

Contents

Editorial 1

News

Changes to the operational forecasting system. 2

New items on the ECMWF web site 2

ECMWF Education and Training Programme for 2007... 3

Lennart Bengtsson receives the prestigious IMO Prize... 3

ECMWF/GEO Workshop on Atmospheric Reanalysis,
ECMWF, 19–22 June 2006 4

The new IBM Phase 4 HPC facility 5

Third International Workshop on Verification Methods .. 7

CTBTO: “Synergies with Science” 7

Summary of ECMWF

Forecasts Product Users’ Meeting, June 2006. 8

NCEP uses ECMWF’s methodology for

humidity analysis and background error generation 9

Annual Meetings of the

European Meteorological Society. 9

Meteorology

Analysis & forecast impact of humidity observations. . 11

Hindcasts of historic storms with the DWD models

GME, LMQ and LMK using ERA-40 reanalyses 16

Hurricane Jim over New Caledonia: a remarkable

numerical prediction of its genesis and track 21

Ice supersaturation in

ECMWF’s Integrated Forecast System. 26

Computing

New features of the Phase 4 HPC facility 32

General

ECMWF Calendar 2006/2007 37

ECMWF publications. 37

Index of past newsletter articles. 38

Useful names & telephone numbers

within ECMWF 40

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, Berkshire RG2 9AX, UK

Fax:+44 118 986 9450

Telephone: National0118 949 9000

International+44 118 949 9000

ECMWF Web site<http://www.ecmwf.int>

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 The ECMWF Newsletter is not peer-reviewed.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: Bob Riddaway

Typesetting and Graphics: Rob Hine

Front cover

Mean analysis differences of total-column water vapour between experiments using conventional observations and satellite radiances and those with all humidity observations withheld. See Andersson et al. on page 11.

Editorial**Monthly Forecasting**

If at some point in the future we look back on current developments at ECMWF, we will potentially consider monthly forecasting to be one of the most successful. This is for two main reasons: it is a breakthrough which was considered impossible only a few years ago, and it addresses requirements which have become very important at the time when the forecasts became available.

The story started in 2001 when two questions were raised.

◆ Given the progress in medium-range forecast, is there any useful skill beyond day 10?

◆ How can we fill the gap between the medium-range forecasts (up to 10 days) and the seasonal forecasts (from one month)?

The answers to these questions were positive.

◆ Yes, there is some skill up to day 15 in some cases.

◆ Yes, we might be able to develop a useful monthly forecasting system by coupling the most recent version of the atmospheric model (as used for medium-range forecasts) with an ocean model (as in seasonal forecasts).

On that basis an experimental system was developed by the seasonal forecasting team, and it was then tested for two years. It was run every other week and showed interesting skill, in particular indicating two weeks in advance several cold spells in winter.

Then this experimental system benefited from two specific events.

◆ For some time the Met Office had been running a monthly forecasting system based on a low-resolution atmospheric model. Because of the constraints imposed by the move to Exeter and considering the results of the ECMWF experimental system, the Met Office decided to stop production of its own monthly forecasts. As a result the ECMWF system immediately benefited from being able to build upon the Met Office’s user requirements and evaluation.

◆ The heat wave which dramatically hit Europe in summer 2003 drew the attention of several Meteorological Services to the potential benefits of the ECMWF monthly forecasts for such events.

The monthly forecasting system became operational in summer 2004. Dissemination of products started in March 2005. A description of the system was provided in the spring 2004 issue of the *ECMWF Newsletter* (No. 100). Major developments are still necessary. In particular it is necessary to develop specific products allowing forecasters to extract information from this new type of forecasts: several are being tested. Next year the medium-range and monthly ensemble forecasting systems will be merged in order to produce both medium-range and monthly forecasts using a single, unified system, based on coupled atmosphere and ocean models from day 10. Also the frequency of the monthly forecasts will be reviewed, based on user requirements.

In the end extending the medium-range forecasts to a month ahead was made possible by the development of the seasonal forecasts which itself was built upon the medium-range forecasting system. Such positive feedbacks are important elements to consider when initiating, and later evaluating, new developments.

Dominique Marbouty

Changes to the operational forecasting system

David Richardson

Implementation of Cy31r1

A new cycle of the ECMWF system, Cy31r1, was introduced on 12 September 2006. This version includes the following changes:

- ◆ Revisions and changes to the cloud scheme including treatment of ice supersaturation and new numerics.
- ◆ Implicit computation of convective transports.

- ◆ Introduction of turbulent orographic form drag scheme and revision to sub-grid scale orographic drag scheme.
- ◆ Gust fix for orography and stochastic physics.
- ◆ Reduction of ocean surface relative humidity from 100% to 98% (due to salinity effects).
- ◆ Revised assimilation of rain-affected radiances.
- ◆ Variational bias correction of satellite radiances.
- ◆ Thinning of low level AMDAR data (mainly affects Japanese AMDAR network).

Technical changes in preparation for the extension of the EPS to day 15 at reduced resolution (VAREPS) are also included.

New items on the ECMWF web site

Andy Brady

Workshop on “Parametrization of clouds in large-scale models”

The Workshop on “*Parametrization of clouds in large-scale models*” will be held from 13 to 15 November 2006. The representation of clouds in large-scale models is still a challenging problem and further development of the cloud scheme is central to the Centre’s plans. The purpose of the workshop is to review the most recent developments in this area of research and to explore new ideas.

www.ecmwf.int/newsevents/meetings/workshops/2006/parametrization_clouds/

ECMWF Global Data Monitoring Report

The ECMWF global data monitoring report is a monthly publication intended to give an overview of the availability and quality of observations from the Global Observing System within the World Weather Watch of the World Meteorological Organization. It should be recognised that the statistics given in this report refer to data as received at ECMWF in time for the appropriate analysis. The annex of the report gives further explanations of the methods applied to compile the statistics and on the reference used to establish the quality of observations.

www.ecmwf.int/products/forecasts/monitoring/mmr/

Phase 4 of the High Performance Computing Facility (HPCF)

The first of the new HPCF Phase 4 clusters (hpce) has passed its operational acceptance test and is used by the Operational Suite to produce the daily forecasts. A full service is now being provided on hpce and all users are invited to start using this machine. There are some major architectural changes between the Phase 3 and Phase 4 systems. Information about Phase 4 can be found in the two articles in this edition of the ECMWF Newsletter by Neil Storer as well as on the ECMWF website.

www.ecmwf.int/services/computing/hpcf/

New Release of the MARS Archive User Guide

MARS is ECMWF’s Meteorological Archival and Retrieval System. This documentation describes facilities for the retrieval of data from ECMWF’s operational and other archives, including the on-line Fields Data Base (FDB).

www.ecmwf.int/publications/manuals/mars/

Forecast Products Users’ Meeting

ECMWF organizes annually a meeting of users of its medium-range and extended-range products. The last meeting was held on June 2006. A report of this meeting by David Richardson appears as a news item in this edition of the ECMWF Newsletter. The presentations are now available from the ECMWF website.

www.ecmwf.int/newsevents/meetings/forecast_products_user/Presentations2006/

ECMWF Calendar 2007

The ECMWF 2007 Calendar of events is now available and will be updated with new events as they are planned. This calendar is also duplicated in the ECMWF Newsletter.

www.ecmwf.int/newsevents/calendar/

Third International Workshop on Verification Methods

The “*Third International Workshop on Verification Methods*” will be held at ECMWF from 29 to 2 February 2007. The workshop will include both a tutorial session and a programme of talks on new developments in verification methodologies. The news item in this edition of the ECMWF Newsletter by Anna Ghelli describes the workshop.

www.ecmwf.int/newsevents/meetings/workshops/2007/jwgv/

Web Search

The ECMWF Web Site Search software has been changed. The Search is now quicker than before, more accurate, covers more areas of the web site and has a better feature set. The Web Site Search can be found at the top right of every ECMWF web page.

www.ecmwf.int/publications/cms/get/ecmwfnews/153

ECMWF Education and Training Programme for 2007

Umberto Modigliani,
Renate Hagedorn, David Richardson

ECMWF has an extensive education and training programme to assist Member States and Co-operating States in the training of scientists in numerical weather forecasting, and in making use of the ECMWF computer facilities. The training courses consist of modules that can be attended separately. A student may decide to attend different modules in different years.

Use of Computing Facilities

The objective of this course is to introduce users of ECMWF's computing and archiving systems to the Centre's facilities, and to explain how to use them. The course is divided into five modules.

COM SMS	8 – 9 Feb	Introduction to SMS/XCdp
COM INTRO	19 – 23 Feb	Introduction for new users/MARS
COM MAG	26 – 27 Feb	MAGICS
COM MV	28 Feb – 2 Mar	METVIEW
COM HPCF	5 – 9 Mar	Use of supercomputing resources

Each module will consist of some lectures and some practical sessions. All the lectures will be given in English. A workbook will be provided for each module, together with basic manuals and other relevant documentation. Information about the course can be found at:

www.ecmwf.int/newsevents/training/2007/computer

Students attending any part of the course should:

- ◆ Have experience of a computer system elsewhere.
- ◆ Be familiar with ANSI Fortran 77 or 90.
- ◆ Know basic UNIX commands.
- ◆ Be able to use an editor (e.g. *vi*).

Numerical Weather Prediction

This course aims to provide advanced training in the field of numerical weather forecasting. The course is divided into four modules which are geared towards research aspects of numerical weather prediction. Preliminary dates for the modules have been agreed.

Dominique Marbouty

At a ceremony in Norrköping on 2 October Prof. Lennart Bengtsson was awarded the 51st IMO Prize by WMO for his pioneering research in numerical weather prediction (NWP). The IMO Prize is given in recognition of outstanding contributions to the advancement of meteorology and hydrology, as well as encouraging scientists active in those fields.

MET NM	19 – 28 Mar	Numerical methods and adiabatic formulation of models
MET PR	16 – 24 Apr	Predictability, diagnostics and seasonal forecasting
MET DA	25 Apr – 4 May	Data assimilation and use of satellite data
MET PA	8 – 18 May	Parametrization of diabatic processes

All the lectures will be given in English. Information about the course can be found at:

www.ecmwf.int/newsevents/training/2007/NWP

Students attending the course should have:

- ◆ A good meteorological and mathematical background.
- ◆ Familiarity with the contents of standard meteorological and mathematical textbooks.

Use and Interpretation of ECMWF Products

The course is designed to assist Member States and Co-operating States in advanced training on the operational aspects of the ECMWF forecasting system. The course will be given twice in 2007, plus an additional module specifically for WMO Member States.

MET OP-I	12 – 16 Mar	Use and interpretation of ECMWF products
MET OP-II	4 – 8 Jun	Use and interpretation of ECMWF products
MET OP-WMO	15 – 19 Oct	Use and interpretation of ECMWF products for WMO Members

All the lectures will be given in English. Information about the course can be found at:

www.ecmwf.int/newsevents/training/2007/Products

Students attending the course should have:

- ◆ A good meteorological background.
- ◆ Practical experience in weather forecasting.

Annual Seminar and Workshops

In addition to the training courses, there will be an Annual Seminar and a number of workshops. Information about these events will be available later in the year.

Lennart Bengtsson receives the prestigious IMO Prize

Lennart has contributed to the development of meteorological science and international cooperation for more than 40 years. As well as making major contributions to the development of NWP his research interests cover a wide range of other topics including climate modelling and prediction, data assimilation and climate data studies. It is recognised that Lennart's tremendous enthusiasm has been a source of inspiration for his colleagues and collaborators worldwide.



Prof. Lennart Bengtsson

As a member of the interim planning staff Lennart was involved in the establishment of ECMWF and in 1976 he became Deputy Director and Head of Research. Under his leadership the first ECMWF model was developed, which became operational on 1 August 1979. In 1982 Lennart was appointed as Director of ECMWF. During his period as Director a new operational model based on the spectral technique was introduced in April 1983. This approach has formed the basis of subsequent operational models. In addition, during Lennart's period of office the Centre expanded its activities and moved into new areas such as ensemble forecasting, wave prediction and seasonal forecasting. More information about the contribution made by Lennart to the development of ECMWF can be found in the book "*Medium-Range Weather Prediction: The European Approach*" published by Cambridge University Press. This book describes the history of ECMWF and how it became acknowledged as the world leader in its field.

After leaving ECMWF in 1990 Lennart became Director of the Max Planck Institute for Meteorology in Hamburg and he is now a Senior Research Fellow at the Environmental Systems Science Centre in the University of Reading.

I am delighted that Lennart's contribution to the development of European meteorology, and ECMWF in particular, has been recognised by the award of the prestigious IMO Prize. I join his many friends and colleagues from around the world in offering my congratulations.

ECMWF/GEO Workshop on Atmospheric Reanalysis, ECMWF, 19–22 June 2006

Sakari Uppala, Adrian Simmons, Dick Dee

Over the past decade, reanalyses have become established as an important and widely utilised resource for the study of atmospheric and oceanic processes and predictability. They are also used in a wide range of applications that require a record of the state of either the atmosphere or its underlying land and ocean surfaces. Whilst high-resolution operational forecasting systems continue to provide what currently are the best possible analyses of recent conditions, the rapid pace of improvements in data assimilation and increases in computer power enables an up-to-date reanalysis system to provide products for all but the last few years that are superior to those available from the archives of past operational products. Moreover, the reanalysis products, unlike their operational counterparts, do not suffer from inhomogeneities due to changes in the data assimilation system. Thus they are in principle more suited to use in studies of low frequency variability and trends in climate that complement studies of climate change based on individual instrumental records and climate-model simulations.

Two recent ECMWF reanalyses have exploited the substantial advances made in the ECMWF forecasting system and technical infrastructure since operations began in 1979. The first project ERA-15 (1979–1993) was launched in 1993 and the second "extended" reanalysis project ERA-40 (1957–2002) in 1998. The funding and support for these projects was shared by the Member and Cooperating States of ECMWF, EU framework programmes and several institutes notably in the USA, Japan and China. Pending resources sufficient to support the next extended reanalysis, ECMWF has recently begun the ERA-Interim reanalysis. This will run from 1989 onwards using 12-hourly 4D-Var and a T255L91 model, with observations and boundary fields from ERA-

40 and operations. Once it reaches the present day, ERA-Interim will continue in close to real time in so-called "Climate Data Assimilation System (CDAS)" mode. It will serve as an intermediate reanalysis between ERA-40 and the next extended reanalysis, and complement CDAS activities in Japan and the USA.

To address future reanalysis requirements a Workshop on Atmospheric Reanalysis was held at ECMWF from 19 to 22 June 2006. Funding was provided by ECMWF and GEO (Group on Earth Observations), and the programme for the meeting was developed in liaison with the WCRP Observation and Assimilation Panel (WOAP) and the GCOS/WCRP Atmospheric Observation Panel for Climate (AOPC). The workshop considered the status of and plans for global reanalysis in Europe, Japan, and North America, and discussed the work needed to prepare for the new generation of multi-decadal global reanalyses to succeed ERA-40, JRA-25 and the NCEP reanalyses. Conclusions of recent workshops were reviewed, including the 2005 NASA/NOAA/NSF Workshop on the Development of Improved Observational Data Sets for Reanalysis, the 2005 ECMWF/NWP-SAF Workshop on Bias Estimation and Correction in Data Assimilation, the 2005 ECMWF Workshop on a potential European Regional Reanalysis project and the 2006 SCAR/CliC/ICPM Workshop on High Latitude Reanalysis.

User views expressed at the workshop confirmed the need for improved climate products, especially as regards the representation of moist processes, surface fluxes, the southern hemisphere in the pre-satellite era, constituent transport and low-frequency variability and trends. More generally, users require the following.


- ◆ Good documentation of reanalysis systems and their performance.

- ◆ Measures of expected accuracy or uncertainty.
- ◆ Continuation of reanalyses in CDAS mode.
- ◆ Capability to access reanalysis data assimilation systems to perform observing-system experiments.
- ◆ Involvement in the planning of new reanalyses through User Advisory Groups.

Many users' needs can be met by reanalyses for the decades for which there is good upper-air data coverage by satellites or at least radiosonde data. The possibility of extending reanalysis to cover earlier periods when only surface observations are available in reasonable numbers (e.g. from the 1850s to the 1930s) is nevertheless of interest, and has been explored in two pilot studies comparing analyses with good coverage of satellite data and other upper-air data with analyses using only surface-pressure observations. For the ECMWF system it was shown that 4D-Var is capable of providing realistic analyses in the troposphere based on surface data only, provided the background error statistics used in the analysis are adapted to reflect the state of the observing system.

Instead of being viewed as a one-off effort, reanalysis has come to be seen as an iterative process, where developments in modelling, data-assimilation techniques, data-rescue efforts and computing power, and data and experience from earlier reanalyses, are utilized to produce successive reanalyses of increasing quality. To illustrate the evolving data assimilation and modelling capability at ECMWF, the verification of tropical wind forecasts for 1989 from the original operations, ERA-15, ERA-40 and a preliminary version of the ERA-Interim system are shown in the figure. New generations of extended "climate reanalyses" are likely also to exploit possibilities not open to daily operational weather prediction. A reanalysis system can, for example, be designed to make use of observations taken over a period after the analysis time as well as before. Utilization of the available observational information may be improved by lengthening the time window of the data assimilation in a weak-constraint 4D-Var system. New approaches for handling observational and model biases also show promise for use in reanalysis.

