TECHNICAL MEMORANDUM

741

All-sky assimilation of microwave humidity sounders

Alan J. Geer, Fabrizio Baordo, Niels Bormann and Stephen English

Research Department

November 2014

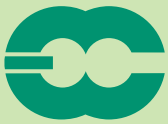

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

Series: ECMWF Technical Memoranda

A full list of ECMWF Publications can be found on our web site under:
http://www.ecmwf.int/en/research/publications

Contact: library@ecmwf.int

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

Abstract

A main aim of humidity, cloud and precipitation assimilation in global weather forecasting is to infer winds and other dynamical variables directly in the data assimilation system. This 'model-tracing' effect helps improve dynamical initial conditions and leads to improved forecasts. Recent improvements in radiative transfer modelling (particularly in the scattering from frozen particles) and in surface emissivity over land and sea-ice mean it is now feasible to assimilate microwave humidity sounding observations in all sky conditions over a large part of the globe, adding observations in many of the most interesting meteorological areas. The next operational version of the ECMWF system will assimilate all-sky humidity sounding observations from four Microwave Humidity Sounders (MHS) and one Special Sensor Microwave Imager / Sounder (SSMIS). These observations give significant improvements to dynamical forecast scores in the medium range in the midlatitudes, along with benefits to wind and temperature fields around the tropical tropopause. Single observation test cases and single-observing system experiments help isolate the mechanism that improves forecasts in the midlatitudes: 4D variational assimilation can indeed infer dynamical initial conditions from the humidity, cloud and precipitation features in the radiance observations. The benefit is greatest in the southern midlatitudes, where the storm-tracks provide ideal conditions for model-tracing. Here, the impact of all-sky humidity channels on upper-tropospheric winds approaches that of the microwave temperature sounding instruments, which can use geostrophic balance to infer winds in the assimilation system. All-sky assimilation has roughly doubled the impact on forecast scores and observation fits compared to clear-sky assimilation. For example, clear-sky assimilation of microwave humidity sounding observations improves day 5 forecasts in the southern hemisphere by about ¾ hour; all-sky assimilation by about 1½ hours. Cloud- and precipitation-affected radiances are now assimilated from both microwave imagers and microwave humidity sounders; it is hoped to add infrared humidity sounding observations in future.

1 Introduction

Satellite humidity observations are sometimes thought to have limited impact on weather forecasts. In the ERA-40 reanalysis, humidity observations gave apparently no benefit (Bengtsson and Hodges, 2005) and worse, biases between model and observations led to excessive tropical precipitation (Uppala et al., 2005). In these pioneering efforts to assimilate satellite humidity data, a large part of the impact (for good or ill) came through systematic changes in the humidity analysis (e.g. Gérard and Saunders, 1999). Perhaps these early difficulties encouraged a pessimistic view of humidity observations that persists to this day. Numerical weather prediction (NWP) is often thought to benefit only from wind, temperature and pressure information (following e.g. Smagorinsky et al., 1970) from in-situ and satellite measurements (e.g. English et al., 2000); in this view wind is mainly inferred from the temperature field by including geostrophic balance in the background error covariances, given a lack of direct observations. However, this significantly underestimates the potential of humidity observations, because free-troposphere humidity features are principally driven by winds. The great hope of humidity assimilation is to infer winds directly in the data assimilation system, using the adjoint of the forecast model (Andersson et al., 1994). For many years, sequences of water vapour images have been used to infer upper tropospheric wind fields (e.g. Velden et al., 1997) and the possibility of inferring winds from tracer fields using data assimilation has been studied with both Kalman-filter and variational approaches (Daley, 1995; Riishøjgaard, 1996; Peuch et al., 2000; Allen et al., 2013, 2014). In the absence of cloud and precipitation, and away from the surface, humidity is a tracer and governed by a continuity equation. It is obvious that wind can be inferred if it is aligned with a strong gradient in the tracer field. But even in a uniform tracer field, the tracer concentration can be increased or decreased through convergence or divergence. In a theoretical setup representative of the stratosphere, it is possible to almost completely constrain the dynamics with error-free tracer observations, even without using dynamical observations (Allen et al., 2014).

