

Ice supersaturation in ECMWF's Integrated Forecast System

Adrian Tompkins, Klaus Gierens, Gaby Rädcl

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 Rädcl: 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:

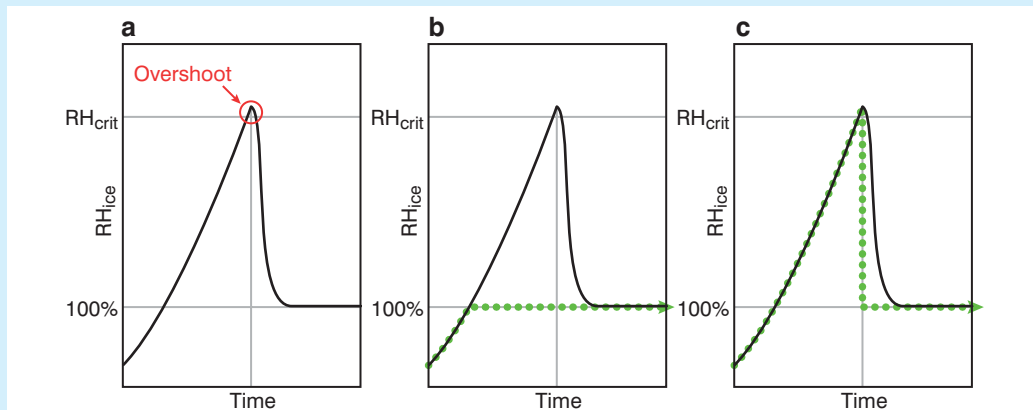
- ◆ Ice nucleation initiates when the supersaturation locally reaches the homogeneous nucleation threshold which is specified as a function of temperature and can reach values exceeding 170% (See box A for details).
- ◆ The clear-sky humidity fluctuations are uniformly distributed with a fixed constant 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.

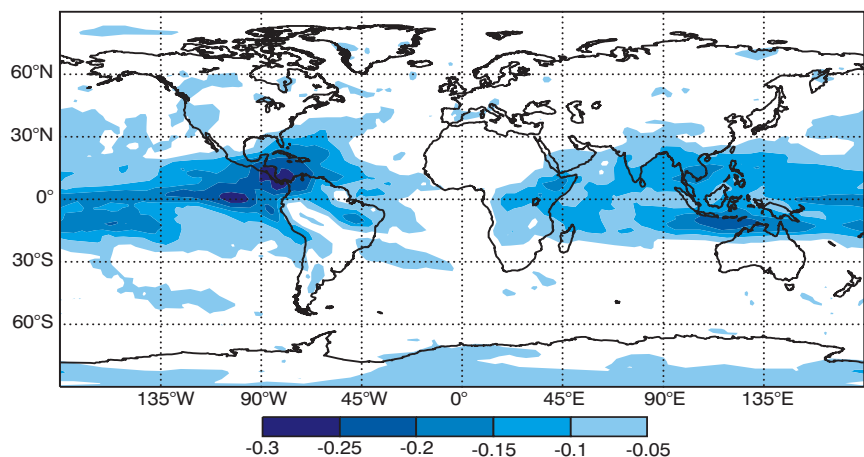
In terms of the model cloud climatology, the largest sensitivity was not in the total column ice water (TCIW, sometimes referred to as the ice water path) but rather the total cloud cover. High cloud cover is reduced by the new scheme, as expected, since humidity must now attain higher values to initiate 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

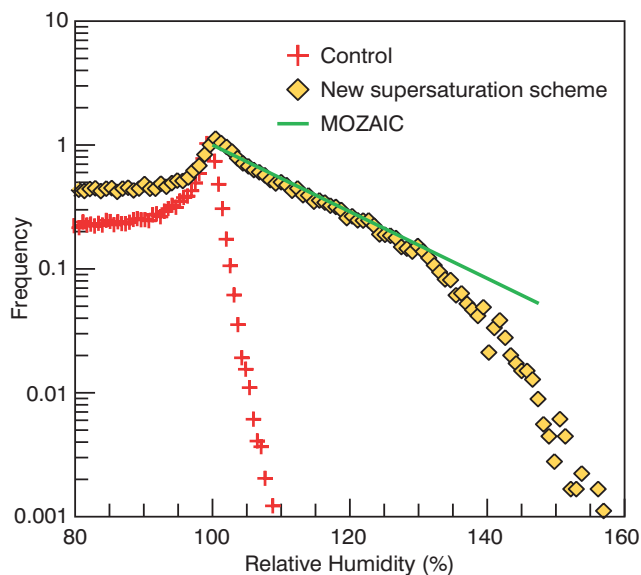


Figure 3 Normalised probability distribution function of RH at the 250 hPa level for the northern hemisphere for the control and new supersaturation scheme using the first member from each set of climate integrations compared to aircraft measurements from the MOZAIC campaign.

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.