Despite being able to exploit the scientific and technical infrastructure developed for weather forecasting, reanalyses



Tropical wind root mean square forecast errors (ms^{-1}) at 850 hPa averaged over forecasts from 1st and 16th of each month of 1989 and 1990 for the ERA-Interim system, ERA-40, ERA-15 and operations. Forecasts from each system are verified against analyses from the same system.

remain large and demanding undertakings. They require multi-institutional collaboration to acquire not only the many types of daily data to be ingested but also the necessary boundary forcing fields and specifications of atmospheric composition. Product continuity has to be maximized, involving the homogenization of input observations in addition to the bias handling built into the assimilation system itself. The production phase of a reanalysis commonly runs to a tight timetable dictated by funding constraints or availability of computing resources. Intensive monitoring of production is needed to ensure satisfactory scientific and technical performance, and documentation, data services and user support are needed long after production is complete. Because of the complexity of these tasks, the potential for improved reanalysis cannot be realized without proper long-term organization, funding and international collaboration.

The report and presentations from the ECMWF/GEO Workshop on Atmospheric Reanalysis can be found at: www.ecmwf.int/newsevents/meetings/workshops/2006/re-analysis/

In addition more information about reanalyses can be found in web reports of the various workshops referred to earlier and articles about ERA-40 in the *Quarterly Journal of the Royal Meteorological Society* (2005, **131**, 2961–3012) and recent editions of the *ECMWF Newsletter*.

The new IBM Phase 4 HPC facility

Neil Storer

The European Centre for Medium-Range Weather Forecast's high-performance computing (HPC) facility is provided via a "service contract" with IBM, which originally ran until the end of March 2007. In the spring of 2004 ECMWF undertook an HPC market survey to decide how best to cater for its computational needs from 2007 onwards. As a result the contract with IBM was extended for a further two years, until the end of March 2009.

The contract extension calls for the two POWER4+ clusters of the Phase 3 system to be replaced in 2006 by two POWER5+ clusters. It was decided to install the two Phase 4 clusters in a phased implementation.

- ◆ One of the clusters would be installed in the new extension to the computer room, while the other would be pre-built in IBM's development labs in Poughkeepsie. This would enable a stable environment to be maintained on the cluster at ECMWF, while possibly disruptive experiments could be tested on the cluster in Poughkeepsie, where the hardware and software specialists reside.
- ◆ Once cluster-4E (i.e. the Phase 4 cluster in ECMWF) passed its acceptance tests, one of the old Phase 3 clusters would be decommissioned and the other cluster would be transported from the USA and installed in the space vacated by the old cluster.
- ◆ After the whole system passed its acceptance tests, the remaining Phase 3 cluster would be decommissioned.

At the time of writing cluster-4E has just passed its acceptance tests, cluster-3C has been de-commissioned and cluster-4F's acceptance tests are about to begin.

The Phase 4 system comprises two clusters, each with approximately 140 p5-575+ SMP (shared memory processors) servers – the actual number will depend on the performance of the total system when running ECMWF's benchmarks, and on the date on which the system as a whole is fully accepted. Each p5-575+ server (or node) has 32 GB of memory (a few have 128 GB) and contains 8 dual-core micro-processor chips, with each of the sixteen 1.9 GHz processors (or cores) having a peak performance of 7.6 GFLOPS. This gives a peak performance in excess of 34 TFLOPS for the total system. Although this is roughly the same as the peak performance of the Phase 3 system, extra features of the POWER5+ based system enable it to sustain a much higher percentage of peak performance on ECMWF's codes, giving it a sustained performance in the region of 1.8 times that of the Phase 3 system. The two clusters use IBM's multi-cluster GPFS (General Parallel File System) to enable them to access filesystems concurrently at much higher data-transfer rates than is available using other filesystem sharing mechanisms, such as NFS, that use IP protocols to transfer data.

Information about the technical aspects of the new features of the Phase 4 system can be found in the article "New features of the Phase 4 HPC facility" which appears in this edition of the *ECMWF Newsletter*. Also further information for users of the system is provided on ECMWF's website at:

www.ecmwf.int/services/computing/hpcf/

Once installation and commissioning is complete ECMWF's user community will have access to state-of-the-art high-performance computing resources for the next two and a half years, with which to progress their research in numerical weather prediction.

The main features of Phases 3 and 4 of ECMWF's IBM HPC systems.

Feature	Phase 3 – POWER4+	Phase 4 – POWER5+
System		
Number of clusters	2 computational clusters 1 test cluster	2 computational clusters 1 test cluster 1 MC-GPFS "owning" cluster
Sustained performance	~2.5 TFLOPS	~4 - 4.5 TFLOPS
Each Computational Cluster		
Operating system	AIX 5.2	AIX 5.3
I/O (VSD) nodes	6x8-way p690+ (LPARs)	8x16-way p5-575+
Network nodes	2x8-way p690+ (LPARs)	2x16-way p5-575+
Compute nodes	70x32-way p690+	> 140x16-way p5-575+
Node interconnect	pSeries High performance switch (dual-plane)	pSeries High Performance switch (dual-plane)
I/O Subsystem		
Disk types	FAStT500 + FAStT900	FAStT900
Amount of disk space	~25 TB per cluster	~50 TB per cluster
Multi-cluster GPFS	No	Yes
Each Server (p690+ or p5-575+)		
Memory	32 GB (2.24 TB per cluster)	32 GB (~4.5 TB per cluster)
Dual-core chips	16	8
Processors (or "cores")	32 (2,240 per cluster)	16 (~2,240 per cluster)
SMP building blocks	4 multi-chip modules (MCMs), 4 dual-core chips on each	8 dual-chip modules (DCMs), 1 dual-core chip on each
Each Processor (or "core")		
Lithography	130 nm	90 nm
No. of transistors	184 M (per dual-core chip)	276 M (per dual-core chip)
Frequency	1.9 GHz	1.9 GHz
Peak performance	7.6 GF (17 TF per cluster)	7.6 GF (~17 TF per cluster)
Level-2 cache	1.44 MB (8-way LRU)	1.92 MB (10-way LRU)
Level-3 cache	32 MB (8-way LRU)	36 MB (12-way LRU)
Memory controller	Off-chip	On-chip
Additional POWER5+ Features		
Upwards binary compatible with POWER4+		
Intra-MCM bus is 100% faster than on POWER4+		
Inter-MCM bus is 50% faster than on POWER4+		
Dynamic CPU electrical power management and reduced power leakage		
Fast barrier synchronisation		
Additional new instructions		
Simultaneous Multi-Threading (SMT)		
Mid-size memory page - 64 KB (a feature of the AIX 5.3 operating system)		

Third International Workshop on Verification Methods

Anna Ghelli

The Third International Workshop on Verification Methods will be held at ECMWF from 29 January to 2 February 2007. It is organised by the WWRP/WGNE Joint Working Group on Verification (JWGV) with ECMWF, WMO and WCRP sponsoring the workshop. The JWGV aims to encourage:

- ◆ Cooperation between users and verifiers.
- ◆ Development and application of improved diagnostic verification methods.
- ◆ Sharing of observational datasets.
- ◆ Awareness of the importance of verification as a vital part of numerical and field experiments.

Included in the workshop are both a tutorial session and a programme of talks on new developments in verification methodologies. The tutorial session will start on 29 January and go through to the morning of the 31st. The programme of talks starts on the morning of the 31st and finishes on the afternoon of 2 February. Note that there are only 30 places for the tutorial session.

The tutorial session will provide a pedagogical introduction and review of verification techniques for both deterministic and probabilistic (including ensemble) forecasts with hand-on laboratory sessions, whereby participants will use methodologies on real case studies. For these laboratory sessions, participants will be invited to bring their own datasets and verification problems.

The second part of the workshop will include keynote addresses focussing on new verification techniques and issues related to the practice of forecast verification, as well as

contributed presentations on verification methodologies applied to a variety of forecasts (including forecasts of phenomena outside of atmospheric sciences, such as economics). Subjects covered will include the verification of ensemble/probability forecasts, extreme events, forecast value and user issues, and spatial forecasts.

Further information about the workshop can be found at:
www.ecmwf.int/newsevents/meetings/workshops/2007/jwgv/

CTBTO: "Synergies with Science"

Manfred Klöppel

At a symposium to celebrate the tenth anniversary of the Preparatory Commission for the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO), the ECMWF Director, Dominique Marbouty, gave a speech on "*Data Assimilation in Weather Forecasting*". He described the important developments in data analysis achieved by the meteorological community during recent years, in particular allowing comprehensive use of satellite data. Emphasis was put on the usefulness of data assimilation not only as necessary tool for initialising weather forecasts, but also to provide direct feedback to operators of observation systems, to allow re-analysis of historic data, and to run Observation System Experiments (OSE) allowing assessments of the impact of various observing systems on the quality of weather forecasts.

One of the aims of the Symposium on "*CTBTO: Synergies with Science: 1996-2006 and beyond*", held in the Hofburg in Vienna from 31 August to 1 September 2006, was to stress

the synergies between the CTBTO Preparatory Commission and the scientific community. The intention was to increase the ability to build global capacity in relevant science and technology related to the activities of the Commission. Mr Tibor Tóth, Executive Secretary of the CTBTO Preparatory Commission, opened the Symposium, and a keynote speech was given by Dr Mohamed ElBaradei, Director General of the International Atomic Energy Agency (IAEA).

Sessions at the Symposium were dedicated to "Imaging the Earth", "Imaging the Atmosphere" and "Data Mining". The session devoted to "Imaging the Atmosphere" had meteorology playing a major role. All the presentations can be found by following the links from:

www.ctbto.org

The CTBTO Preparatory Commission is an international organisation related to the UN and based in Vienna. It carries out the preparations for the effective implementation of the Comprehensive Nuclear-Test-Ban Treaty. The Commission's main task is the establishment and operation

of an International Monitoring System and an International Data Centre; these activities assist in the verification of the Treaty. In order to get a good picture of the state of the atmosphere, the Commission has a capability to run atmospheric transport and dispersion models, which rely on data from established meteorological centres.

After approval by the ECMWF Council at its 56th session in June 2002, an agreement with the CTBTO Preparatory Commission was signed in October 2003. Since then, ECMWF has made available its real-time products for use in CTBTO's own transport modelling system to compute back-trajectories from radionuclide detections.

Summary of ECMWF Forecasts Product Users' Meeting, June 2006

David Richardson

The annual Forecast Products Users' Meeting was held at ECMWF during 14-16 June 2006. In the first part of the meeting the development and performance of the operational forecasting system was reviewed by ECMWF and plans for future changes, including product development, were presented. The subsequent sessions gave the opportunity to report on their experiences with the full range of ECMWF products (including the ensemble prediction system, and monthly and seasonal forecasts). Requirements for additional products in support of users' operational activities were consolidated in a final discussion session.

Participants included operational forecasters from National Meteorological Services, commercial providers, and academic users involved in developing products and assessing predictability at different timescales.

The range of applications (both public and commercial) continues to increase. Examples were presented from a variety of sectors including energy, media, transport, health and water management (drought planning and flood-risk assessment). Use of the ensemble prediction system (EPS) continues to grow, with several examples of new applications using EPS data to provide information on confidence and risk assessment. Several users reported positively on the monthly forecast system which was found to provide useful skill, especially for temperature up to three weeks ahead.

The general high quality of ECMWF products was noted in many presentations. Although relatively soon after the resolution upgrade, forecasters reported improvements for severe events and for the EPS. Two specific problems of winter forecasts were reported:

- ◆ Prediction of convective snow (extent of inland penetration).
- ◆ Prediction of temperature and low cloud in anticyclonic periods.

It was noted that improvements would lead to significant benefits for users.

Despite the wide range of user-developed applications, it is also clear that many users appreciate and benefit from the products developed at ECMWF. The ease of access and wide range of products available through the ECMWF website is widely appreciated.

The programme for the Users' Meeting, together with the presentations and the conclusions from the final discussion, can be found on the ECMWF website at:

www.ecmwf.int/newsevents/meetings/forecast_products_user/Presentations2006/

ECMWF Operational Forecast System

Atmospheric data assimilation

- ◆ 4-dimensional variational data assimilation at 25/80 km resolution and 91 levels

Atmospheric global forecasts

- ◆ Forecast to ten days from 00 and 12 UTC at 25 km resolution and 91 levels
- ◆ 50-member ensemble to ten days from 00 and 12 UTC at 50 km resolution, 62 levels

Ocean wave forecasts

- ◆ Global forecast to ten days from 00 and 12 UTC at 40 km resolution
- ◆ European waters forecast to five days from 00 and 12 UTC at 25 km resolution
- ◆ 50-member global ensemble to ten days from 00 and 12 UTC at 110 km resolution

Monthly forecasts:

Atmosphere-ocean coupled model

- ◆ Global ensemble to one month (weekly) 125 km, 62 level atmosphere; 1°, 29 level ocean

Seasonal forecasts:

Atmosphere-ocean coupled model

- ◆ Global ensemble to six months (monthly): 200 km, 40 level atmosphere; 1°, 29 level ocean

ECMWF Forecast Products

Global forecasts (deterministic)

- ◆ Fields (Model variables, GRIB)
- ◆ Time series (BUFR Meteograms)
- ◆ Direct Model Output (DMO) only

Ensemble Prediction System (EPS)

- ◆ Fields (GRIB) + Time Series (BUFR EPSGRAMs)
- ◆ Post-processed products (Clusters, Tubes, EFI)

Monthly and seasonal forecasts:

atmosphere-ocean coupled model

- ◆ Anomalies (Fields + Time Series) from model climate

NCEP uses ECMWF's methodology for humidity analysis and background error generation

Elías Valur Hólm

On 20 June 2006 NCEP (National Centers for Environmental Prediction, USA) implemented a new national forecast model and analysis system replacing the over decade old ETA model and 3D-Var. The new Weather Research and Forecasting (WRF) system includes a non-hydrostatic model and the Gridpoint Statistical Interpolation (GSI) analysis.

One of the main model changes is the change in vertical coordinates. The new WRF model uses hybrid sigma-pressure layers, like the ECMWF model, which replace the step-mountain eta-layers in the ETA model, and the models top layer has been extended to 2 hPa from 25 hPa.

In the new GSI analysis, two of the main new features are built on ECMWF developments, namely the background error covariance generation and humidity analysis variable. According to NCEP scientists, both of these features individually improve forecasts.

The background error covariances, which were previously generated from lagged forecast differences (a method developed at NCEP and previously also used at ECMWF) are now

generated from a Monte-Carlo method (ensemble of forecast differences) developed and used at ECMWF (see *ECMWF Technical Memo No. 347*). The new background errors are more compact and localized than the previous version.

The humidity analysis variable, which used to be pseudo-relative humidity (*Dee & Da Silva, 2002, Monthly Weather Review, 131,155–171*), is now normalized relative humidity, which was previously developed and implemented at ECMWF (see *ECMWF Technical Memo No. 383*). This humidity variable takes into account the heterogeneous nature of the humidity field by dividing analysis increments by a relative humidity and height dependent variance, derived from the forecast differences used in the Monte-Carlo background error calculations. The result is a flow-dependent estimate of the humidity background errors, which better accounts for transitions between dry and moist regions and avoids negative humidity and excessive supersaturation in the analysis. Both pseudo-relative humidity and normalized relative humidity are an option in the GSI analysis, and future global analysis developments can use either of these.

For further information go to:

www.emc.ncep.noaa.gov/WRFinNAM/.

Annual Meetings of the European Meteorological Society

David BurrIDGE (EMS President)

The sixth Annual Meeting of the European Meteorological Society (EMS) and the sixth European Conference on Applied Climatology (ECAC) were held jointly in the Cankarjev dome conference centre in Ljubljana, Slovenia, from 4 to 8 September 2006. The joint conference brought together over 500 scientists, service providers and manufacturers from across Europe. It provided a forum for the exchange of ideas on future strategies in climatology and meteorology, and provided a stimulating framework for developing increased collaboration in the fields of climate monitoring, research and prediction and the provision of climate data for the research community.

ECMWF was represented at the meeting with several staff making contributions. In particular Dominique Marbouty (ECMWF Director) gave a presentation during the opening session on “*European developments in Numerical Weather Prediction*” and Horst Böttger (Head of Meteorological Division) chaired the session on “*Instruments and Methods of Observations*”. ECMWF supports EMS by being an Associate Member.



El Escorial – the location of next year's EMS conference which will be held jointly with the European Conference on Applied Meteorology (ECAM).

As in previous years, the EMS employed the Copernicus organisation (<http://meetings.copernicus.org>) to manage the event and the Environmental Agency of the Republic of Slovenia and the Slovenian Meteorological Society provided valuable local support.



Opening session speakers and local hosts. From left to right: Franco Einaudi (President of the AMS), David Burridge (President of the EMS), José Achache (Director GEO Secretariat), Silvo Žlebir (Director General, Environmental Agency of the Republic of Slovenia), Georgios Amanatidis (EC, DG Research), Mikael Rattenborg (EUMETSAT), Tanja Cegnar (Environmental Agency of the Republic of Slovenia), Janez Podobnik (Ministry of the Environment and Spatial Planning), Michel Jarraud (Secretary General of WMO), Marjan Vezjak (Director General, Directorate for European Affairs and Investments), Dominique Marbouty (Director of ECMWF), Jožef Rošar (Director of Slovenian Met Service), Mark Žagar (President of Slovenian Meteorological Society).