After many developments in data assimilation, models and observations, including the move from three to four dimensional variational assimilation (3D-Var to 4D-Var), Andersson et al. (2007) demonstrated a combined impact from all humidity observing systems of 2%–4% on temperature and wind forecasts at day 4 in the extra-tropics. In the context of a state of the art NWP system, that is a substantial beneficial impact. However, that study did not examine the mechanism behind the dynamical impact of humidity observations and discussed mainly: (a) direct improvements to the humidity field that could lead to improvements in the dynamical forecast, for example, through modifying precipitation and changing latent heat release and (b) deriving temperature information from background error correlations between temperature and humidity (e.g. Hólm et al., 2002). However, Peubey and McNally (2009) were able to separate the effects of 4D-Var humidity tracing, background error correlations and the impact of improved humidity fields on subsequent forecasts. They identified humidity tracing as the main mechanism by which geostationary water vapour observations improve the wind fields in the ECMWF system.

What applies to water vapour can also apply to clouds and precipitation. Inferring winds from the motion of clouds has a longer history than humidity tracing though it has always been recognised that clouds are far from pure tracers, i.e. they are driven by more than a continuity equation (e.g Fujita et al., 1975). However, modern data assimilation systems represent many of the mechanisms of cloud and precipitation in the forecast model. In the ECMWF system, convective and large-scale cloud and precipitation is represented in the tangent linear and adjoint models used by 4D-Var (Tompkins and Janisková, 2004; Lopez and Moreau, 2005; Janisková and Lopez, 2013). So just as the 4D-Var tracer effect follows from the adjoint form of the continuity equation, a much broader 'model tracing' effect comes from all the physical processes modelled inside 4D-Var. This should be able to infer many aspects of the dynamical description of the atmosphere. For example, an observation of convection should be able to generate convection in a model where it is lacking, perhaps through changes in humidity, temperature profile (i.e decreased stability) or low-level convergence. ECMWF has been using this broader 'model tracing' operationally for over five years in the assimilation of all-sky microwave imager radiances (Bauer et al., 2010; Geer and Bauer, 2010). In testing that work, a single cloud-affected observation could be used to shift the position of a front, generating increments to surface pressure and winds. One of the biggest remaining issues with assimilating cloud and precipitation-affected observations has parallels with the early days of humidity assimilation: large biases between model and observations mean that data usage has to be quite cautious in certain areas, e.g. maritime low cloud situations (Kazumori et al., 2014).

Microwave water vapour sounding observations, like microwave imaging observations, are ideal for all-sky assimilation because of their relatively smooth and linear response to water vapour, cloud and precipitation (e.g. Bauer et al., 2010; English et al., 2013). However, operational all-sky assimilation has had to wait for improved fast radiative transfer models that can more accurately simulate scattering effects from the frozen particles that dominate cloudy radiative transfer at the high microwave frequencies used for humidity sounding (Liu, 2008; Geer and Baordo, 2014). With this in place, all-sky assimilation of Special Sensor Microwave Imager/Sounder (SSMIS) humidity sounding channels over ocean produced beneficial impacts on wind forecasts and is now part of the operational system at ECMWF (Geer, 2013). The impact appeared to come from the broader 'model tracing' effect of 4D-Var, from humidity, cloud and precipitation, with the all-sky giving roughly twice the impact of a clear-sky approach. This motivated work to transfer all microwave humidity sounding channels to the all-sky framework at ECMWF. The bulk of that information comes from the Microwave Humidity Sounders (MHS) on four operational polar meteorological satellites. These are already assimilated over land and sea-ice and this capability was missing from the all-sky framework at ECMWF and needed to be added (Baordo and Geer, 2014). This report describes the all-sky assimilation of one SSMIS and four MHS sensors over ocean, land and sea-ice. With such an amount of data it is easier to see its impact, to explore the importance of 'model tracing' of humidity, cloud and precipitation, and to critically examine benefits the all-sky approach.

