working Group
Apr 25	Policy Advisory Committee
Apr 26–27	Finance Committee
May 9–13	NWP Training Course: Predictability and Ocean–Atmosphere Ensemble Forecasting

May 16–20	NWP Training Course: Parametrization of Subgrid Physical Processes
May 17–18	Security Representatives' Meeting
May 18–20	Computing Representatives' Meeting
Jun 6–8	Using ECMWF's Forecasts (UEF)
Jun 20–24	Workshop on Assimilation of High-Resolution Satellite Data
Jun 30 – Jul 1	Council
Sep 5–9	Annual Seminar
Sep 12–16	Workshop on Drag Processes and Links to Large-Scale Circulation
Sep 26–30	Computer User Training Course: Data Analysis and Visualisation using Metview
Oct 3–6	Workshop on Numerics
Oct 24–28	Workshop on High-Performance Computing in Meteorology
Nov 7–11	Workshop on Tropical Modelling and Assimilation
Nov 14–17	EUMETSAT Satellite Application Facility on Climate Monitoring (CM SAF) Workshop
Nov 29–30	ECOMET GA and EUMETNET Assembly
Dec 1–2	Council

Contact information

ECMWF, Shinfield Park, Reading, Berkshire RG2 9AX, UK

Telephone National 0118 949 9000

Telephone International +44 118 949 9000

Fax +44 118 986 9450

ECMWF's public website <http://www.ecmwf.int/>

E-mail: The e-mail address of an individual at the Centre is firstinitial.lastname@ecmwf.int. For double-barrelled names use a hyphen (e.g. j-n.name-name@ecmwf.int).

Problems, queries and advice	Contact
General problems, fault reporting, web access and service queries	calldesk@ecmwf.int
Advice on the usage of computing and archiving services	advisory@ecmwf.int
Queries regarding access to data	data.services@ecmwf.int
Queries regarding the installation of ECMWF software packages	software.support@ecmwf.int
Queries or feedback regarding the forecast products	forecast_user@ecmwf.int

ECMWF publications

(see <http://www.ecmwf.int/en/research/publications>)

Technical Memoranda

- 761 **MacLeod, D.A., H.L. Cloke, F. Pappenberger & A. Weisheimer:** Improved seasonal prediction of the hot summer of 2003 over Europe through better representation of uncertainty in the land. *August 2015*
- 758 **Hogan, R. & S. Hirahara:** Effect of solar zenith angle specification on mean shortwave fluxes and stratospheric temperatures. *August 2015*

ERA Report Series

- 21 **Kobayashi, S., P. Poli & V. John:** CM-SAF Visiting Scientist Activity CM_VS14_01 Report: Characterisation of SSM/T-2 radiances using ERA-Interim and other reanalyses. *August 2015*

- 19 **Poli, P., C. Peubey, K. Fennig, M. Schroeder, R. Roebelling & A. Geer:** Pre-assimilation feedback on a Fundamental Climate Data Record of brightness temperatures from Special Sensor Microwave Imagers: A step towards MIPs4Obs? *August 2015*

EUMETSAT/ECMWF Fellowship Programme Research Reports

- 39 **Lawrence, H., N. Bormann & S. English:** Scene-dependent observation errors for the assimilation of AMSU-A. *August 2015*
- 38 **Baordo, F. & A.J. Geer:** All-sky assimilation of SSMIS humidity sounding channels over land within the ECMWF system. *July 2015*