I am delighted that the EMS aim of establishing an international European conference that focuses on applications and involves the whole meteorological and climate community became a reality in 2006.

Next year's EMS conference will be held jointly with the European Conference on Applied Meteorology (ECAM) from 1 to 5 October 2007. It will take place in El Escorial which is about 41 km to the northwest of Madrid. The preliminary list of possible themes for the 2007 conference is:

- ◆ Meteorology and customer value
- ◆ Data policy
- ◆ The forecasting process
- ◆ Atmosphere and the water cycle
- ◆ Climatology
- ◆ Computing in atmospheric sciences
- ◆ Instruments and observations
- ◆ Information provision, education, history
- ◆ WMO programme THORPEX
- ◆ Socio-economic impact
- ◆ The ENSEMBLES project
- ◆ COST
- ◆ User sessions
- ◆ A satellite plenary session (ESA, EUMETSAT)

User sessions, for which there will be a call, will be considered as special themes in the programme; they will get special attention and support of the entire conference. Where different themes have a common sub-theme, joint sessions will be organised (e.g. extreme weather).

Participants of EU projects will be invited to hold their working group meetings at the conference and organise sessions to inform the European community involved in meteorology, climatology and hydrology about the subject and progress of their work – examples are ENSEMBLES, ACCENT, GMES, EMMA, EURORISK (early warning system) and AMMA.

The Instituto Nacional de Meteorología (INM) and the Asociación Meteorológica Española (AME) have agreed to be the local organisers. These institutions will co-operate closely with the EMS to provide a suitable environment for a successful EMS/ECAM 2007 event. In addition the incoming President of the AMS, Rick Anthes, is keen to support the 2007 EMS conference and the AMS is planning to sponsor and contribute to the organisation of some of the sessions.

With the strong support of INM and the excellent organization provided by Copernicus I expect the biennial EMS/ECAM conference to be a highlight in the European meteorological calendar.

Analysis and forecast impact of humidity observations

Erik Andersson, Elías Hólm,
Peter Bauer, Anton Beljaars, Graeme Kelly,
Tony McNally, Adrian Simmons,
Jean-Noël Thépaut, Adrian Tompkins

We have found that humidity observations have a significant impact on analyses and forecasts extending into the medium range (5–6 days), with a marked impact also on the wind field. This is in contrast to Bengtsson & Hodges (2005) who found only small-scale, unstructured temperature differences in the tropics, and no noticeable effect on the skill of the wind forecasts. They explained that the model is capable of forming realistic humidity fields through assimilation of temperature and wind data, to the extent that the addition of humidity observations has negligible effect. Furthermore, in the ERA-40 reanalyses the assimilation of satellite data led to a more poorly balanced global hydrological cycle. ECMWF has improved its humidity assimilation and moist-physics parametrizations (namely clouds, convection and vertical diffusion), and demonstrated that the hydrological cycle is now significantly better balanced with respect to the ERA-40 system. Consequently we decided it was now time to reassess the impact of humidity observations on analyses and forecasts.

In October 2003 an improved statistical model of humidity background errors was implemented (Hólm *et al.*, 2002); this is an important component of the analysis scheme as it determines how the humidity information is distributed away from observation points. Subsequently, data from several additional instruments on satellites in both geostationary (METEOSAT and GOES) and polar orbits (AMSU-B and AIRS) have been introduced. Also the use of radiosonde and surface humidity data has been revised based on current data quality statistics. This concerted effort on humidity analysis is motivated by the increased availability of humidity observations, and by the need to improve the assimilation in cloudy and precipitating regions. Some of the areas where moisture is particularly important are as follows.

- ◆ The latent heat release from strong convective events can modify the jet-stream aloft and influence subsequent down-stream developments.
- ◆ The moisture content of the air on the warm side of a frontal zone can influence the rate of development of baroclinic systems.

- ◆ In the tropics, the supply of low-level humidity can affect the intensity of the tropical convection, and hence the intensity of the Hadley circulation.

The humidity observing systems

Our humidity impact experiment was run from 1–31 July 2003, using the 4D-Var assimilation system with 12-hourly cycling at T319 horizontal resolution (~60 km), 60 model levels and analysis increments at T159 (~120 km) resolution. The October 2004 version of the forecast system was used throughout. Its standard configuration, which uses a large variety of conventional observations and satellite radiances, provided the 'Control'. In the experiment 'Noq' all humidity observations were withheld. See Andersson *et al.* (2006) for a more detailed report on the separate impacts of each of the main humidity observing systems.

First, we describe the main humidity observing systems, in the order they appear in Table 1.

- ◆ **Polar orbiting microwave radiances.** The SSMI microwave instruments provide radiance data in seven channels. These channels are sensitive to the integrated atmos-

Humidity measurement dataset	Daily number of data	Number of channels (available)	Main humidity information	Typical data coverage of used data
SSMI, DMSP-13, 14, 15 Polar orbiting microwave radiances	220,000	7 (of 7)	Total column, except in clouds/rain	Ice free ocean
TEMP Radiosonde specific humidity	21,500	–	Tropospheric humidity profiles with high vertical resolution	Concentrations over North America, Europe, Eastern Asia and Australia
SYNOP 2m relative humidity	13,200	–	Boundary layer humidity	Irregular with concentrations in populated regions
GEOS, Meteosat 5,7, GOES-9, 10, 12 Geostationary IR radiances	141,500	1 (of 2)	Upper troposphere, clear air	Within 50° of the equator, cloud free only
AMSUB, NOAA-16, 17 Polar orbiting microwave radiances	131,500	3 (of 5)	Upper and mid troposphere	Irregular, emissivity dependent over land and ice; good over ice free ocean
HIRS, NOAA-16, 17 Polar orbiting IR radiances	120,000	6 (of 19)	Mostly upper troposphere, clear air	Cloud free ocean and ice
AIRS, AQUA Polar orbiting IR radiances	280,000	230 (of 2378)	Upper and mid troposphere, clear air	Cloud free ocean and ice

Table 1 Humidity observing systems used in ECMWF data assimilation. Approximate data counts, the main humidity information and typical data coverage are also given.

pheric water-vapour content and to wind-induced sea-surface roughness.

◆ **Radiosonde specific humidity.** Radiosondes provide dew-point temperature observations which are converted to specific humidity using the observed temperature and pressure. The humidity profiles from Vaisala's RS90 and RS80 sondes are used at all reported levels up to 100 hPa, subject to a temperature threshold which is -80°C for RS90 and -60°C for RS80. All other sondes are used only up to 300 hPa, subject to the temperature being higher than -40°C .

◆ **Two-metre relative humidity.** The SYNOP two-metre dew-point observations are converted to relative humidity. These data are used over land, but not over sea, during local daytime only (i.e. solar elevation angle greater than zero). The removal of night-time SYNOP data is motivated by poor representativity of the data in stable surface-layer conditions.

◆ **Geostationary IR radiances.** From geostationary satellites, clear-sky radiance data from the $6.3\text{ }\mu\text{m}$ water-vapour channel of each of five platforms (GOES-9, 10 and 12 and METEOSAT-5 and 7) are used. These data provide a complete and frequent coverage within about 50° latitude of the equator.

◆ **Polar orbiting microwave radiances.** From AMSUB, three of the five available 183 GHz microwave channels are assimilated depending on the land/sea mask and the height of the terrain.

◆ **Polar orbiting IR radiance.** From the infrared sounding instruments (HIRS and AIRS) it is primarily the $6.3\text{ }\mu\text{m}$ band that carries humidity information. For HIRS this corresponds to channels 11 and 12, and for AIRS it comprises channels 1290–1843. The infrared satellite data carry ambiguous temperature and humidity information.

The satellite systems provide very good coverage over the oceans, with gaps in cloudy and precipitating regions depending of the sensitivity to clouds in the infrared, and to water clouds and rain in the microwave measurements. The geostationary radiances and some of the higher-peaking channels of AIRS, HIRS and AMSUB are used also over land. The conventional data (SYNOPs and radiosondes) provide an uneven coverage over land, with dense concentrations over parts of North America, Europe, East Asia and Australia.

Analysis bias differences

The analysis impact of any data type depends on the data coverage, the frequency of the data and their accuracy. The impact also depends on the specification of background errors in the assimilation scheme, and on the existence of any systematic deficiencies and biases in the forecast model and observations. In these experiments, satellite radiance data are corrected for air-mass dependent and scan-angle




Figure 1 Zonal-mean monthly mean (2–31 July 2003) cross-sections of relative-humidity analysis differences (%). The contour interval is 0.2 % with red (blue) indicating that the *Control* assimilation is moister (drier) than the *Noq* experiment which withholds the humidity data.

dependent biases. There have been several reports on persistent observation bias also in radiosondes. However, radiosonde and SYNOP humidity data are currently not subject to bias correction at ECMWF.

Differences between the *Control* and the *Noq* experiments are shown in Figure 1, in the form of north-south cross-sections, averaged over the study period. This shows that the humidity observations substantially modify the moisture analysis: in the northern hemisphere and the tropics there is an increase in moisture in the lower troposphere and a decrease at higher levels. Geographical maps of the differences (not shown) indicate that the moisture is added in the subsidence regions, where the background fields are biased dry. The added moisture is advected to the ITCZ region by the trade winds, leading to an impact on precipitation in that region (see Figure 3 which is described later). Detailed investigations showed that over many continental areas the radiosondes and SYNOPs contribute biases in the boundary layer with opposite sign. Over parts of Europe and North America ($40\text{--}70^{\circ}\text{N}$) radiosondes contribute to a drying of the analysis in the upper troposphere, which is consistent with published literature showing that several radiosondes have upper-tropospheric dry biases. While absolute amounts are small, upper-tropospheric humidity is




Figure 2 Mean analysis difference (*Control*–*Noq*) in total-column water vapour. The contour interval is 0.5 kg m^{-2} with red (blue) indicating that the *Control* assimilation is moister (drier) than the *Noq* experiment which withholds the humidity data.




Figure 3 Average daily precipitation difference (2–31 July 2003) in 12-hour forecasts from analyses at 06 and 18 UTC, showing *Control* minus *Noq*. The contour interval is 1 mm/day with red (blue) contours indicating that the assimilation of observed humidity has increased (decreased) precipitation.

important for accurate radiative transfer calculations, and also for the indirect influence on cirrus cloud formation.

Figure 2 shows the impact of the humidity observations on the total-column water vapour (TCWV) in terms of the mean analysis difference between *Control* and *Noq*. TCWV is dominated by the moisture content of the warm air in the lower troposphere and at lower latitudes. The mid- and upper-tropospheric differences in relative humidity seen in Figure 1 contribute very little in terms of TCWV. In Figure 2 we can see that the net effect of all assimilated humidity observations is to add moisture. Over sea, the TCWV differences are almost entirely due to SSMI data, whereas over land, the differences are due to the combined effects of radiosonde and SYNOP data. Although we have seen that radiosondes have a drying effect in the boundary layer and in the upper troposphere, the net effect in terms of TCWV is a slight moistening in most regions with good radiosonde coverage, including parts of North America. Also SYNOP assimilation adds significant moisture to the Sahel and Sub-Sahara regions of Africa. As the humidity increments at the lowest model level are used as input to the soil water analysis, it is likely that the impact of SYNOPS in this region is through its interaction with the soil moisture analysis.

Short-range precipitation forecast

The impact of the humidity observations on precipitation for parts of the tropics and North America is shown in Figure 3. This shows that the rainfall in the first 12 hours of the forecasts is increased in the Western Pacific due to the assimilation of humidity data. Within the ITCZ there is no significant difference in precipitation intensity in these experiments, but it is evident from the differences that the assimilation of humidity data has modified the location of the ITCZ in the East Pacific. Nearly all the differences seen over the ocean in Figure 3 are due to SSMI. The spreading of humidity impact from clear-sky to precipitating areas is due to transport and physical processes in the model, and is also due to extrapolation by the analysis. The latter effect is determined by the analysis structure functions (i.e. the background error covariance matrix) which currently does not recognize the two regimes, and whether there should be decorrelation of humidity increments across cloud and precipitation boundaries.

More detailed investigations showed that assimilation of radiosondes and SYNOP humidity data locally increases the rainfall in the regions where Figure 2 showed that these data on average add moisture: parts of North America, Europe, India, and in particular central Africa where SYNOP data add moisture to the boundary layer and the soil moisture analyses. In the mid-latitude storm-track regions, the net impact of observed humidity on precipitation is smaller than in the tropics but may nevertheless be important in relative terms. Also Figure 3 shows reduced precipitation in the western parts of the Atlantic and Pacific oceans, which is due to assimilation of SSMI data.

In ERA-40 and in earlier versions of the ECMWF forecasting system there was a rapid adjustment during the first day of forecasts of the tropical rate of precipitation, becoming almost constant at lower rates thereafter. This so-called ‘spin-down’ problem has now been significantly reduced, through changes to the moist physics parametrizations and the assimilation system. Due to the spin-down problem the assimilation of observed humidity observations produced up to 50% more tropical precipitation early in the forecasts than assimilations without humidity data. It is evident from Figure 3 that the current *Control* humidity analyses are in good balance with the forecast model, and compared to the *Noq* analyses they do not result in excessive amounts of precipitation in the early stages of the forecasts.

Forecast verification

Differences in latent heat release associated with the above mentioned differences in precipitation result in temperature differences throughout the troposphere. In forecasts, the large-scale fields of geopotential and wind are influenced where the evolution of weather systems is affected by changes in the moisture distribution. Furthermore, the convection parametrization directly changes the wind through its vertical momentum transport. To assess the benefit to the forecasts of humidity assimilation, ten-day forecasts have been run daily for the *Control* and *Noq* experiments. The scores shown here are root mean square errors (RMSE), computed with the operational ECMWF analysis as reference.

Figure 4 shows the forecast impact of observed humidity in relative terms: values of $\text{RMSE}(\text{Noq}) - \text{RMSE}(\text{Control})$ normalized by the mean of the two RMSEs. Positive values

of this score mean that the *Noq* forecast errors are larger than those of the *Control*, which indicates a beneficial impact of assimilating humidity observations. Conversely, negative values of the score would indicate deterioration. The scores are plotted as a function of forecast range from day 1 to day 10, with the error bars indicating 90% two-sided confidence intervals. Scores are shown for the northern hemisphere extratropics, the tropics within 20° of the equator, and the southern hemisphere extratropics, for 300 hPa vector wind and 850 hPa relative humidity. We see that there is a clear positive impact from assimilating humidity observations in

all three regions. The humidity forecast impact is initially very large (>15 % in the northern hemisphere and tropics) but falls off rapidly during the first four days of the forecasts. This indicates that the humidity field is strongly forced by the dynamics and exchanges with the surface. The improved humidity initial conditions lead to modified precipitation, which in turn affects the dynamical fields (e.g. through latent heat release and convective momentum transfer).

The results in Figure 4 show that the dynamic impact of observed humidity is largest in the tropics (6 to 9%), and it is larger in the southern (3 to 6%) than in the northern




Figure 4 Forecast impact of observed humidity in terms of normalized RMSE for 300 hPa vector wind (left) and 850 hPa relative humidity (right) in the northern hemisphere extratropics, tropics and southern hemisphere extratropics. The error bars indicate 90% two-sided confidence intervals. The normalised RMSE is given by $\{RMSE(Noq) - RMSE(Control)\} / \{0.5[RMSE(Noq) + RMSE(Control)]\}$.

hemisphere (2 to 4%). Where the lower bound of the error bar is above the zero line, the positive impact is significant with at least 95% confidence, which is the case for the first 4 to 5 days of our experiments. In the tropics we see a very significant impact on the upper-tropospheric wind field (8% at day 1) and temperature field (6% at day 1) which remains significant and positive to days 7 or 8 of the forecast.

Perspectives

Several additional sources of moisture information can be exploited in the near future to bring further benefit to the humidity analysis.

- ◆ **Infrared radiances (HIRS, AIRS and MSG).** The use of infrared radiances (HIRS and AIRS) continues to be developed (e.g. to better account for the influence of clouds and aerosol on the measurements). The long-term provision of high-spectral resolution infrared data (like AIRS) is ensured through U.S. and European programmes (the CrIS and IASI instruments, respectively). Furthermore, the first in the Meteosat Second Generation (MSG) series of geostationary satellites was launched in August 2002 providing frequent data in one additional water vapour channel. These data have been assimilated in ECMWF operational system since June 2005. Currently, infrared radiances are given reduced weight in the analysis; the observation errors are significantly inflated (up to 2 K) to offset known deficiencies in the use of the data. For example, correlations of observation error between channels and within swaths of data are ignored. The distribution in the vertical of the analysis increments is another delicate problem. In step with improvements in these areas, more weight will be assigned to what are intrinsically high-quality radiance measurements, resulting in a more robust and important influence upon the humidity analysis.
- ◆ **Microwave radiances from new instruments.** Microwave radiances from new instruments will be assimilated in a way that complements the existing imager-sounder combination. These are the Special Sensor Microwave Imager Sounder (SSMIS) that combines SSMI with AMSU-A/B type channels, the Tropical Rainfall Measuring Mission (TRMM) Microwave Imager (TMI) and the Advanced Microwave Scanning Radiometer (AMSR-E) onboard the AQUA satellite; all these instruments have specifications similar to the SSMI. While these instruments will not provide additional information on the vertical distribution of moisture, their orbit configuration will greatly improve the data coverage. The data from these instruments is currently monitored and is expected to be actively assimilated later in 2006.
- ◆ **SSMI radiances.** Since June 2005, SSMI radiances affected by clouds and precipitation are assimilated at ECMWF through a 1D+4D-Var analysis method. As with the SSMI data in clear areas, the impact on the moisture analysis is significant. The bias impact of the clear-sky data (as seen in Figure 2) has in later experimentation been partly compensated by the assimilation of rain-affected SSMI

data. The optimal combined use of clear and cloudy data is strongly dependent on the definition of moisture background errors inside and outside clouds, the horizontal structure functions, and the current methodology of first retrieving the cloud/rain-affected TCWV through 1D-Var. Developments towards a direct 4D-Var assimilation of rain-affected SSMI radiances are underway.