The successful assimilation of cross-track and conical scanners over ocean, land and sea-ice and in clear-skies, cloud and precipitation relies on a large amount of scientific infrastructure. In this technical report we examine and briefly justify much of this, providing detail that should be useful to ourselves and colleagues working on similar activities at other NWP centres. We hope to provide a briefer summary of the work in a journal publication.

2 Method

2.1 Overview

All-sky microwave radiances are assimilated directly alongside many other conventional and satellite observation types in the ECMWF 4D-Var data assimilation system (Rabier et al., 2000) with background error covariances coming from an ensemble of data assimilations (EDA, Bonavita et al., 2012). The atmospheric control variables are transforms of surface pressure (a 2D field), and humidity, the two horizontal wind components and the temperature (3D fields). There is no cloud or precipitation control variable, but in the minimisation, cloud and precipitation are diagnosed from the dynamic and humidity fields every timestep including the first (see Sec. 2.8). The current approach allows the analysis to fit all-sky observations by modifying cloud and precipitation in the model but ultimately by adjusting the dynamical and humidity control variables at the beginning of the window.

This study examines the combined results from summer and winter experiments covering six months: August–October 2013 and January–March 2014. A month's spinup has already been excluded from the beginning of each experiment. The resolution is T511 (around 40 km) and though this is lower than the operational T1279, it is sufficient to generate realistic cloud and precipitation fields and is the normal resolution for testing at ECMWF. The assimilation window is 12 h long with two non-overlapping windows per day: 21 UTC to 09 UTC and 09 UTC to 21 UTC. The experiments use the cycle 40r2 version of the ECMWF system but with the satellite assimilation configuration of cycle 40r3, including the all-sky assimilation of SSMIS and MHS. From this configuration, some other experiments have been derived (for example by denying the microwave humidity observations, or reverting to the old clear-sky assimilation) but these will be introduced as we go along.

2.2 Observations

In this work we are examining the assimilation of 183 GHz humidity sounding channels from two instruments, MHS (Robel, 2009) on two American and two European satellites and SSMIS (Kunkee et al., 2008) on Defense Meteorological Satellite Programme satellite F17 (DMSP-F17). The former are cross-track scanners with a variable size field of view and zenith angle, whereas the latter is a conical-scanning imager, with fixed field of view and zenith angle. Table 1 gives more details of data usage, which is affected by instrument differences in some areas (such as thinning, superobbing and cloud identification, mentioned later) but at a broader level the processing of the two instruments is unified. Both instruments have similar channels, summarised in the table, though the exact details of frequency and polarisation differ. The 90 GHz and 150 GHz window channels are used for emissivity retrievals over land and sea-ice but are not actively assimilated. Note that channel 17 on SSMIS (91 GHz, V) is assimilated over oceans as part of the all-sky microwave imager assimilation (e.g. Bauer et al., 2010; Geer and Bauer, 2011; Kazumori et al., 2014), but that is not affected or tested in this work on the humidity sounding channels, where channel 18 (91 GHz, H) is used for emissivity retrievals. The three 183 GHz channels,

Table 1: Satellites, instruments and channels used in this work. Polarisation: V = vertical; H = horizontal. NOAA (National Oceanic and Atmospheric Administration) satellites are US operational polar orbiters; Metop satellites are the operational European polar orbiters.

Instrument	**SSMIS**		**MHS**	
Satellites	DMSP-F17		NOAA-18 NOAA-19 Metop-A Metop-B	
Channel usage	Channel number	Frequency [GHz] and polarisation	Channel number	Frequency [GHz] and polarisation
Emissivity retrievals over land	18	91.655 H	1	89.0 V
Emissivity retrievals over sea-ice	8	150.0 H	2	157.0 V
Upper tropospheric humidity	11	183.31±1 H	3	183.311±1 H
Mid-upper tropospheric humidity	10	183.31±3 H	4	183.311±3 H
Mid-lower tropospheric humidity	9	183.31±6.6 H	5	190.311 V

with weighting functions spanning the troposphere, are what we want to assimilate actively and will often be referred to as ‘upper’, ‘mid’ and ‘lower’ tropospheric humidity channels, unless we are referring to a specific instrument, in which case the channel number or frequency may be used.