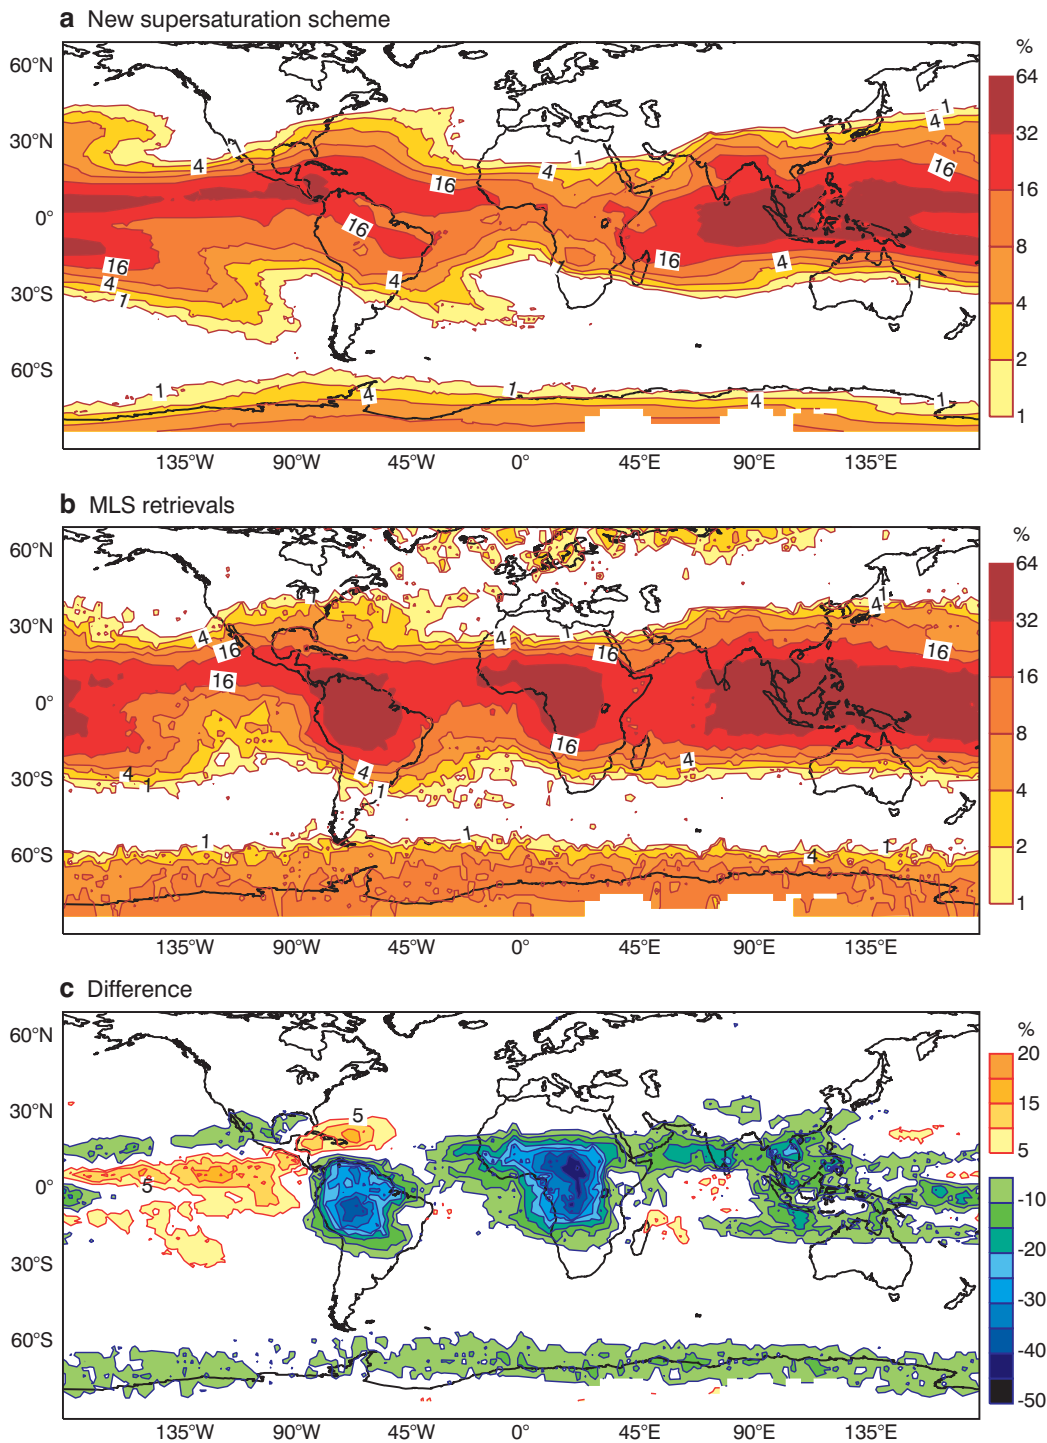


Figure 4 Annual mean frequency (%) of occurrence of ice supersaturation for (a) the model level at 144 hPa for the seven-member climate ensemble using the new supersaturation scheme and (b) the MLS retrieval climatology at 147 hPa. (c) Difference between (a) and (b).

	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		
Predicted	Observed		Predicted	Observed	
	Yes	No		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