- ◆ **Radiosonde and aircraft humidity sensors.** There has been a gradual improvement in the accuracy of radiosonde humidity sensors, and recent years have seen the introduction of the Vaisala RS90 and RS92 sondes into operational use. The number of stations using the latest types of sonde is increasing. However, there remains a clear need for bias correction, which can be the result of calibration errors, time-lag errors, sensor icing errors, sensor aging or contamination, or radiative sensor heating effects. Some of the causes have been tackled by improved radiosonde design; for instance, the RS90 and RS92 introduced a pulse-heated twin sensor design to eradicate sensor icing and improve sensor time-lag errors. Other error sources have been addressed by post-processing bias correction techniques. Humidity sensors suitable for commercial aircraft are being developed in order to add humidity to AMDAR reports in the future, and enhance the humidity profiling capability over land.
- ◆ **Global Positioning System humidity data.** GPS radio occultation techniques are being developed that may improve humidity information in the upper troposphere and tropopause regions. Ground-based GPS measurements provide total-column humidity information that is used experimentally in NWP assimilations. Near real-time GPS networks are being coordinated in Europe, North America and Japan, and real-time European data has been received at ECMWF since March 2004. The use of these new data types in assimilation will be explored in the coming years.

Good quality humidity analyses are becoming even more important as the resolution of the forecasting system increases. The emphasis is on convective and severe precipitation events, which requires accurate depiction of the global distribution of moisture.

FURTHER READING

- Andersson, E., E. Hólm, P. Bauer, A. Beljaars, G.A. Kelly, A.P. McNally, G.A. Kelly, A. Simmons & J.-N. Thépaut, 2006: Analysis and forecast impact of the main humidity observing systems. *ECMWF Tech. Memo. No. 493*.
- Bengtsson, L. & K.I. Hodges, 2005: On the impact of humidity observations in numerical weather prediction, *Tellus*, **57A**, 701–708.
- Hólm, E., E. Andersson, A. Beljaars, P. Lopez, J.-F. Mahfouf, A.J. Simmons & J.-N. Thépaut, 2002: Assimilation and modelling of the hydrological cycle: ECMWF's status and plans. *ECMWF Tech. Memo. No. 383*.
-

Hindcasts of historic storms with the DWD models GME, LMQ and LMK using ERA-40 reanalyses

Helmut P. Frank, Detlev Majewski
Deutscher Wetterdienst Research and
Development, Germany

Understanding the circumstances in which coastal floods occur is very important as preparations can then be made to reduce the impact of these events.

The main forcing for the coastal floods is the surface wind.

To analyze historic floods NLWKN (Forschungsstelle Küste des Niedersächsischen Landesbetriebs für Wasserwirtschaft, Küsten- und Naturschutz, www.nlwkn.de), the Coastal Research Station of Lower Saxony Water Management, Coastal Defence and Nature Conservation Agency, contracted Deutscher Wetterdienst (German Weather Service, DWD) to provide high-resolution winds for the Ems river estuary in northwestern Germany. NLWKN chose 22 storms from the famous Hamburg Storm in February 1962 to a storm in October 2002. Most of these storms occurred in winter, and only one in summer (19 August 1990). Table 1 lists the dates chosen for the hindcasts and the strongest winds observed at the east Friesian island Norderney off the coast of Lower Saxony.

To perform these hindcasts it was necessary to have a modelling capability with adequate resolution. This was provided by the three linked models GME, LMQ and LMK of DWD. In addition access was required to high-quality analyses from the last forty years. The ERA-40 reanalyses prepared at ECMWF were used for this purpose.

The results show that the combination of the three DWD models starting from the ERA-40 reanalyses provide information of the kind required by NLWKN.

Models and set-up

DWD produced wind fields during these storms using its model chain GME, LMQ and LMK starting from ERA-40 reanalysis data. LMQ and LMK are versions of the Lokal-Modell (LM). Details about these models are given in Box A and there is a short summary of their characteristics in Table 2.

A series of 18 hour forecasts starting from 00 and 12 UTC analysis data was used to obtain high-resolution hourly wind fields along the German North Sea coast with a grid spacing as small as 2.8 km. To allow for an adaptation of the models to the initial fields only forecasts from 6 to 18 hours are provided to NLWKN. The 6 and 18 hour forecasts for the same verification time are averaged.

The ERA-40 reanalyses were produced with ECMWF's Integrated Forecasting System (IFS) using a spherical harmonics representation TL159 and a reduced Gaussian grid corresponding approximately to a mesh size of 125 km (Uppala *et al.*, 2005). Strong cyclones cannot be well resolved with such a resolution. Therefore, it is scaled down in three steps: first, to the GME model with 40 km resolution, then to LMQ with 7 km mesh size, and finally to LMK with 2.8 km mesh size.

Start and end dates of the run	Maximum wind (ms ⁻¹)	Time of maximum wind
15 February to 18 February 1962	–	No data
22 February to 24 February 1967	26.7	18 UTC on 23 February 1967
2 January to 4 January 1976	33.4	06 UTC on 3 January 1976
19 January to 22 January 1976	25.7	00 UTC on 21 January 1976
29 December to 31 December 1977	23.1	19 UTC on 30 December 1977
22 November to 26 November	23.6	07 UTC on 24 November 1981
3 December to 5 December 1988	14.9	05 UTC on 5 December 1988
12 February to 15 February 1989	22.6	07 UTC on 14 February 1989
24 January to 27 January 1990	20.5	00 UTC on 26 January 1990
19 August to 21 August 1990	23.1	04 UTC on 21 August 1990
18 December to 21 December 1991	14.9	10 UTC on 18 December 1991
21 January to 23 January 1993	17.5	00 UTC on 23 January 1993
8 December to 10 December 1993	19.0	19 UTC on 9 December 1993
18 December to 20 December 1993	20.1	22 UTC on 19 December 1993
25 January to 29 January 1994	24.2	07 UTC on 28 January 1994
12 March to 14 March 1994	21.1	20 UTC on 13 March 1994
31 December to 2 January 1994	19.5	11 UTC on 1 January 1995
7 January to 11 January 1995	25.2	04 UTC on 10 January 1995
29 October to 30 October 1996	22.1	11 UTC on 29 October 1996
3 February to 6 February 1999	23.0	06 UTC on 5 February 1999
2 December to 5 December 1999	25.0	18 UTC on 3 December 1999
25 October to 29 October 2002	26.0	20 UTC on 27 October 2002

Table 1 List of the analysed storms and the maximum observed 10 m wind at the island Norderney.

Box A**Global model GME**

The global model GME is a hydrostatic weather prediction model (Majewski *et al.*, 2002). It operates on the icosahedral-hexagonal grid. The model has a mesh size of approximately 40 km and 40 layers up to 10 hPa.

For operational runs at DWD an intermittent data assimilation based on optimum interpolation is used to determine the initial state of the model. For these hindcasts the initial state was interpolated from the ECMWF ERA-40 reanalysis to the GME grid. A digital filter initialization over 1 hour with a cut-off period of 3 hours was performed prior to the forecast to reduce noise from the interpolation of the ERA-40 analysis.

LMQ and LMK versions of Lokal-Modell

The Lokal-Modell (Steppeler *et al.*, 2003) is a non-hydrostatic limited-area atmospheric prediction model operating on the meso-beta and meso-gamma scale. It uses a regular C-grid in rotated geographical coordinates. It is applied here in two different setups.

The first setup is called LMQ. In this case the mesh size is 7 km (0.0625°). A leap-frog scheme with a time step of 40 s is used for time integration. Deep and shallow convection are parametrized using a mass flux approach. The model domain consists of 333×333 grid points. The initial state for LMQ is interpolated from the GME initial state. LMQ is nested within GME with an hourly update of lateral boundary values.

Within LMQ the high-resolution version LMK is nested. Its mesh size is only 2.8 km (0.025°). Hence, it is assumed that deep convection can be explicitly resolved by the model. Only shallow convection is parametrized through moisture convergence in the planetary boundary layer. For time stepping a Runge-Kutta scheme of third order with time step 30 s is used. Initial values are interpolated from LMQ. Lateral boundary values, taken from LMQ, are updated hourly.

Both, LMQ and LMK have the same rotated geographical coordinates with the north pole at 40°N, 170°W.

Characteristic	GME	LMQ	LMK	ERA
Mesh size (km)	40	7	2.8	125
Domain (km ²)	Global	2324×2324	1176×1288	Global
Grid points	10×192 ² +2 = 368642	333×333 = 110889	421×461 = 194081	
Levels	40	40	50	60
Top height (km)	~31	23.6	22.0	~65
Top pressure (hPa)	10	~20	~39	0.1
Time step (s)	133.3	40.0	30.0	
Time stepping	LF	LF	RK	

Table 2 Characteristics of the models and domains (LF = Leap Frog, RK = Runge-Kutta).

The nesting of LMK in LMQ and of LMQ in GME is shown in Figure 1. For this plot the GME orography was interpolated to the LMQ rotated geographical grid using the nearest neighbour as interpolation method. NLWKN is interested in wind fields in an area around the Ems river estuary at the North Sea coast along the Dutch-German border. The LMK domain cuts through the North Sea and crosses the Alps. It is not optimal for this application. However, the simulations served also as test bed for LMK which was put in daily pre-operational suite in summer 2006.


As there is no data assimilation for the models we allow 6 hours of adaptation to the interpolated initial conditions before the forecasts are used. During these hours the deepening of the central pressure in cyclones of the high-resolution model compared to that for the lower-resolution driving model can take place.

Results for some storms

We present results for three arbitrarily chosen storms. The storms selected are one from the 1970s, the only summer storm and the very strong storm in December 1999 (which is remembered by one of the authors for the broken trees and news about the destructions in Denmark).

For the plotted maps the GME data was interpolated from the original hexagonal grid to the rotated coordinates of LM with a mesh width of 0.25°. The ERA-40 forecast data had been interpolated to a regular grid of 0.5° using ECMWF's MARS software. For the maps it was interpolated from this regular grid to the LM coordinates with a mesh width of 0.25°. As the reduced Gaussian grid has a coarser mesh width this interpolation should not reduce the winds significantly.

Figure 1 Model orography and nesting of LMQ in GME (black box) and LMK in LMQ (red box).



Storm on 30 December 1977

On 30 December 1977 a cyclone passed from the north Atlantic across southern Norway and Sweden to the Baltic Sea. In the lee of the Norwegian mountains a secondary low developed over Denmark. The strongest winds occurred west of this low over the North Sea – see Figure 2. The highest winds are predicted by LMQ. LMK shows more details such as the realistic streaks of high winds over the North Sea.

Time series for the island Helgoland in the German Bight are compared with observations in Figure 3. The breaks in the curves come from missing observations or from the jump from an 18 hour forecast of one run to the 6 hour forecast of the following run. The second pressure minimum at 18 UTC on 30 December 1977 is over-estimated by GME, LMQ and LMK, and there is a phase shift of the maximum wind. However, the strength of the maximum wind of approximately 20 ms^{-1} is well predicted by all three models. The reanalysis cannot capture it. Wind direction is very similar for all models.

LMK is the only model to simulate strong fluctuations of the wind speed between 4 and 11 UTC on 30 December 1977 similar to the observations.




Figure 2 Mean wind speed at 10 m in ms^{-1} at 21 UTC on 30 December 1977 simulated by GME, LMQ, LMK, and the ERA-40 forecast (ERF) initialised at 12 UTC on 30 December 1977. The mean sea level pressure in hPa is shown by isobars. The LMK domain does not cover the whole area shown in the maps.




Figure 3 Time series of sea level pressure (hPa), wind speed (ms^{-1}) and wind direction (degrees) observed at DWD station Helgoland and predicted by the models for the storm on 30 December 1977. ERF is ERA-40 forecast data every 3 hours, and ERA40 is the ERA-40 reanalysis every 6 hours. Observations are drawn as full bold lines.

This can be an indication that the wind streaks seen in Figure 2 are realistic.

Storm on 20 August 1990

The summer storm moved from Scotland to Denmark. For almost one day it stayed over the North Sea hardly advancing. During this period the wind direction was westerly at Norderney while the wind speed increased from 10 ms⁻¹ to 20 ms⁻¹ as shown in Figure 4. There was a band of strong winds along the cold front. However, the strongest winds followed in the cold sector. Observations at Helgoland and on the east Friesian island Norderney directly off the coast of Lower Saxony showed a very pronounced jump in wind direction from south to west in the morning of 20 August 1990. LMQ and LMK show this jump very well though one to two hours too late. Wind speed increased almost continuously behind the front until the following night. The magnitude and direction is similar for all DWD models, but LMQ and LMK show lower surface pressure which is closer to observations.

Figure 5 shows the wind speed and pressure at 03 UTC on 21 August 1990 at the Ems estuary. LMK is the only model where the east Friesian islands are resolved at least partly by the grid. The reduced wind speed on and downstream of the resolved islands is clearly visible. Also the acceleration of the wind from northwest to southeast in the Jadebusen Bay, on the east side of the section shown, can only be modelled by LMK.

Storm on 3 December 1999

Similar to the storm in August 1990, the storm on 3 December 1999 crossed the North Sea from Scotland to Denmark. However, it was very fast and traversed the North Sea in less than 12 hours. Sea level pressure dropped to 952 hPa. This was the strongest observed storm in Denmark (Rosenørn, 2000).

The small cyclone cannot be captured well with the ERA-40 resolution. This storm has been investigated using ECMWF's Ensemble Prediction System (EPS). Buizza & Hollingsworth (2001) showed that a high-resolution EPS (T255) predicted this storm much better than the then operational EPS (T159) which has the same resolution as the ERA-40 reanalysis.

Figure 6 shows the time series of pressure and wind at Helgoland. The minimum pressure occurred at 15 UTC on 3 December. The mean wind speed at 10 m at that time is given in Figure 7. The higher the resolution of the model the lower the minimum pressure of the storm centre. The strongest winds occur in the cold sector south and southwest of the centre. The models GME, LMQ and LMK predict approximately the same maximum speeds. The ERA-40 reanalysis and forecast give much weaker 10 m winds because the pressure gradient is weaker. In addition, the sea surface roughness, z_0 , is greater – over 5 cm compared to approximately 2 mm for the DWD models. LMK is the model which best captures the short backing of the wind at 23 UTC on 2 December and 18 UTC on 4 December.




Figure 4 Time series of sea level pressure (hPa), wind speed (ms⁻¹) and wind direction (degrees) observed at DWD station Norderney and predicted by the models for the storm on 20 August 1990.

Value of high-resolution hindcasts

The short comparison with observations shows satisfactory results. The high-resolution model LMK simulates much more detail (e.g. rain bands and gust streaks) than the other models. Also, it is the only one to resolve at least some of the Friesian islands off the coast of Germany and the Netherlands. Further verification will be done by NLWKN which will use the wind fields for coastal protection studies.

This study shows that a combination of the DWD models GME, LMQ and LMK, and the ERA-40 reanalyses provides the capability to produce high-resolution hindcasts that can be used to investigate extreme weather events from the past. Such hindcasts may be of value to many organisations that need to carry out studies to help reduce the impact of extreme weather.




Figure 5 Mean wind speed at 10 m in ms⁻¹ at 03 UTC on 21 August 1990 simulated by GME, LMQ, LMK, and the ERA-40 forecast (ERF) initialised at 12 UTC on 20 August 1990. The mean sea level pressure in hPa is shown by isobars. The Ems estuary is in the centre of the domain, the Jadebusen Bay is near the eastern boundary.




Figure 6 Time series of sea level pressure (hPa), wind speed (ms⁻¹) and wind direction (degrees) observed at DWD station Helgoland and predicted by the models for the storm on 3 December 1999.




Figure 7 Mean wind speed at 10 m in ms⁻¹ at 15 UTC on 3 December 1999 simulated by GME, LMQ, LMK, and the ERA-40 forecast (ERF) initialised at 00 UTC on 3 December 1999. The mean sea level pressure in hPa is shown by isobars.

We would like to acknowledge the contributions made to this study by Nils Kaiser and Katharina Klein from the University of Bonn, and Kai Sven Radtke from the University of Leipzig who performed the calculations and the first analysis.

FURTHER READING

Buizza, R. & A. Hollingsworth, 2001: Severe weather prediction using the ECMWF EPS: The European storms of December 1999. *ECMWF Newsletter No. 89*, 2–12.

Majewski, D., D. Liermann, P. Prohl, B. Ritter, M. Buchhold, T. Hanisch, G. Paul & W. Wergen, 2002: The operational global icosahedral-hexagonal gridpoint model GME: Description and high-resolution tests. *Mon. Wea. Rev.*, **130**, 319–338.