Note that there are many other sources of free-troposphere humidity data in the global observing system, including geostationary and polar orbiting infrared sensors, aircraft and radiosonde data. There is also one remaining microwave sensor, Advanced Technology Microwave Sounder (ATMS) with actively assimilated 183 GHz channels that have not yet been transferred to the all-sky approach. All of these sensors, including ATMS, remain actively assimilated in all of our experiments based on the full observing system. The only sensors for which data usage is altered in our full-observing system experiments are SSMIS and MHS, and only the 183 GHz channels on those sensors.

2.3 Observation operator

Radiative transfer simulations are provided by RTTOV-SCATT, which is a fast model designed for assimilating microwave radiances in all-sky conditions (Bauer et al., 2006). It is a component of the wider RTTOV package (Radiative Transfer model for Television Infrared Observation Satellite Operational Vertical sounder; Eyre, 1991; Saunders et al., 2012). The radiative transfer equation, including scattering from cloud and precipitation, is solved using the delta-Eddington approximation (Joseph et al., 1976). Transmittances for oxygen and water vapour are computed from regression tables driven by atmospheric predictors. Bulk optical properties of hydrometeors are taken from lookup tables, with the optical properties of cloud water, cloud ice and rain hydrometeors generated using Mie theory (Bauer, 2001) and the optical properties of snow hydrometeors from discrete dipole calculations (Liu, 2008; Geer and Baordo, 2014). Ocean surface emissivity is computed by version 6 of FASTEM (English and Hewison, 1998; Liu et al., 2011; Kazumori and English, 2014). Land-surface and sea-ice emissivity comes from a mixture of retrievals and atlas that is described in Sec. 2.7. The all-sky brightness temperature is computed as the weighted average of the brightness temperature from two independent sub-columns, one clear and one cloudy. The weighting is done according to the effective cloud fraction of Geer et al. (2009) which accounts for the effects of sub-grid variability in cloud and precipitation

2.4 Observation errors

In cloudy and precipitating situations, the dominant source of random error in the first guess (FG) departures is (broadly speaking) representivity, not instrument noise or radiative transfer inaccuracies (Geer and Bauer, 2011). Practically, the model cannot simulate cloud and precipitation with a precise enough intensity or location to give a good fit to observations. There is some debate as to whether this could be treated as background or model error, but in practice operational all-sky assimilation works when these 'representivity' errors are treated as observation error. For microwave imagers, the problem of spatial representivity is partly addressed through superobbing, i.e. spatially averaging the observations to scales more representative of the model's cloud and precipitation (see e.g. Roberts and Lean, 2008). The SSMIS observations used here are superobbed in roughly 80 km by 80 km boxes (Geer and Bauer, 2010). MHS has a raw field of view running from about 17 km diameter at nadir (Robel, 2009) to around 30 km at the edge of swath (see Fig. 3 of Bennartz, 2000). A decision was made not to superob MHS, mainly due to the technical difficulty of combining observations with different zenith angles. Following the logic of the microwave imagers, the MHS observations would ideally have been superobbed, but this has been left for future work.

The second stage in treating representivity error, as well as the other smaller sources of error, is to apply an observation error model. Here we use the error model of Geer and Bauer (2010, 2011) which inflates observation errors as a function of symmetric 'cloud' amount, i.e. the average amount of cloud or precipitation in the model and observations. As a measure of cloud amount over land, we use a scattering index (SI, e.g. Baordo et al., 2012):

$$\mathrm{SI} = \mathrm{TB}_{90} - \mathrm{TB}_{150}, \tag{1}$$

where TB_{90} or TB_{150} is the brightness temperature of the nearest channel to 90 GHz or 150 GHz listed in Tab. 1. Brightness temperatures at 150 GHz are strongly depressed by scattering from precipitation-sized ice particles but the effect of scattering is less at 90 GHz. Hence, SI greater than a few Kelvin tends to indicate the presence of scattering. An SI of 50 K indicates very strong scattering, usually associated with deep convection. SI is computed separately from observations ($\mathrm{SI}_{\mathrm{obs}}$) or from model first guess brightness temperatures that have been bias-corrected towards the observations ($\mathrm{SI}_{\mathrm{FG}}$) so that the symmetric cloud predictor C_{SYM} can be computed as:

$$C_{\mathrm{SYM}} = (\mathrm{SI}_{\mathrm{obs}} + \mathrm{SI}_{\mathrm{FG}})/2 \tag{2}$$