	No.	Date	Page		No.	Date	Page
Describing ECMWF's forecasts and forecasting system	133	Autumn 2012	11	Improving the representation of stable boundary layers	138	Winter 2013/14	24
COMPUTING				Interactive lakes in the Integrated Forecasting System	137	Autumn 2013	30
Supercomputing at ECMWF	143	Spring 2015	32	Effective spectral resolution of ECMWF atmospheric forecast models	137	Autumn 2013	19
SAPP: a new scalable acquisition and pre-processing system at ECMWF	140	Summer 2014	37	Breakthrough in forecasting equilibrium and non-equilibrium convection	136	Summer 2013	15
Metview's new user interface	140	Summer 2014	42	Convection and waves on small planets and the real Earth	135	Spring 2013	14
GPU based interactive 3D visualization of ECMWF ensemble forecasts	138	Winter 2013/14	34	Global, non-hydrostatic, convection-permitting, medium-range forecasts: progress and challenges	133	Autumn 2012	17
RMDCN – Next Generation	134	Winter 2012/13	38				
A new trajectory interface in Metview 4	131	Spring 2012	31				
METEOROLOGY				PROBABILISTIC FORECASTING & MARINE ASPECTS			
OBSERVATIONS & ASSIMILATION				Have ECMWF monthly forecasts been improving?	138	Winter 2013/14	18
GEOVOW project boosts access to Earth observation data	145	Autumn 2015	35	Closer together: coupling the wave and ocean models	135	Spring 2013	6
CERA: A coupled data assimilation system for climate reanalysis	144	Summer 2015	15	20 years of ensemble prediction at ECMWF	134	Winter 2012/13	16
Promising results in hybrid data assimilation tests	144	Summer 2015	33	Representing model uncertainty: stochastic parametrizations at ECMWF	129	Autumn 2011	19
Snow data assimilation at ECMWF	143	Spring 2015	26	METEOROLOGICAL APPLICATIONS & STUDIES			
Assimilation of cloud radar and lidar observations towards EarthCARE	142	Winter 2014/15	17	Improvements in IFS forecasts of heavy precipitation	144	Summer 2015	21
The direct assimilation of principal components of IASI spectra	142	Winter 2014/15	23	New EFI parameters for forecasting severe convection	144	Summer 2015	27
Automatic checking of observations at ECMWF	140	Summer 2014	21	The skill of ECMWF cloudiness forecasts	143	Spring 2015	14
All-sky assimilation of microwave humidity sounders	140	Summer 2014	25	Calibration of ECMWF forecasts	142	Winter 2014/15	12
Climate reanalysis	139	Spring 2014	15	Twenty-five years of IFS/ARPEGE	141	Autumn 2014	22
Ten years of ENVISAT data at ECMWF	138	Winter 2013/14	13	Potential to use seasonal climate forecasts to plan malaria intervention strategies in Africa	140	Summer 2014	15
Impact of the Metop satellites in the ECMWF system	137	Autumn 2013	9	Predictability of the cold drops based on ECMWF's forecasts over Europe	140	Summer 2014	32
Ocean Reanalyses Intercomparison Project (ORA-IP)	137	Autumn 2013	11	Windstorms in northwest Europe in late 2013	139	Spring 2014	22
The expected NWP impact of Aeolus wind observations	137	Autumn 2013	23	Statistical evaluation of ECMWF extreme wind forecasts	139	Spring 2014	29
Winds of change in the use of Atmospheric Motion Vectors in the ECMWF system	136	Summer 2013	23	Flow-dependent verification of the ECMWF ensemble over the Euro-Atlantic sector	139	Spring 2014	34
New microwave and infrared data from the S-NPP satellite	136	Summer 2013	28	iCOLT – Seasonal forecasts of crop irrigation needs at ARPA-SIMC	138	Winter 2013/14	30
Scaling of GNSS radio occultation impact with observation number using an ensemble of data assimilations	135	Spring 2013	20	Forecast performance 2013	137	Autumn 2013	13
ECMWF soil moisture validation activities	133	Autumn 2012	23	An evaluation of recent performance of ECMWF's forecasts	137	Autumn 2013	15
Forecast sensitivity to observation error variance	133	Autumn 2012	30	Cold spell prediction beyond a week: extreme snowfall events in February 2012 in Italy	136	Summer 2013	31
FORECAST MODEL				The new MACC-II CO2 forecast	135	Spring 2013	8
An all-scale, finite-volume module for the IFS	145	Autumn 2015	24	Forecast performance 2012	134	Winter 2012/13	11
Reducing surface temperature errors at coastlines	145	Autumn 2015	30	Teaching with OpenIFS at Stockholm University: leading the learning experience	134	Winter 2012/13	12
Atmospheric composition in ECMWF's Integrated Forecasting System	143	Spring 2015	20	Uncertainty in tropical winds	134	Winter 2012/13	33
Towards predicting high-impact freezing rain events	141	Autumn 2014	15	Monitoring and forecasting the 2010-11 drought in the Horn of Africa	131	Spring 2012	9
Improving ECMWF forecasts of sudden stratospheric warmings	141	Autumn 2014	30	Characteristics of occasional poor medium-range forecasts for Europe	131	Spring 2012	11
				A case study of occasional poor medium-range forecasts for Europe	131	Spring 2012	16
				The European Flood Awareness System (EFAS) at ECMWF: towards operational implementation	131	Spring 2012	25



Newsletter | Number 145 – Autumn 2015
European Centre for Medium-Range Weather Forecasts

www.ecmwf.int