Rosenørn, S., 2000: De kraftigste storme i det tyvende århundrede i Danmark. *Vejret*, **82**, 15–18.

Steppeler, J., G. Doms, U. Schättler, H.W. Bitzer, A. Gassmann, U. Damrath & G. Gregoric, 2003: Meso-gamma scale forecasts using the nonhydrostatic model LM. *Meteorol. Atmos. Phys.*, **82**, 75–96.

Uppala, S.M., P. Kallberg, A. Simmons, U. Andrae, V. da Costa Bechtold, M. Fiorino, J. Gibson, J. Haseler, A. Hernandez, G. Kelly, X. Li, K. Onogi, S. Saarinen, N. Sokka, R. Allan, E. Andersson, K. Arpe, M. Balmaseda, A. Beljaars, L. van de Berg, J. Bidlot, N. Bormann, S. Caires, F. Chevallier, A. Dethof, M. Dragosavac, M. Fisher, M. Fuentes, S. Hagemann, E. Holm, B. Hoskins, L. Isaksen, P. Janssen, R. Jenne, A. McNally, J.-F. Mahfouf, J.-J. Morcrette, N. Rayner, R. Saunders, P. Simon, A. Sterl, K. Trenberth, A. Untch, D. Vasiljevic, P. Viterbo & J. Woollen, 2005: The ERA-40 re-analysis. *Q. J. R. Meteorol. Soc.*, **131**, 2961–3012.

Hurricane Jim over New Caledonia: a remarkable numerical prediction of its genesis and track

Thierry Lefort
Météo-France, New Caledonia

When communicating with members of the public, journalists and even scientifically educated users, we still encounter the popular belief that tropical cyclones are “capricious and unpredictable beings”. And yet using a set of global deterministic forecasts from vari-

ous centres and an ensemble of forecasts from a single centre now often give very valuable information, not only in the short range, but also into the medium range. Furthermore, during the last few years, global models have been able to generate increasingly realistic tropical cyclones. Consequently National Meteorological Services are now able to issue useful medium-range public forecasts for the tropics as well as mid-latitudes (Van der Grijn et al., 2004).

In January 2006, the information provided by the numerical models was clear enough to enable Météo-France's New Caledonian forecasters to mention in their outlook the high risk of tropical cyclogenesis in the Coral Sea (north-east of Australia). And it happened! This is the story of Hurricane Jim – a highly predictable tropical cyclone.

Medium-range bulletins in the tropics

Medium-range forecasts are no longer restricted to extra-tropical climates. Météo-France's New Caledonia Office has been producing a 5-8 day outlook twice a week since May 2005. This bulletin describes the expected weather pattern: southeasterly trades, easterly flow, westerlies, or shift of the InterTropical/South Pacific Convergence Zone (Lefort, 2004). It is mostly based on products from the ECMWF Ensemble Prediction System (EPS) – ensemble means, probability thresholds, EPSgrams, spaghetti charts etc. Also the French Polynesia Office of Météo-France issues a 5 day text bulletin which can be found at:

www.meteo.pf/previsions.php?carte=me

Another source of tropical medium-range forecasts is the Hydrometeorological Prediction Center at Camp Springs (US National Weather Service, NWS) which issues a daily extended outlook for Hawaii valid until day 8 – see the example in Box A and the information at:

www.hpc.ncep.noaa.gov/discussions/fxpa.html

In addition NWS's Samoa Office, Pago-Pago, issues a 6-day narrative synoptic discussion for US Samoa and the independent state of Samoa, as well as a 6-day marine forecast. These can be obtained from:

www.prh.noaa.gov/samoa/

Other tropical island states in the South West Pacific Ocean issue 5 to 7 day city or island forecasts.

Public outlooks for tropical cyclone usually cover a three-day period. Nevertheless, in some more technical documents, such as the example from Hawaii, there is an indication of favourable conditions for the potential development of tropical disturbances beyond that period.

When the rainy season comes, there are now times when global models, often supported by ensembles and statistical prediction methods, generate and/or move a tropical cyclone in an area covered by a tropical cyclone watch in the medium-range. Such a case happened in January 2006 in the South West Pacific Ocean: Hurricane Jim was predicted even before it existed.

Tropical Cyclone Jim in New Caledonia

In New Caledonia, a tropical cyclone (TC) watch is established as soon as a TC reaches tropical storm force (47 knots) in our area of national responsibility. Then, a warning at Level 1 is issued for the region which might be affected by the TC within 24 hours. A Level 2 warning means that the effect of a TC will be apparent within 6 hours.

When a TC has entered our national responsibility area, the major challenge for Météo-France's forecasters in New Caledonia is to deliver to the authorities the most accurate forecast track valid for 48/72 hours ahead.

Box A Example of a tropical medium-range forecast issued by the NOAA-NWS-Hydrometeorological Prediction Center

Hawaii extended forecast discussion NWS Hydro-meteorological Prediction Center Camp Springs MD 732 am EDT Tue Jul 04 2006 valid 00Z Wed Jul 05 2006 – 00Z Wed Jul 12 2006 models again cont to show a weak mid level trof E of the islands that extends well NWD to a large closed mid level low off B.C. retrograding to over the islands this week contg to give a break or cull in the east to west ridging S of 30N. This weakness fills by late week. At the SFC fairly brisk trades will prevail well into next week. High pressure strengthens to the N and NE of the islands. Weak tropical disturbances will track well SWD of the islands under unfavorable conditions for development. Models are showing more favorable conds in the EPAC for potential development next week. Rosenstein

Let us go back to Monday 23 January 2006: it is the rainy season in the South Pacific Ocean. Some rather cool air has just spread over New Caledonia after a heavy rain event in the Coral Sea. The monsoon trough is located around 15°S: to the north there are light to moderate north-westerlies with light to moderate easterly trade winds to the south. A small-scale convective cluster can be seen long the Queensland coast on the trough, but nothing impressive. In New Caledonia, the weather is expected to be very pleasant during the coming week.

For several runs now, most global models were consistently generating a TC in the Coral Sea, off Cairns. The development was supported by the ECMWF EPS. Even more exciting, the models predicted that once the TC formed it




Figure 1 T+168 h forecast of the mslp from the ECMWF T511 model, based on 12 UTC run on 23 January 2006.




Figure 2 T+156 h forecasts of the probability of mslp less than 1000 hPa from the ECMWF EPS, based on runs from (a) 00 UTC on 24 January and (b) 00 UTC on 25 January 2006.




Figure 3 T+156 h forecast of locations of mslp minima less than 995 hPa from the ECMWF EPS, based on the run from 00 UTC on 25 January 2006.

would then move quickly towards New Caledonia as it intensified. Figure 1 shows the T+168 h deterministic forecast from the ECMWF T511 model. There is an indication of a modest TC with a centre of 996 hPa located north of New Caledonia. The probability fields of sea level pressure from later runs (Figures 2) show that the EPS supports this scenario: 35 out of the 50 runs show a sign of TC over or near New Caledonia (Figure 3). It really does seem that New Caledonia is the target of future TC Jim!

These developments were associated with an active phase of a Madden-Julian-Oscillation (MJO) approaching from the west. The MJO has been proved to strongly influence TC genesis in the South Pacific Ocean. Mostly based on the sea surface temperature and MJO, weekly statistical predictions of TC genesis and occurrence are made for a large domain roughly from the coast of Queensland to Fiji (Leroy, 2004). The results show a peak of probability of about 60% compared with a value of about 40% from climatology (Figure 4a), and a chance of the presence of a TC of up to 80% with the climatological value being about 60% (Figure 4b). Also note that this peak – although not as sharp – was present in the week 3 forecast (blue lines in Figures 4a and 4b).

With such a consistent signal in the NWP products, the duty forecasters could expect a high risk of a TC watch being required for next Sunday or Monday, and a significant risk of a Level 1 or Level 2 warning for the Monday or Tuesday.

On Monday 23rd, in the 5–8 day outlook bulletin, the favourable situation for the formation of a TC in the Coral Sea was mentioned. Subsequently a vigorous westerly monsoon surge took place over the Arafura Sea on Tuesday 24th. This was captured by the models.

On Wednesday 25th, even though there is no operational medium-range warning system in place, the Head of Forecasting Division in New Caledonia informed the authorities by phone that they should be aware of the likelihood of a TC threatening the Territory by the end of the week.

The westerlies reached the Coral Sea on Thursday 26th. It is likely that this inflow contributed to bringing low-level vorticity into the old low east off the Queensland coast. The convection developed suddenly, and at 00 UTC on Saturday 28th the disturbance was designated as Tropical Cyclone Jim by Brisbane RMSC. This was quite an emotional moment for the forecasters: the TC had formed where predicted by the NWP models five to six days ahead. At that time, the track appeared to be much more likely to be to the east of the main island (Figure 5), and the forecast tracks from different Meteorological Services were in good agreement.

At 06 UTC on Monday 30th, Jim was already a hurricane, with wind speeds estimated to have reached 80 kn and with a central pressure of 955 hPa. In the morning the TC watch was established. As a result there were requests for more information from radio and TV stations and newspapers. In the evening, a Level 1 warning was issued for the north of the main island and the Loyalty islands, and it was planned to upgrade the warning to Level 2 on Tuesday 31st at 0600 local time. The forecaster who watched Hurricane Jim on the satellite picture still had in mind the model charts he had seen more than 7 days earlier: “models were incredibly

right!”. While everybody could enjoy the calm, pleasant, and rather cool weather at a time when there was no hint of TC threat of any kind on the satellite pictures and short-range forecasts, the forecasters had the medium-range NWP products that predicted the formation of a TC which would threaten New Caledonia seven days later.

Eventually, Jim weakened on Tuesday 31st and it passed, as predicted, 50 km to the east of the Loyalty islands (Figures 6 and 7). There was only very little damage on most exposed areas of New Caledonia and the islands nearby.

It is not only important to predict the track – predicting the intensity is also important. However, for Hurricane Jim it was difficult to predict the intensity, even at 24 hour lead time. Even intensity analyses from different centres differed significantly. Clearly the consistency in intensity forecasts is rather poor.

It seems that the situation was highly predictable, probably because of the well-analysed vigorous westerly flow that spread over the Coral Sea.

The accurate forecast of the formation and track of Hurricane Jim was an important event in the development of tropical numerical prediction: advanced warning of the formation of a TC was given seven days ahead.

Discussion and perspectives

What is the risk of the presence of a tropical cyclone within a certain domain, or along a certain route, during the coming week, or the coming two weeks? This is a frequent question asked by sailors who need five to seven days to cross the Coral Sea from Noumea to Cairns. All kinds of vessels need this kind of information. As has been shown with Hurricane Jim, NWP models and the associated ensemble forecasts can sometimes provide valuable information which can answer questions of that kind.

The predicted formation of tropical cyclones is one of the many aspects that will benefit from the multi-ensemble approach that is being developed within the TIGGE component of WMO’s THORPEX programme.




Figure 4 Predicted probability of (a) TC genesis and (b) TC occurrence in the South West Pacific Ocean based on forecasts made each Monday (black: week 1; red: week 2; blue: week 3). The last point of each line shows the prediction made on 23 January 2006. The smooth curve is the probability according to climatology.

Probability that Jim will pass within 120km radius during the next 120 hours




Figure 5 Probability that Jim will pass within 120 km radius during the next 120 hours plus the individual tracks based on the ECMWF EPS run from 12 UTC on 28 January 2006.




Figure 6 MTSAT IR picture of Jim at 21 UTC on 30 January 2006 (source Météo-France).




Figure 7 Analysed track of TC Jim with TC watch area indicated by the green square (source Météo-France).

Even if the forecast is increasingly precise when shifting from medium to short range, there is still often much uncertainty about track and intensity at 24 hour lead time. However, there might be some decision-making processes or mitigation actions for which advanced warnings five to seven days ahead based on probabilistic information may be of value, especially if a high false alarm rate can be tolerated. It is clear that the provision of advanced warnings going into the medium range should be considered in order to help mitigation actions. But some questions remain:

- ◆ Should we validate the formation of tropical cyclones by the models?
- ◆ What should the forecasters do with the information about tropical cyclones available from the models?

FURTHER READING

Lefort, T., 2004: Starting up medium-range forecasting in New Caledonia. *ECMWF Newsletter No. 102*, 1–7.

Leroy, A., 2004: Statistical prediction of the weekly tropical cyclone activity in the southern hemisphere. *Rapport de stage de fin d'études*, 939, Météo-France, Met College-ENM, Toulouse, France.

Van der Grijn, G., J.E. Paulsen, F. Lalaurette & M. Leutbecher, 2004: Early medium-range forecasts of tropical cyclones. *ECMWF Newsletter No. 102*, 7–14.

Ice supersaturation in ECMWF's Integrated Forecast System

Adrian Tompkins, Klaus Gierens, Gaby Rädel

As well as being a key forecast product themselves, clouds are also important for the ECMWF forecasts since they determine the precipitation falling to the surface and are key modulators of the radiation budget. Accurate prediction of the two-metre surface temperature would be impossible without a reasonable representation of clouds. In the early days of ECMWF the key goal regarding clouds was simply to represent their basic macroscopic properties such as how much of a grid-cell they cover and how much liquid and ice they contain. Now that ground-

based and satellite validation shows that this task has been fulfilled with a certain degree of success, it is possible to consider refining the cloud model to include neglected physics. One of these refinements is the inclusion of ice supersaturated states.

The meteorology enthusiast will know from their own amateur weather stations at home that relative humidity (RH) can change dramatically from day to day, but that it is always subject to an upper limit of 100%. This is because the atmosphere has a plentiful supply of aerosols, known as cloud condensation nuclei, upon which any vapour in excess of this limit can quickly and efficiently condense to form (or “nucleate”) cloud drops.

The situation is not the same for cloud ice crystals at very cold temperatures such as those found in the upper troposphere. Here much higher values of RH, significantly exceeding 100%, must be attained before an ice crystal can be nucleated. This implies that airmasses can become super-

Affiliations

Adrian Tompkins: ECMWF, Reading, UK

*Klaus Gierens: Institut für Physik der Atmosphäre,
DLR Oberpfaffenhofen, Germany*

Gaby Rädel: Department of Meteorology, University of Reading, UK

Box A Ice nucleation mechanisms

Ice particles in cirrus can form by both homogeneous and heterogeneous nucleation.

Homogeneous nucleation

Homogeneous nucleation is a process where an ice crystal is formed from a liquid droplet. Pure water droplets can only exist down to approximately -38°C ; below this limit they freeze spontaneously. In the ECMWF model, the present temperature-based diagnostic division of cloud water between liquid and ice in the mixed phase necessarily restricts the use of the new scheme to the assumed pure ice phase at temperatures below -23°C . Below this temperature, if a moistening clear air-mass attains the liquid water saturation mixing ratio, the forming liquid water droplets are assumed to freeze instantaneously. Since the liquid water saturation exceeds that for ice, the air mass will be ice-supersaturated, and the excess humidity will cause the rapid depositional growth of the newly formed ice crystals. The mismatch between the -23 and -38°C temperature thresholds and the treatment of ice nucleation in mixed phase clouds will be addressed in a planned future implementation of a separate prognostic ice variable.

The above picture is complicated by the presence of aqueous solution droplets (haze particles). The foreign molecules impede the formation of an ice crystal lattice and thus haze particles can remain unfrozen at very cold temperatures. If relative humidity increases, the foreign molecules become increasingly rarefied (dilute) by the uptake of water molecules. At high supersaturations, this dilution can be sufficient to allow the formation of a germ crystal lattice and the haze particle freezes. At very cold temperatures, this may occur at humidities significantly

lower than the liquid water saturation mixing ratio. Thus in the ECMWF scheme we adopt a threshold function for ice nucleation that is the minimum of the liquid water saturation mixing ratio and the analytical form for the freezing threshold of aqueous solution drops given by *Kärcher & Lohmann* (2002). For a temperature of -38°C a RH of 145% is necessary to nucleate ice crystals homogeneously, while at -83°C this threshold increases to approximately 165% (or equivalently a supersaturation of 65%).

Heterogeneous nucleation

The alternative process, heterogeneous nucleation, relies on solid particles to initiate the freezing event by contact with or immersion in a supercooled water droplet, or to act as a nucleus upon which ice forms by direct deposition of water vapour. The subset of aerosols that have these properties are termed ice nuclei. The ice nucleating properties of aerosols are a function of the ambient relative humidity, and aerosols usually become “activated” (take on ice nucleating abilities) at significant supersaturations; a typical assumed threshold in modelling studies being 30%. While this threshold is lower than the threshold for homogeneous ice nucleation, heterogeneous ice nucleation is considered to play a more minor role in cirrus formation since ice nuclei are much rarer than cloud condensation nuclei. However, if ice nuclei are present in sufficient numbers, heterogeneous nucleation has the potential to dominate as the cloud forming pathway.

The IFS model currently represents basic aerosol types as a monthly-mean climatology. In the near future, the ongoing GEMS project will result in a daily analysis and ten-day forecast of many different aerosol species which would potentially allow the heterogeneous nucleation pathways to be locally represented.

saturated with respect to the ice saturation mixing ratio. Ice nucleation is complicated by the fact that it can occur both homogeneously and heterogeneously (see Box A for details). This article only considers the former process of spontaneous freezing of liquid water or aqueous solution drops.