Over ocean, Eq. 1 is sensitive not just to scattering, but also to water vapour absorption, which causes up to 50 K variation between 90 GHz and 150 GHz brightness temperatures. For SSMIS it is most straightforward to continue computing the symmetric cloud amount from the 37 GHz polarisation difference, as used for assimilation of the imager channels (Geer, 2013). This is not possible for MHS, so we use the 'ocean SI':

$$\mathrm{SI}_{\mathrm{ocean}} = (\mathrm{TB}_{90} - \mathrm{TB}_{150}) - (\mathrm{TB}_{90}^{\mathrm{clr}} - \mathrm{TB}_{150}^{\mathrm{clr}}). \tag{3}$$

Here the first term is just the SI from Eq. 1 and it can be computed from modelled or observed brightness temperatures. The second term is the clear-sky SI, which removes the effect of water vapour absorption. It is always simulated from the model and we assume that lower tropospheric moisture is reasonably well forecast, at least compared to the cloud and precipitation. With most of the water vapour signal removed, $\mathrm{SI}_{\mathrm{ocean}}$ has similar properties to SI computed over land and is mostly a measure of scattering from frozen particles. A symmetric cloud predictor is constructed from observed and simulated $\mathrm{SI}_{\mathrm{ocean}}$ in the same way as over land.

Over sea-ice, the frequency dependence of surface emissivity depends strongly on the depth and properties of the snow cover, so it is impossible to use a scattering index to uniquely identify frozen particles

Table 2: Parameters of the observation error model for MHS. All units are Kelvin.

	Channel	$g_{\rm clr}$	$g_{\rm cld}$	$C_{\rm clr}$	$C_{\rm cld}$
Ocean	3	2.0	10.0	0.0	37.0
	4	2.0	20.0	0.0	34.0
	5	2.2	32.0	0.0	30.0
Land	3	3.0	18.0	0.0	24.5
	4	3.0	40.0	0.0	24.5
	5	3.0	60.0	0.0	22.4

in the atmosphere. Instead, for both SSMIS and MHS, we abandon the adaptive observation error model and just used a fixed observation error across clear or cloudy skies. This is justified by the relative lack of strong cloud and precipitation effects on the brightness temperatures over sea ice. The lowest-peaking channel we assimilate over sea ice is 183±3 GHz. In this channel, over sea-ice, the effect of cloud and precipitation on model simulated brightness temperatures is greater than 1 K in only 6% of cases; it is greater than 5 K in only 0.2% of cases. This can be contrasted with the signals of 150 K in convective precipitation that are the main driver for the variable observation error model at lower latitudes. A fixed observation error is a reasonable starting point over sea-ice.

The land and ocean observation error models for SSMIS 183 GHz channels have been examined in Geer and Bauer (2011); Geer (2013); Baordo et al. (2012, 2013); Baordo and Geer (2014) so here we concentrate on the specification of MHS observation errors. Figure 1a and b show the number of observations over ocean and land binned by the appropriate symmetric cloud predictor: as usual the vast majority of microwave observations are associated with little scattering, i.e. they are mostly unaffected by cloud and precipitation. It is also clear that the ocean SI and the land SI are not identical. By removing the majority of clear-sky (i.e. water vapour) and surface influences from the difference between 90 GHz and 150 GHz brightness temperatures, the ocean SI gives a much tighter distribution around zero.