The result of this high nucleation threshold is that it is quite common to have upper-tropospheric air masses that are supersaturated with respect to the ice saturation mixing ratio. The presence of such layers has been documented both in ground launched sonde soundings, in situ aircraft measurements and also from satellite-based instruments. But if you want to start your own observational database of supersaturated layer occurrence do not despair; you do not need to build your own satellite! If you live near a busy flight corridor, you only need to observe the behaviour of the contrails (the linear condensation clouds that high altitude aircraft produce) to determine the state of the upper atmosphere.

A high-flying aircraft producing no contrail, or a contrail that disperses, is a clear sign that the ambient air through which it is flying is subsaturated, causing the sublimation of the ice crystals from the jet exhaust as they mix with the ambient air. On the other hand, permanent contrails tell you that the ambient air is supersaturated. Just such a database of contrail observations has been built up by researchers at the University of Reading and will be used later in this article. Sometimes these contrails can spread out to form cirrus cloud of substantial coverage, significantly altering the local radiative balance.

Many models, including the ECMWF operational forecast model, do not allow supersaturation to occur – they treat

ice cloud nucleation in the same fashion as liquid clouds. Why is it desirable to correct this and represent this physics more accurately? One reason is that models which neglect this will obviously suffer from a dry bias in the upper atmosphere, but perhaps more importantly they are likely also to overestimate the incidence of upper-level cirrus clouds.

A new supersaturation scheme

A new scheme to represent ice supersaturated layers has been introduced into model Cycle 31r1 of the ECMWF Integrated Forecast System which became operational on 13 September 2006. The new scheme only affects the forecasts for the moment, and the analysis system is still prevented from producing supersaturated states.

The scheme is based on three assumptions:


- ◆ Ice nucleation initiates when the supersaturation locally reaches the homogeneous nucleation threshold which is specified as a function of temperature and can reach values exceeding 170% (See box A for details).
- ◆ The clear-sky humidity fluctuations are uniformly distributed with a fixed constant variance. Thus nucleation can occur when the grid-mean RH exceeds a threshold that is lower than the local criterion.
- ◆ Once ice is present, the ice growth (deposition) is sufficiently rapid relative to a model time-step that it can be approximated by an instant adjustment to exactly saturated conditions inside the cloud.

We contrast the new scheme and the operational ECMWF scheme in Box B. The new scheme is simple, and cannot represent the presence of ice crystals existing in subsaturated

Box B

Panel (a) is a schematic (adapted from *Kärcher & Lohmann, 2002*) of the evolution of relative humidity in a hypothetical parcel subject to adiabatic cooling at cold temperatures ($T < 232$ K). The green dotted line in panel (b) indicates the approximation of this process by the current operational ECMWF IFS model. The model does not allow RH to exceed 100%, and any excess humidity is instantaneously converted to ice. The impact of the new supersaturation scheme is shown by the green dotted line in panel (c). This scheme allows supersaturation to occur and thus represents

the hysteresis in the ice cloud behaviour. However, the scheme does not attempt to explicitly model the nucleation and depositional growth processes, and converts all humidity exceeding the saturation value to ice instantaneously once the nucleation threshold is attained. It should be emphasised that the scheme attempts to ascertain how much of the grid box has reached this threshold and the adjustment only occurs inside this cloudy region of the gridbox. Thus grid-mean supersaturated states can be maintained even in gridboxes that contain some ice.



Box C Numerical considerations

The deposition timescale, the time taken for the ice crystal growth to reduce the *RH* back to 100% after cloud initiation, can vary from a few minutes to a few hours and depends on the number of ice crystals nucleated. This number in turn can vary by orders of magnitude, and depends critically on the “overshoot” of the *RH* past the critical threshold for ice nucleation (see box B), both in terms of the magnitude of the excess achieved and the period for which *RH* exceeds this threshold. This is measured in seconds, since once ice nucleation initiates, depositional ice growth prevents further significant increases in *RH*. Parcel models, which aim to accurately model cirrus nucleation by placing the emphasis on representing the microphysics in a simple dynamical framework, resort to using timesteps of fractions of a second, four orders of magnitude shorter than the typical IFS timestep!

The matter is further complicated by the fact that the overshoot is also governed by the updraught velocity; fast updraughts overshoot more and nucleate more ice crystals. However, the velocities that matter occur on the cloud scale and are not resolved even by the T799 operational model. While recent efforts have successfully derived analytical approximations to the nucleation process using timesteps of many minutes, these must still rely on uncertain parametrizations for the cloud-scale vertical velocity spectrum.

conditions, nor supersaturated conditions within clouds. However, the simplicity is necessary due to the long timesteps between 12 minutes and one hour that the various operational configurations use. Box C provides more details on this aspect of the numerical complications.

Impact on aspects of the model climate

Two 7-member ensembles of 13-month model integrations at T159 (approximately 125 km equivalent resolution) are conducted to investigate the effect of the new scheme on the model climate. For reasons of brevity, we only examine the direct impact on the humidity and ice cloud properties

here, although there are obviously further ramifications for the simulated hydrological cycle and radiative budgets. We discard the first month of each integration, and average statistics over the last 12 months of each ensemble member, giving a total of seven years. The control integration uses the cloud scheme of model Cycle 31r1, with the supersaturation scheme switched off.


In terms of the model cloud climatology, the largest sensitivity was not in the total column ice water (TCIW, sometimes referred to as the ice water path) but rather the total cloud cover. High cloud cover is reduced by the new scheme, as expected, since humidity must now attain higher values to initiate clouds (Figure 1). Compared to the relative reduction in high cloud cover of 15% (or 6% absolute decrease), the global reduction in TCIW is much more modest at only 3.5%.

The lower sensitivity of the ice water content relative to that of the high cloud cover is easy to understand. Since the new scheme has a higher (but correct) threshold for ice nucleation, cirrus clouds will occur more rarely. However, if humidity increasing in a gridbox does manage to attain this higher nucleation threshold, all of the water vapour exceeding the saturation mixing ratio is instantly converted to ice. Thus it is clear that the difference in sensitivity between cloud cover and cloud ice content is a result of this “hysteresis”: ice clouds that are relatively optically thick in the default model will also be thick in the new model (but with a delayed onset), while some thinner clouds will simply not occur.

The impact of the supersaturation scheme on the humidity field is shown as a zonal mean difference in *RH* in Figure 2. As expected the scheme increases *RH* in the upper troposphere and also influences the lower stratosphere through increased cross-tropopause transport. The absolute increases in zonal mean *RH* are mostly in the range of 5 to 10%, with the peak of approximately 20% occurring in the troposphere/stratosphere transition zone above the main deep convective detrainment level in the tropics.

The net effect of the changes in cloud properties in the tropics (latitudes less than 30°) is an increase in net outgoing longwave radiation of 2.1 Wm⁻² and a similar compensating increase of 2.2 Wm⁻² in net incoming shortwave radiation, while the corresponding global figures are 1.3 and 1.5 Wm⁻², respectively.

Figure 1 Difference in high cloud cover (pressure < 450 hPa approximately) between experiments using the new supersaturation scheme and the control, respectively, based on 7-member ensemble mean 12-month averages.






Figure 2 Difference in zonal mean relative humidity between experiments using the new supersaturation scheme and the control (shaded contours) and the zonal mean relative humidity in the control forecast (line contours with 5% intervals) based on the 7-member 12-month averages.

Comparison to aircraft observations

The humidity fields in the climate integrations are compared to aircraft data from the MOZAIC campaign, which compiled a number of years of in situ aircraft measurements taken by commercial aircraft carrying research quality instrumentation (see *Gierens et al.*, 1999). The model integrations are validated in terms of normalized probability density functions (PDFs) of RH (with respect to ice), which are




Figure 3 Normalised probability distribution function of RH at the 250 hPa level for the northern hemisphere for the control and new supersaturation scheme using the first member from each set of climate integrations compared to aircraft measurements from the MOZAIC campaign.

compared to the MOZAIC climatological PDFs. Results are shown for the 250 hPa level for the northern hemispheric latitudes greater than 30° (Figure 3).

The control behaves as expected, with a sharp peak (or “mode”) around $RH = 100\%$ due to the fact that this limit can not be exceeded in cloudy air masses. There is a slight overshoot with RH attaining values up to almost 10% supersaturation due to numerical reasons.

After the smaller 100% mode, the new scheme produces a PDF of humidity that follows the best-fit line to the observations made in the MOZAIC campaign. However, at a supersaturation of around 30% the PDF becomes steeper than the observations, that is, relative to the frequency of occurrence of 100% RH , fewer grid-points attain higher values of RH . This scale-break is a result of the model taking into account the variability of humidity on scales smaller than the grid-scale. Consequently when the grid-mean relative humidity has a value of 130%, the assumed subgrid fluctuations imply that some part of the gridbox has a RH of 150%, perhaps sufficient to initiate nucleation, depending on the temperature. The ice cloud formation reduces humidity, hence leading to the steeper humidity PDF.

Comparison to microwave limb sounder retrievals

Another source of information concerning ice supersaturation is provided by microwave limb sounder (MLS) retrievals (see *Spichtinger et al.*, 2003). Although the horizontal and vertical resolution is coarser than the in situ aircraft data, the retrievals have the advantage of global coverage, in particular allowing inspection of the tropical upper troposphere. We compare the model with the supersaturation scheme to an MLS retrieval climatology obtained between September 1991 and

June 1997. The retrievals are gridded with a horizontal resolution of approximately 300 km, and have a vertical resolution of 3 km. We only use the pressure levels of 147 hPa and 215 hPa due to the reduction in reliability lower in the troposphere.

The MLS retrievals at 147 hPa shown in Figure 4(b) give the maximum frequency of occurrence of supersaturation occurring in the tropics, as expected, with peaks over the deep convective regions of the tropical Western Pacific, Africa and the Americas. Figure 4(a) shows that the model

reproduces the general frequency of occurrence statistic well, with peak values in the tropics and the secondary peak over the South Pole. The model biases given in Figure 4(c) can be summarised as a slight underestimation of ice supersaturation frequency over the Western Pacific, a corresponding overestimation over the Indian Ocean and the warm SST tongue of the Eastern Pacific, and a significant underestimation over continental Africa and the Americas.




Figure 4 Annual mean frequency (%) of occurrence of ice supersaturation for (a) the model level at 144 hPa for the seven-member climate ensemble using the new supersaturation scheme and (b) the MLS retrieval climatology at 147 hPa. (c) Difference between (a) and (b).

	147 hPa		215 hPa	
	Model	MLS	Model	MLS
Global	7.8	10.5	8.6	4.3
Northern Hemisphere (lat > 30°)	0.7	0.9	5.9	2.2
Tropics (-30° < lat < 30°)	14.4	18.9	11.0	6.0
Southern Hemisphere (-55° < lat < -30°)	0.5	0.3	6.3	2.5
Antarctic (lat < -55°)	2.2	6.1	6.0	6.1

Table 1 Frequency of occurrence of ice supersaturation in percent for the model using the supersaturation scheme and MLS retrievals, which are taken from *Spichtinger et al. (2003)*, at pressure levels of 147 hPa and 215 hPa. The model results are taken at the nearest model levels at pressures of 144 and 211 hPa, respectively.

The striking aspect of the ice supersaturation frequency biases is that they reflect precisely the known biases in the model's over- or under-estimation of tropical deep convection. Comparison of the model's rainfall with data from the Global Precipitation Climatology Project (GPCP), top of atmosphere infrared and solar net fluxes with CERES, and ice mixing ratios with MLS retrievals, all point to the same major errors in the tropics; in particular the underestimation of deep convective activity over tropical America and Africa. This indicates a mechanism for creating the tropical ice supersaturated regions whereby deep convection locally transports water vapour to the region of the radiative tropopause (where radiative heating rates are zero) which is then transported upwards towards the temperature tropopause by slow mean ascent which balances the small net radiative warming. Where the model locally produces too little deep convection, the radiative tropopause is too dry, reducing the instance of ice supersaturated regions in the tropopause transition layer.

Table 1 shows the model and data statistics for both 147 hPa and 215 hPa to reveal the model's ability to predict the geographical variability region by region of the supersaturation occurrence.

Comparison to contrail observations

If the new supersaturation scheme improves the representation of supersaturation in the humidity field at cold temperatures, more accurate predictions of persistent contrail occurrence should result. We test this conjecture using a six-month database of contrail observations, taken over Reading in Southern UK; weather permitting of course! Reading is particularly well suited to establish such a contrail database due to its geographical location at the entrance of the North-Atlantic flight corridor.

The contrail observations are divided into two categories: 'yes'-events, where permanent/non-dissipating contrails were observed indicating supersaturated conditions, and 'no'-events, where contrails either dissipated or were not observed. These observations are compared to T511 model integrations at a forecast range of 20 to 30 hours to coincide with each

Control			New supersaturation scheme		
		Observed		Observed	
Predicted	Yes	No	Predicted	Yes	No
Yes	31	38	Yes	41	28
No	12	71	No	16	67

Table 2 Contingency tables for comparisons of visual contrail observations with predictions using the control (left) or the new supersaturation scheme (right).

contrail observation. The corresponding model event is whether or not $RH > 99\%$ is predicted anywhere through the standard upper-tropospheric flight levels. For instance, if the model predicts $RH = 120\%$ at 300 hPa and a permanent contrail is observed, this is considered a 'hit'. The control model without the new scheme predicts permanent contrails if the gridbox is saturated, that is $RH = 100\%$.

From the contrail information we form a 2×2 contingency table, yielding the number of correct hits (a), false alarms (b), number of misses (c) and lastly the number of correct rejections (d). We test the model forecasting skill using the 'odds ratio', which is the ratio of odds of making a hit given that the event occurred to the odds of making a false alarm given that the event failed to occur. In terms of the entries in the contingency table the odds ratio is given by $(a \times d) / (b \times c)$, with larger values representing improved performance. We find that the odds ratio increases from 4.8 to 6.1 when the supersaturation scheme is used (see Table 2), which is a statistically significant improvement.

Summary and future work

We have presented results from a simple new scheme that represents ice supersaturation and the homogeneous ice nucleation process. Using this scheme the ECMWF forecast model allows supersaturation with respect to ice for the first time.

We have attempted to validate the new scheme by comparing relative humidity related model fields to in situ aircraft measurements, remotely sensed data, and ground-based permanent contrail observations. These all showed that the scheme has the ability to predict the location, magnitude and frequency of supersaturation events with reasonable fidelity.

Some readers may have noted an apparent omission from the validation exercise; namely a systematic comparison to the radiosonde network. The reason for this was the current lack of a general humidity bias correction methodology for the diverse sonde types assimilated into the ECMWF system. In addition to the various Vaisala sonde types, ECMWF also currently ingests VIS sondes, and the various platforms made in France, China, Russia, India and Japan. While some of these platforms have documented bias characteristics, the Vaisala RS80 dry biases being a good example, many do not. ECMWF is thus currently developing a methodology for systematic humidity bias correction of the radiosonde

network. This will not only help by providing an additional source of information for forecast validation exercises, but will also improve the quality of the forecasts themselves through an improved humidity analysis.

FURTHER READING

Gierens, K. U. Schumann, M. Helten, H. Smit & A. Morenco, 1999: A distribution law for relative humidity in the upper troposphere

and lower stratosphere derived from three years of MOZAIC measurements. *J. Geophys. Res.*, **105**, 22,743–22,753.

Kärcher, B. & U. Lohman, 2002: A parameterization of cirrus cloud formation: Homogeneous freezing of supercooled aerosols. *J. Geophys. Res.*, **107**, DOI: 10.1029/2001JD000470.

Spichtinger, P., K. Gierens & W. Read, 2003: The global distribution of ice-supersaturated regions as seen by the Microwave Limb Sounder. *Q. J. R. Meteorol. Soc.*, **129**, 3391–3410.

New features of the Phase 4 HPC facility

Neil Storer

A earlier news item in this newsletter entitled “*The new IBM Phase 4 HPC facility*” starting on page 5 described the new IBM Phase 4 system that is currently being installed and commissioned at ECMWF. The aim of the current article is to look in more technical detail at three new features of this system: simultaneous multi-threading, mid-size memory pages and multi-cluster GPFS.

It is important that users of ECMWF’s HPC facility are aware of these new features as it will allow them to gain the full benefits of the new system. For example, effective use of Phase 4 will allow a higher percentage of peak performance to be sustained.

Simultaneous multi-threading (SMT)

SMT is probably the most noticeable new feature of the POWER5 micro-architecture; around 20% of the additional transistors on the chip (over and above those on the POWER4 chip) are dedicated to this feature. SMT enables two threads to execute concurrently on a single processor (physical CPU), i.e. in any given CPU cycle it is possible for each of the two threads to execute an instruction on that processor. The effect is to make every processor look as if it were in fact two virtual CPUs. This in turn enables the processor to be kept busy for a larger proportion of the time, thus producing a higher *sustained* performance than would otherwise be the case. However, because only certain parts of the hardware (such as the program counter, instruction buffer and rename registers) are duplicated, but none of the functional units, the peak performance of the processor is not affected by SMT. The processor can intelligently assign priorities to the two threads to dynamically balance access to the shared resources, thus ensuring that one thread does not “hog” these thereby starving the other thread of CPU cycles.

It is difficult to gauge the benefit in performance that SMT brings. Some applications, in particular those that may be limited by being functional-unit bound, or memory-bound, may not see any benefit from using SMT. It is possible to switch the chip into single-threaded mode and by doing so allocate all of the rename-registers that are shared in SMT mode to that thread. In theory this should enable some applications that do not benefit much from SMT to achieve higher levels of performance, but in practice this does not

appear to be the case for any of the applications that have been tested in this way at ECMWF.

A p5-575+ server can be run either with SMT functionality enabled or with it disabled. When SMT is disabled, each of the 16 processors can execute only a single thread, and as such, at any instant in time, each thread is guaranteed to run on a processor which is executing no other threads. When SMT is enabled, each of the 16 processors can execute two threads concurrently and as such, at any instant each thread may be sharing a processor with another thread (possibly from a different, unrelated job). In the first instance the operating system has 16 virtual CPUs on which to schedule threads for execution, while in the second instance it has 32 virtual CPUs.

ECMWF has decided that all of the p5-575+ servers in the two Phase 4 clusters will be run with SMT enabled at all times. This simplifies the operation of the clusters and even when there are only 16 threads associated with user jobs in a node, the additional virtual CPUs can be used by the operating system and system daemons that need to execute from time to time; in this way they do not have to interrupt user threads by switching execution context in order to gain access to the processor, reducing the effect of what is termed “operating system jitter”.

Use of SMT by parallel applications

Parallel applications can make use of SMT in several ways. Take for instance a parallel application that has 64 threads (MPI tasks, or a mixture of MPI tasks and OpenMP threads) which runs using 64 processors on one of ECMWF’s POWER4+ clusters using 2 nodes with 32 threads/node on that cluster. One can envisage several scenarios for running this on the POWER5+ cluster. Consider three scenarios and their impact on running two NWP applications from Member States.

1) **2 nodes, 32 threads/node.** If this were to run on the same number of nodes of the POWER5+ cluster it would still be using 64 CPUs, but this time they would be virtual CPUs. That is the application would be making use of SMT (2 virtual CPUs per processor) so that it would only be using 32 processors (not 64 as it did on the POWER4+ system). With half of its threads sharing the processors with the other half, it is likely run slower than on the POWER4+ cluster. However the slowdown may be much less than one would surmise and in fact in some

cases it might indeed run faster. We have seen NWP applications run in this way on the POWER5+ system that have taken much less than 10% longer to run than on the POWER4+ system.

- 2) **4 nodes, 16 threads/node.** Another way of running this application on the POWER5+ cluster would be to go back to using 64 processors but not using the SMT feature. This would almost certainly run faster than on the POWER4+ cluster, mainly owing to improvements in the memory access capability of the POWER5+ system and the fewer interruptions to the execution of application's user-threads needed in order to accommodate kernel and daemon access to the processor. Nevertheless, in this instance, no use would be made by the application of the additional virtual CPUs. In this respect the application would be wasting CPU cycles that it could otherwise make use of. The NWP applications ran between 1.4 and 1.6 times faster than on the POWER4+ system.
- 3) **4 nodes, 32 threads/node.** Yet another way of running the application on the POWER5+ cluster would be still to use 64 processors (4 nodes), but this time doubling the number of user-threads, to use all 32 virtual CPUs in each node. The gain in doing this mainly depends on the scalability of the application, but should in general be greater than that using the second paradigm. The NWP applications ran between 1.6 and 1.8 times faster than on the POWER4+ system.

Table 1 summarises these scenarios for the two NWP applications.

The question of how much gain an application will experience using SMT is a difficult one. In general this question can only be answered empirically by running the application twice, once using SMT and once without. For the two NWP applications one can see a gain of over 10% in using scenario (3) over scenario (2); for the IFS model gains in the region of 20% have been seen. Besides the wall-clock time taken to run the application, the user should also take into consideration the relative cost of the two runs (see the next section).

SMT and job-scheduling

The ways that users can make use of SMT inevitably complicates matters for job-scheduling, particularly since certain

features in LoadLeveler (IBM's batch scheduler) have yet to be fully implemented. In order to cater for the various scenarios some ground-rules have been set for the Phase 4 clusters. It is clear that jobs which prevent resources from being fully utilised should pay for those resources, so for the Phase 4 system ECMWF's LoadLeveler job-submit filter and charging algorithm have been modified to take this into consideration. The rules are as follows.

- 1) Any serial or parallel job that a user submits with the LoadLeveler directive:
#@ node_usage = not_shared
will not share the node(s) allocated to it with any other job. Also, regardless of the memory requirement specified by the job, it will be given all of the available memory in the nodes. Furthermore, regardless of the CPU requirement specified by the job, it will be accounted as if it used all of the processors in the node(s), since no other jobs can make use of any CPUs (virtual or physical) that it does not actually use.
- 2) Any serial job that specifies the LoadLeveler directive:
#@ node_usage = shared (the default)
or does not specify:
#@ node_usage = not_shared
will run in a node where up to 31 other serial jobs may be executing. This job will use a virtual CPU and may end up from time to time sharing a processor with other serial jobs in the node and consequently may have inconsistent wall-clock times each time it executes.
- 3) Any parallel job that uses more than one node will be forced to run as if the user had specified the directive
"#@ node_usage = not_shared"
and regardless of the memory requirement specified by the job, it will be given all of the available memory in the nodes. In this way, if a user decides to execute only 16 user-threads on a node, the 16 remaining virtual CPUs will not be allocated to any other job and consequently its wall-clock time should be consistent each time it executes. Regardless of the CPU requirement specified by the job, it will be accounted as if it used all of the processors in the nodes, since no other jobs can make use of any CPUs that it does not use.
- 4) Any parallel job that uses a single node and specifies the ECMWF site-specific pseudo LoadLeveler directive:

Application	System	Scenario (SMT use)	Nodes used (processors)	Threads per node	Total user-threads	Wall-clock time (seconds)	Gain over POWER4+
X	P4+		4 (128)	32	128	196	
	P5+	1 (yes)	4 (64)	32	128	210	-7%
		2 (no)	8 (128)	16	128	139	1.41x
		3 (yes)	8 (128)	32	256	122	1.61x
Y	P4+		2 (64)	32	64	1950	
	P5+	1 (yes)	2 (32)	32	64	2022	-4%
		2 (no)	4 (64)	16	64	1185	1.65x
		3 (yes)	4 (64)	32	128	1076	1.81x

Table 1 The performance of two specific applications (labelled X and Y) under different scenarios. These applications are NWP models developed by two of ECMWF's Member States.

#@ ec_smt = yes (the default)

or does not specify:

#@ ec_smt = no

will run in a node where up to 32 user-threads may execute. Such a job may end up from time to time sharing the processors assigned to it with other parallel jobs in the node and consequently may have inconsistent wall-clock times each time it executes.

LoadLeveler does not recognise this (or any other) ECMWF site-specific pseudo-directive; it is interpreted by ECMWF's job-submit filter and is not passed on to LoadLeveler itself. Be aware that if users put such ECMWF pseudo-directives in jobs run elsewhere, Loadleveler running at that site will reject the job.

- 5) Any parallel job that uses a single node and specifies the pseudo LoadLeveler directive:

#@ ec_smt = no

will run in a node where at maximum 16 user-threads may execute concurrently on the 16 processors. In this way such a job will not share the processors assigned to it with any other jobs and consequently its wall-clock time should be consistent each time it executes. The job will be accounted at a higher cost than if this directive had been omitted, since virtual CPUs are essentially forced to be idle by this job.

The following rule has been implemented to try to take into account those jobs that ask for so many (but not all) resources (CPUs or memory) in a node that they effectively prevent other jobs from using what is left over.

- 6) Any job that specifies a requirement for more than 75% of the memory in a node (i.e. more than 18750 MB), or that specifies 29, 30 or 31 (virtual) CPUs in a node (or 15 CPUs if “#@ ec_smt = no” is specified) will be forced to run as if the user had specified the LoadLeveler directive “#@ node_usage = not_shared” and regardless of the memory requirement specified by the job, it will be given all of the available memory in the node(s).

The following case, which is based on actual runs of a CPU-bound job (a coupled ocean + IFS model), is an example of how accounting differs for the scenarios in rules (4) and (5).

- ◆ Four copies of an 8-way parallel job were run on a single node with the default: “#@ ec_smt = yes”.
- ◆ The same four jobs were then run on two nodes using “#@ ec_smt = no”.

The wall-clock times, CPU times and the relative billing units charged are shown in Table 2. It is interesting to note that 25% more CPU time was used in the case of “ec_smt = no” than in the case where “ec_smt = yes” was used. This is due to the way in which the system itself accounts for use of the proces-

sor when two threads are executing on it at any one instant, as opposed to when only a single thread is running on it.

As you can see in Table 2, the user was charged over 40% more for reserving virtual CPUs and not using them. This is because in the one case the four jobs occupied a single node for 62,350 seconds, while in the other case the four jobs occupied two nodes for 44,070 seconds (i.e. 88,070 “node-seconds”) and the billing units express this ratio of the “node-seconds” used. The percentage will vary from job to job, depending upon their characteristics, but the message is clear – do not use “#@ ec_smt = no” unless it is absolutely necessary. This applies, for example, to time-critical work that would not finish in time if its threads were to share processors.

SMT and processor binding

To make the most out of SMT and memory affinity it is advisable to bind threads to CPUs. There are several ways to accomplish this, which are described in ECMWF's training material. It is difficult to give a general recommendation on the best CPU configuration to bind to, especially for applications using a mixed MPI and OpenMP programming paradigm; it mainly depends on the amount of work performed by each thread. In the same way as ascertaining how much gain an application will experience using SMT, the choice of which binding configuration to use is in general a question that can only be answered empirically by running the application with several configurations to see which best suits that application.

It is important that binding threads to virtual CPUs should only be done for those parallel applications that do not share a node. For applications that share a node it is difficult for users to know which CPUs to choose to bind to, because it is almost impossible to ascertain which CPUs the other applications running on the node are also bound to. For this reason it is essential not to bind threads to virtual CPUs for parallel applications that share a node. Rather let the operating system allocate the CPUs instead, otherwise it may happen that several applications overload the same processors within the node while other processors lie idle.

Table 3 summarises some experiments using a T399 resolution IFS model running a one-day forecast on 3 POWER5+ nodes (48 processors). These were performed early in the lifetime of the POWER5+ system. The IFS model is programmed using a hybrid MPI-OpenMP paradigm. In one set of experiments four OpenMP threads were used in each MPI task, while in the other set OpenMP was not used.

	Max user-threads per node	Nodes occupied by the 4 jobs	Wall-clock time (seconds)	Total CPU time (seconds per job)	Relative Billing Units
ec_smt = yes	32	1	62,350	247,100	1
ec_smt = no	16	2	44,070	309,200	1.414

Table 2 The extra cost incurred by not using SMT is illustrated in this example of two runs of an ocean model coupled with IFS; the first run uses SMT, while the second (more costly one) does not (though of course SMT was still enabled on the nodes on which it ran).

OpenMP threads = 4					OpenMP not used				
MPI tasks	SMT use	Processor binding	Memory Affinity	Wall-clock (seconds)	MPI tasks	SMT use	Processor binding	Memory Affinity	Wall-clock (seconds)
12	No	No	No	202	48	No	No	No	195
			Yes	195				Yes	190
		Yes	No	193			Yes	No	187
			Yes	182				Yes	188
24	Yes	No	No	159	96	Yes	No	No	159
			Yes	155				Yes	155
		Yes	No	155			Yes	No	160
			Yes	152				Yes	155
Test run with 1 processor per node idle (*) →					90	Yes	No	Yes	163

Table 3 The cumulative effect of SMT, processor binding and memory affinity on the job wall-clock times of two configurations of the IFS model – one using OpenMP, the other not.

(*) The run using 90 MPI tasks verified that leaving 1 processor idle on each of the 3 nodes was not beneficial for the IFS model, unlike some applications that do seem to benefit from this owing to the subsequent reduction in “operating system jitter”.

As you can see, the fastest wall-clock times were produced when using SMT with threads bound to virtual CPUs and memory affinity enabled. The best times – using all of the options (Yes) – were 26% (using OpenMP) to 33% (not using OpenMP) faster than using none of the options (No). Be aware that although the use of processor binding and memory affinity works well for the IFS model, it may not improve the performance of other codes; in fact there can be cases where it degrades performance, so this is another example of “try it and see”.

Page size considerations

On the POWER4+ clusters of ECMWF’s old Phase 3 system the operating system (AIX 5.2) supports two sizes of pages – the normal default 4 KB small page-size and the 16 MB large page-size. It has been demonstrated that large pages can increase the performance of some applications, while for others there is little or no gain. However, on the Phase 3 system it was necessary to decide upon the number of large pages prior to booting the system, and following the system boot the number was fixed. The practical and operational difficulties in setting up and using large pages meant that for all intents and purposes they were impossible to use at ECMWF, and as such large pages were never made available for use by applications.

The POWER5+ clusters run the AIX 5.3 operating system and this provides better support for large pages. A system utility is available that enables the number of large pages to be altered dynamically while the system is running, rather than having to re-boot it. Even so, although operating system support for large pages is better than in AIX 5.2, it is still very poor (e.g. in many instances this utility can take a very long time to run and often fails to execute properly due to memory fragmentation). This in effect renders large pages unusable by applications at ECMWF yet again.

All is not lost however, because AIX 5.3 (64-bit kernel only) also supports 64 KB mid-size pages. These pages are not

handled in the special way that large pages are, but are fully integrated into the operating system in the same way as small pages. Also, the gains that are seen by applications that make use of large pages are by and large seen if mid-size pages are used instead. This means that mid-size pages have all the advantages of large pages, without any of the disadvantages. There are options to the “ld” and “ldedit” commands to create “mid-size page” binaries, or you can over-ride the default mode of execution using the environment variable: “LDR_CNTRL=DATAPSIZE=64K@STACKPSIZE=64K”.

Barring unforeseen consequences it has been decided to set this by default on the Phase 4 system to ensure that all applications are able to take advantage of the possible gains in performance that mid-size pages bring. The IFS model experiences a gain in performance of about a 3% when using mid-size pages, while gains in excess of 10% have been seen using other applications.

Multi-cluster GPFS

The disk space on the old Phase 3 system was split evenly between the two POWER4+ clusters. Each cluster had its own set of disks and filesystems, connected to its own storage area network (SAN). A cluster could only access the data on the other one using techniques such as FTP, RCP, RSYNC and NFS, all of which transfer the data using IP protocols over ECMWF’s local-area network (LAN).

On the Phase 4 system this is done differently. All of the disks have been connected to the same SAN, which is also connected to both clusters. The SAN consists of two IBM “directors”, which are highly-reliable fibrechannel switches, which enable any disk to be accessed by any I/O server on either cluster. This opens up the possibility of having filesystems that are shared between the two clusters. But rather than have the data transferred through the I/O servers of one cluster over the LAN then through the I/O servers of the other cluster, it can be read directly by the I/O servers of both clusters simply via the fibrechannel SAN. The software that

enables this to be done is called the multi-cluster general parallel filesystems (MC-GPFS).

There are many advantages to using MC-GPFS, the main one being speed. Because the disks are connected via a SAN that is directly connected to both clusters, both clusters are able to read the same data at speeds comparable to those associated with disks that are directly connected to the cluster. While techniques such as FTP and such-like are limited to transferring data over the LAN at a few hundred megabytes per second maximum, data can be read via the SAN at a rate of several gigabytes per second. Other advantages include usability and less need for data replication.

With MC-GPFS no longer is it necessary for the user to use IP methods, such as those mentioned earlier, to access data on the other cluster; the data can be accessed as if it were local to the cluster on which his job is running. In fact for all intents and purposes there is no difference in this respect between data that is shared by the clusters in this way and data that is local to the cluster on which the job runs.

On the Phase 3 system certain filesystems needed to be replicated on both clusters so as to ensure that jobs could be submitted to either cluster without having to transfer files from one to the other, such as libraries, binaries and utilities. This replication of data brings its own set of problems. One has to ensure that both copies are kept in-sync and of course twice as much disk space is needed than would be the case for a single copy of the data. These problems vanish with MC-GPFS, because there is only one copy of the data in a single filesystem that is accessible by both clusters.

Having looked at the advantages, it has to be said that there are also some disadvantages. A filesystem that is shared by both systems can be seen as a single point of failure. If something happens to that filesystem, then both clusters are affected. Also, even though a filesystem can be shared by both clusters MC-GPFS has the concept of an “owning cluster”. That is one of the clusters is responsible for managing the filesystem metadata, and should that cluster crash then the filesystem would become unavailable to the other one.

To alleviate these problems the MC-GPFS setup at ECMWF is as shown in Figure 1.

A third cluster (HPCI) has been installed, which is the owning cluster for the MC-GPFS filesystems. This cluster comprises four IBM p5-575 servers, and its equipment is set up in a highly redundant fashion. There is little data traffic between HPCI and the filesystems it “owns” and no user work is allowed to run on it. In this way it is expected that this owning cluster will be highly-available and is unlikely to suffer problems that systems running a varied mixture of user work tend to see. This will mean that the filesystems should always be available on both of the production clusters (HPCE and HPCF), and if one of the production




Figure 1 The MC-GPFS configuration on the Phase 4 system, showing the computational clusters (HPCE and HPCF), the “owning” cluster (HPCI), the fibrechannel storage area network(SAN) and the 10 gigabit-ethernet local area network (LAN).

clusters should crash, or needs to be taken for a system session, this should have no effect on the other production cluster. It is expected that judicious use of MC-GPFS will increase the productivity of the users of ECMWF’s high-performance computing facility.