The other panels on Fig. 1 show the standard deviation of FG departures as a function of symmetric SI with small values for clear-sky scenes (small SI) and larger values for cloudy scenes (large positive SI, see the dot-dash lines). The adaptive error model makes a fit to this distribution and uses it to predict the total error of FG departures as a function of cloud amount (the dashed lines). Until now the increase in error has been modelled as a linear function of the symmetric cloud predictor. That would still be possible over land (panels d,f, and h) but over ocean the increase as a function of ocean SI is better modelled by a quadratic fit. For simplicity a quadratic fit has been used for MHS observations over both land and sea. For land observations this is possible because observations with $C_{\rm SYM} > 20$K, which would have benefitted from a linear error model, are removed by quality control (see later). The quadratic formulation is:

$$g(C_{\rm SYM}) = \begin{cases} g_{\rm clr} & \in C_{\rm SYM} \leq C_{\rm clr} \quad (4) \\ g_{\rm clr} + (g_{\rm cld} - g_{\rm clr}) \left(\frac{C_{\rm SYM} - C_{\rm clr}}{C_{\rm cld} - C_{\rm clr}} \right)^2 & \in C_{\rm clr} < C_{\rm SYM} < C_{\rm cld} \quad (5) \\ g_{\rm cld} & \in C_{\rm SYM} \geq C_{\rm cld}. \quad (6) \end{cases}$$

Table 2 gives the fixed parameters of the MHS error model. The Geer and Bauer (2011) model acknowledges that the standard deviation of FG departures, referred to as the 'total error', is explained by both background errors and observation errors (representivity errors are included within the observation error). The model gives a possibility to remove a background error contribution from the total error so that the assigned observation errors are smaller than the standard deviation of FG departures. However,