Getting the benefits

There are several features of the Phase 4 system that differentiate it from the Phase 3 system. In particular SMT and mid-size pages will enable the system to sustain a higher percentage of the peak performance on users’ codes. For the user there is a small price to pay – the concepts are slightly more complex to understand. However, even though ECMWF has tried to set up the system defaults in such a way as to hide some of the complication, users who take the time to understand the issues should benefit from that understanding.

Understanding the new features of Phase 4 and how to make best use of them has been a joint activity involving George Mozdzyński, Deborah Salmond and Oliver Treiber from ECMWF, John Hague from IBM, and members of the ECMWF User Support Section.

FURTHER READING

For more detailed information on the POWER5 micro- and macro-architecture see “POWER5 and Packaging”, the *IBM Journal of Research and Development* (Vol 49, number 4/5) at:

<http://www.research.ibm.com/journal/rd49-45.html>

ECMWF Calendar 2006/2007

2006		2007	
Oct 30–Nov 3	Workshop – High performance computing in meteorology (12 th Workshop)	Apr 23–24	Finance Committee (78 th Session)
Nov 13–15	Workshop – Parametrization of clouds in large-scale models	Apr (tbd)	Advisory Committee on Data Policy (8 th Session)
Dec 4–5	Workshop – HALO (3 rd Workshop)	Apr 26–27	Policy Advisory Committee (25 th Session)
Dec 7–8	Council (66 th Session)	May (tbd)	Security Representatives' Meeting
2007		May (tbd)	Computer Representatives' Meeting
Jan 29–Feb 2	Workshop – Third International Workshop on Verification Methods	Jun 4–8	Training Course – Use and interpretation of ECMWF products
Feb 8–Mar 9	Computer User Training Course	Jun (tbd)	Forecast Products – Users' Meeting
Feb 8–9	Introduction to SMS/XCDP	Jun (tbd)	Workshop – Flow-dependent aspects of data assimilation
Feb 19–23	Introduction for new users/MARS	Jun 28–29	Council (67 th Session)
Feb 26–27	MAGICS	Sep 3–7	Seminar – Recent developments in the use of satellite observations in numerical weather prediction
Feb 28–Mar 2	METVIEW	Oct 8–10	Scientific Advisory Committee (36 th Session)
Mar 5–9	Use of supercomputing resources	Oct 10–12	Technical Advisory Committee (37 th Session)
Mar 12–16	Training Course – Use and interpretation of ECMWF products	Oct 15–19	Training Course – Use and interpretation of ECMWF products for WMO Members
Mar 19–May 18	Training Course – Numerical Weather Prediction	Oct 15–16	Finance Committee (79 th Session)
Mar 19–28	Numerical methods and adiabatic formulation of models	Oct 17–18	Policy Advisory Committee (26 th Session)
Apr 16–24	Predictability, diagnostics and seasonal forecasting	Oct (tbd)	Advisory Committee of Co-operating States (13 th Session)
Apr 25–4 May	Data assimilation and use of satellite data	(tbd)	Workshop – Meteorological Operational Systems (10 th Workshop)
May 8–18	Parametrization of diabatic processes	(tbd)	Workshop – Ensemble prediction
		Dec 4–5	Council (68 th Session)

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

Technical Memoranda

- 501 **Lalaurette, F., J. Bidlot, L. Ferranti, A. Ghelli, F. Grazzini, M. Leutbecher, D. Richardson & G. van der Grijn:** Verification statistics and evaluations of ECMWF forecasts in 2004–2005. *August 2006*
- 500 **Malguzzi, P., G. Grossi, A. Buzzi, R. Ranzi & R. Buizza:** The 1966 'century' flood in Italy: a meteorological and hydrological revisitation. *August 2006*
- 496 **Bormann, N., S.B. Healy & M. Hamrud:** Assimilation of MIPAS limb radiances in the ECMWF system. Part II: Experiments with a 2-dimensional observation operator and comparison to retrieval assimilation. *August 2006*

- 495 **Bormann, N. & J.-N. Thépaut:** Assimilation of MIPAS limb radiances in the ECMWF system. Part I: Experiments with a 1-dimensional observation operator. *August 2006*

EUMETSAT/ECMWF Fellowship Programme Research Reports

- 16 **Healy, S.B., J.R. Eyre, M. Hamrud & J.-N. Thépaut:** Assimilating GPS radio occultation measurements with two-dimensional bending angle observation operators. *August 2006*

Index of past newsletter articles

This is a selection of articles published in the ECMWF Newsletter series during the last five years.

Articles are arranged in date order within each subject category. Articles can be accessed on the ECMWF public web site

www.ecmwf.int/publications/newsletter/index.html

	No.	Date	Page		No.	Date	Page
NEWS				COMPUTERS, NETWORKS, PROGRAMMING, SYSTEMS FACILITIES AND WEB			
ECMWF education and training programme for 2007	109	Autumn 2006	3	Migration of the high-performance computing service to the new IBM supercomputers	97	Spring 2003	20
Lennart Bengtsson receives the prestigious IMO Prize	109	Autumn 2006	3	ECaccess: A portal to ECMWF	96	Winter 2002/03	28
ECMWF/GEO Workshop on Atmospheric Reanalysis, ECMWF, 19–22 June 2006	109	Autumn 2006	4	ECMWF's new web site	94	Summer 2002	11
The new IBM Phase 4 HPC facility	109	Autumn 2006	5	Programming for the IBM high-performance computing facility	94	Summer 2002	9
Third International Workshop on Verification Methods	109	Autumn 2006	7				
CTBTO: "Synergies with Science"	109	Autumn 2006	7				
Summary of ECMWF Forecasts Product Users' Meeting, June 2006	109	Autumn 2006	8	METEOROLOGY			
NCEP uses ECMWF's methodology for humidity analysis and background error generation	109	Autumn 2006	9	OBSERVATIONS AND ASSIMILATION			
Annual Meetings of the European Meteorological Society	109	Autumn 2006	9	Analysis and forecast impact of humidity observations	109	Autumn 2006	11
Book about the predictability of weather and climate	108	Summer 2006	4	Surface pressure bias correction in data assimilation	108	Summer 2006	20
ECMWF's contribution to EUMETSAT's H-SAF	108	Summer 2006	2	A variational approach to satellite bias correction	107	Spring 2006	18
A celebration of David Anderson's career	108	Summer 2006	5	"Wavelet" J_b – A new way to model the statistics of background errors	106	Winter 2005/06	23
Norbert Gerber Mumm prize	108	Summer 2006	6	New observations in the ECMWF assimilation system: satellite limb measurements	105	Autumn 2005	13
Retirement of Dr Gerd Schultes	107	Spring 2006	2	CO ₂ from space: estimating atmospheric CO ₂ within the ECMWF data assimilation system	104	Summer 2005	14
A new Head of Administration for ECMWF	107	Spring 2006	3	Sea ice analyses for the Baltic Sea	103	Spring 2005	6
A real application of seasonal forecasts – Malaria early warnings	107	Spring 2006	3	The ADM-Aeolus satellite to measure wind profiles from space	103	Spring 2005	11
A kick-off workshop for THORPEX	107	Spring 2006	4	An atlas describing the ERA-40 climate during 1979–2001	103	Spring 2005	20
ECMWF's plans for 2006	106	Winter 2005/06	2	Planning of adaptive observations during the Atlantic THORPEX Regional Campaign 2003	102	Winter 2004/05	16
ECMWF/NWP-SAF workshop on bias estimation and correction in data assimilation	106	Winter 2005/06	4	ERA-40: ECMWF's 45-year reanalysis of the global atmosphere and surface conditions 1957-2002	101	Sum/Aut 2004	2
Tenth ECMWF workshop on meteorological operational systems	106	Winter 2005/06	5	Assimilation of high-resolution satellite data	97	Spring 2003	6
Co-operation Agreement with Estonia	106	Winter 2005/06	8	Assimilation of meteorological data for commercial aircraft	95	Autumn 2002	9
Workshop on the representation of subgrid processes using stochastic-dynamic models	105	Autumn 2005	2	ENSEMBLE PREDICTION			
ECMWF Forecast Products Users Meeting	105	Autumn 2005	5	The ECMWF Variable Resolution Ensemble Prediction System (VAREPS)	108	Summer 2006	14
Long-term co-operation established with ESA	104	Summer 2005	3	Limited area ensemble forecasting in Norway using targeted EPS	107	Spring 2006	23
ECMWF's highlights for 2005	103	Spring 2005	2	Ensemble prediction: A pedagogical perspective	106	Winter 2005/06	10
ECMWF and THORPEX: A natural partnership	103	Spring 2005	4	Comparing and combining deterministic and ensemble forecasts: How to predict rainfall occurrence better	106	Winter 2005/06	17
Collaboration with the Executive Body of the Convention on Long-Range Transboundary Air Pollution	103	Spring 2005	24	EPS skill improvements between 1994 and 2005	104	Summer 2005	10
Co-operation Agreement with Lithuania	103	Spring 2005	24	Ensembles-based predictions of climate change and their impacts (ENSEMBLES Project)	103	Spring 2005	16
25 years since the first operational forecast	102	Winter 2004/05	36	Operational limited-area ensemble forecasts based on 'Lokal Modell'	98	Summer 2003	2
ECMWF external policy	95	Autumn 2002	14	Ensemble forecasts: can they provide useful early warnings?	96	Winter 2002/03	10
				Trends in ensemble performance	94	Summer 2002	2
COMPUTING				ENVIRONMENTAL MONITORING			
ARCHIVING, DATA PROVISION AND VISUALISATION				Progress with the GEMS project	107	Spring 2006	5
A simple false-colour scheme for the representation of multi-layer clouds	101	Sum/Aut 2004	30	A preliminary survey of ERA-40 users developing applications of relevance to GEO (Group on Earth Observations)	104	Summer 2005	5
The ECMWF public data server	99	Aut/Win 2003	19	The GEMS project – making a contribution to the environmental monitoring mission of ECMWF	103	Spring 2005	17
COMPUTERS, NETWORKS, PROGRAMMING, SYSTEMS FACILITIES AND WEB							
New features of the Phase 4 HPC facility	109	Autumn 2006	32				
Developing and validating Grid Technology for the solution of complex meteorological problems	104	Summer 2005	22				
Migration of ECFS data from TSM to HPSS ("Back-archive")	103	Spring 2005	22				
New ECaccess features	98	Summer 2003	31				

	No.	Date	Page		No.	Date	Page
ENVIRONMENTAL MONITORING				METEOROLOGICAL STUDIES			
Environmental activities at ECMWF	99	Aut/Win 2003	18	A snowstorm in North-Western Turkey 12–13 February 2004 – Forecasts, public warnings and lessons learned	102	Winter 2004/05	7
FORECAST MODEL				Exceptional warm anomalies of summer 2003	99	Aut/Win 2003	2
Ice supersaturation in ECMWF's Integrated Forecast System	109	Autumn 2006	26	Record-breaking warm sea surface temperatures of the Mediterranean Sea	98	Summer 2003	30
Towards a global meso-scale model: The high-resolution system T799L91 and T399L62 EPS	108	Summer 2006	6	Breakdown of the stratospheric winter polar vortex	96	Winter 2002/03	2
The local and global impact of the recent change in model aerosol climatology	105	Autumn 2005	17	Central European floods during summer 2002	96	Winter 2002/03	18
Improved prediction of boundary layer clouds	104	Summer 2005	18	OCEAN AND WAVE MODELLING			
Two new cycles of the IFS: 26r3 and 28r1	102	Winter 2004/05	15	Progress in wave forecasts at ECMWF	106	Winter 2005/06	28
Early delivery suite	101	Sum/Aut 2004	21	Ocean analysis at ECMWF: From real-time ocean initial conditions to historical ocean analysis	105	Autumn 2005	24
Systematic errors in the ECMWF forecasting system	100	Spring 2004	14	High-precision gravimetry and ECMWF forcing for ocean tide models	105	Autumn 2005	6
A major new cycle of the IFS: Cycle 25r4	97	Spring 2003	12	MERSEA – a project to develop ocean and marine applications	103	Spring 2005	21
METEOROLOGICAL APPLICATIONS				Towards freak-wave prediction over the global oceans	100	Spring 2004	24
Recent developments in extreme weather forecasting	107	Spring 2006	8	Probabilistic forecasts for ocean waves	95	Autumn 2002	2
Early medium-range forecasts of tropical cyclones	102	Winter 2004/05	7	MONTHLY AND SEASONAL FORECASTING			
European Flood Alert System	101	Sum/Aut 2004	30	Monthly forecasting	100	Spring 2004	3
Model predictions of the floods in the Czech Republic during August 2002: The forecaster's perspective	97	Spring 2003	2	DEMETER: Development of a European multi-model ensemble system for seasonal to interannual prediction	99	Aut/Win 2003	8
Joining the ECMWF improves the quality of forecasts	94	Summer 2002	6	The ECMWF seasonal forecasting system	98	Summer 2003	17
METEOROLOGICAL STUDIES				Did the ECMWF seasonal forecasting model outperform a statistical model over the last 15 years?	98	Summer 2003	26
Hindcasts of historic storms with the DWD models GME, LMQ and LMK using ERA-40 reanalyses	109	Autumn 2006	16				
Hurricane Jim over New Caledonia: a remarkable numerical prediction of its genesis and track	109	Autumn 2006	21				
Starting-up medium-range forecasting for New Caledonia in the South-West Pacific Ocean – a not so boring tropical climate	102	Winter 2004/05	2				

Useful names and telephone numbers within ECMWF

Telephone

Telephone number of an individual at the Centre is:

International: +44 118 949 9 + three digit extension

UK: (0118) 949 9 + three digit extension

Internal: 2 + three digit extension

e.g. the Director's number is:

+44 118 949 9001 (international),

(0118) 949 9001 (UK) and 2001 (internal).

E-mail

The e-mail address of an individual at the Centre is:

firstinitial.lastname@ecmwf.int

e.g. the Director's address is: D.Marbouty@ecmwf.int

For double-barrelled names use a hyphen

e.g. J-N.Name-Name@ecmwf.int

Internet web site

ECMWF's public web site is: <http://www.ecmwf.int>

	Ext		Ext
Director		Meteorological Division	
Dominique Marbouty	001	<i>Division Head</i>	
Deputy Director & Head of Research Department		Horst Böttger	060
Philippe Bougeault	005	<i>Meteorological Applications Section Head</i>	
Head of Operations Department		Alfred Hofstadler	400
Walter Zwiefelhofer	003	<i>Data and Services Section Head</i>	
Head of Administration Department		Baudouin Raoult	404
Ute Dahremöller	007	<i>Graphics Section Head</i>	
		Jens Daabeck	375
Switchboard		<i>Meteorological Operations Section Head</i>	
ECMWF switchboard	000	David Richardson	420
Advisory		<i>Meteorological Analysts</i>	
Internet mail addressed to Advisory@ecmwf.int		Antonio Garcia Mendez	424
Telefax (+44 118 986 9450, marked User Support)		Anna Ghelli	425
Computer Division		Claude Gibert (web products)	111
<i>Division Head</i>		Fernando Prates	421
Isabella Weger	050	Meteorological Operations Room	426
<i>Computer Operations Section Head</i>		Data Division	
Sylvia Baylis	301	<i>Division Head</i>	
<i>Networking and Computer Security Section Head</i>		Adrian Simmons	700
Rémy Giraud	356	<i>Data Assimilation Section Head</i>	
<i>Servers and Desktops Section Head</i>		Erik Andersson	627
Richard Fisker	355	<i>Satellite Data Section Head</i>	
<i>Systems Software Section Head</i>		Jean-Nöel Thépaut	621
Neil Storer	353	<i>Re-Analysis Project (ERA) Head</i>	
<i>User Support Section Head</i>		Saki Uppala	366
Umberto Modigliani	382	Probabilistic Forecasting & Diagnostics Division	
<i>User Support Staff</i>		<i>Division Head</i>	
Paul Dando	381	Tim Palmer	600
Anne Fouilloux	380	<i>Seasonal Forecasting Section Head</i>	
Dominique Lucas	386	Franco Molteni	108
Carsten Maaß	389	Model Division	
Pam Prior	384	<i>Division Head</i>	
Computer Operations		Martin Miller	070
<i>Call Desk</i>		<i>Numerical Aspects Section Head</i>	
<i>Call Desk email: calldesk@ecmwf.int</i>	303	Mariano Hortal	147
<i>Console - Shift Leaders</i>		<i>Physical Aspects Section Head</i>	
<i>Console fax number +44 118 949 9840</i>	803	Anton Beljaars	035
<i>Console email: newops@ecmwf.int</i>		<i>Ocean Waves Section Head</i>	
<i>Fault reporting - Call Desk</i>	303	Peter Janssen	116
<i>Registration - Call Desk</i>	303	GMES Coordinator	
<i>Service queries - Call Desk</i>	303	Anthony Hollingsworth	824
<i>Tape Requests - Tape Librarian</i>	315	Education & Training	
		Renate Hagedorn	257
		ECMWF library & documentation distribution	
		Els Kooij-Connally	751

© Copyright 2006

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. 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, omission and for loss or damage arising from its use.