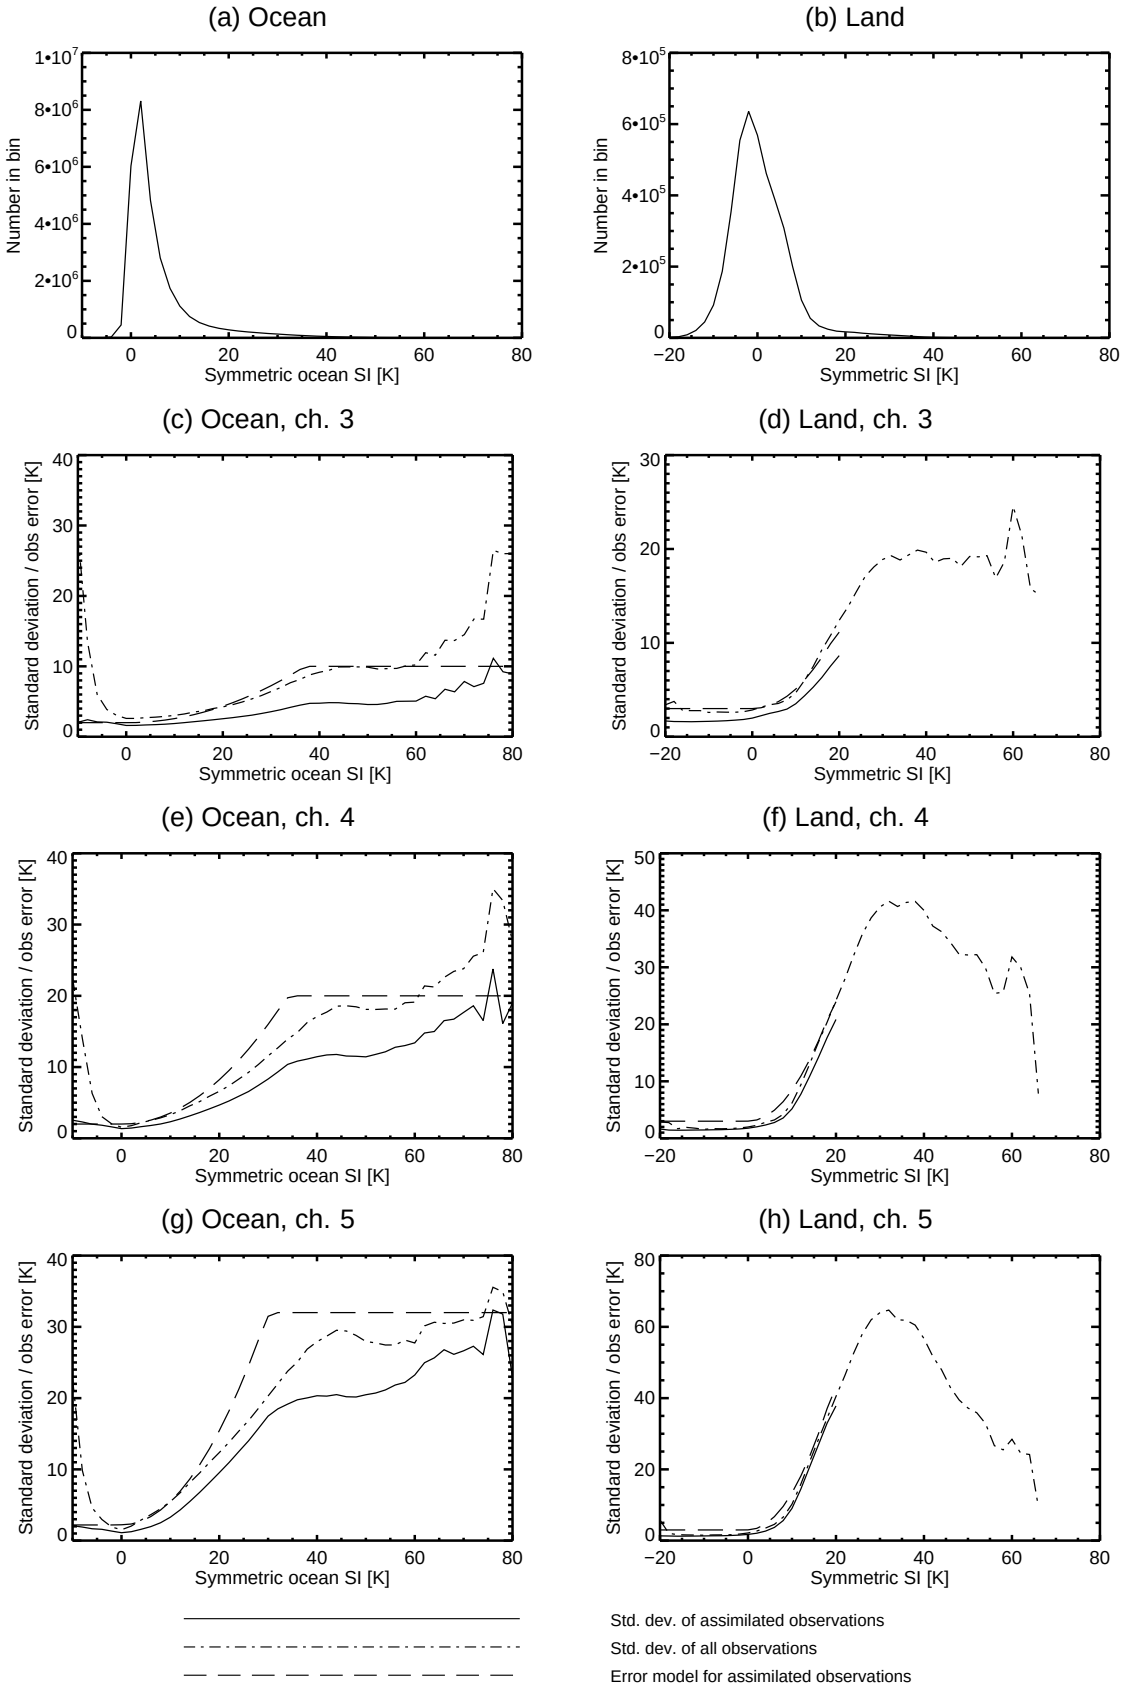

Figure 1: (a-b) Number of observations binned by symmetric 'cloud' amount; (c-j) Standard deviation of MHS FG departures, along with the applied observation errors. Based on a sample of 6 months of observations.

given that representivity error dominates the error budget, this aspect has been ignored for the MHS and SSMIS error models and the assigned observation error is equal to $g(C_{\mathrm{SYM}})$, the total error. As will be seen later, this is fine for cloudy-skies but less good in clear-sky situations, where large background errors can still occur (Bormann and Bonavita, 2013) and assigned observation errors could potentially be smaller. Fig. 1 shows the observation errors resulting from the quadratic fit (the dashed line) are not perfect at higher SI, for example in panels e and g. However, the number of observations with high SI is small, the observation errors are large, so we need not worry too much about that. It is more important to ensure the error model works well for the bulk of observations, which are found at lower values of SI, and for this the quadratic fit works well. The fit has been made by hand, erring on the side of caution, so for the lower channels over land the modelled observation error is 3 K in clear skies, larger than the standard deviation of FG departures (panels f and h).

Figure 1 also distinguishes the standard deviation of FG departures for all observations after basic quality control (for e.g. snow surfaces, high altitude orography etc., dot-dash line) and for the sample actually assimilated, after the first guess check and Variational QC (VarQC, solid line - see next section). The application of the first guess check and VarQC create a sample that has lower standard deviation of FG departures than the sample used to fit the error model. The calculation of standard deviation used for the error model is certainly influenced by outliers that are removed by QC, so an alternative fit based on the filtered observations might produce smaller standard deviations. Nevertheless, the differences that would result in observation error would not be large. The results of the error model are examined further in the next section.

2.5 Quality control

Before any QC, the observations are thinned to a spacing of around 100 km, to avoid the possibility of horizontally correlated errors but more pragmatically to reduce the computational cost of assimilation. This is achieved by keeping only the nearest observation to the grid-point of an 80 km reduced Gaussian grid (T255-N128) and then eliminating every second remaining observation. Quality control is then applied to remove situations where (a) it is impossible to accurately simulate observations, due to problems in the observing system, the forecast model or in the radiative transfer that can be predicted in advance; (b) where a first guess departure is larger than expected according to the predicted observation and background error; (c) where an observation does not agree with other observations. These steps can be referred to respectively as 'Basic' QC, first guess check (Järvinen and Unden, 1997), and VarQC (Andersson and Järvinen, 1998). Table 3 summarises the screening of 6 months of MHS channel 5 observations. This is the lowest peaking humidity sounding channel that we assimilate and it is partially sensitive to the surface. Quality control of MHS channel 5 is dominated by the need to avoid snow-covered surfaces and sea-ice, where surface emissivities are prone to error. Higher peaking channels are less affected by the surface and some QC checks are relaxed, but channel 5 gives an example of the full range of checks.

Snow-covered surfaces and sea-ice are removed through checks on latitude, model sea-ice and surface temperature (rejecting temperatures lower than 274 K over ocean and 278 K over land). Being sensitive to boundary layer cloud, channel 5 is also affected by a longstanding model deficiency over high latitude oceans: a lack of supercooled liquid water cloud known as the 'cold sector' problem (e.g. Geer et al., 2009; Kazumori et al., 2014). None of these checks are applied to channels 3 and 4, which peak high enough to be mostly unaffected by snow, sea-ice or boundary layer cloud problems. Over land, there is also an orography check intended to avoid situations where the surface emission becomes too important (rejection for altitudes higher than 1,500 m, 1,000 m and 800 m in channels 3, 4, and 5). The final

Table 3: Quality control of MHS channel 5 for six months of observations. Basic QC checks are not exclusive: many can be failed simultaneously.

	Remaining		Removed	
	Total	Percentage	Total	Percentage
After thinning	82,149,578	100		
Surface temperature too low			28,538,104	34.7
Latitude greater than 60°			27,493,436	33.5
Cold sectors 2 (TCWV less than 8 kg m^{-2})			13,455,416	16.4
Inhomogeneous surface			11,376,632	13.8
High orography			10,919,939	13.3
Sea-ice			9,057,382	11.0
Cold sectors 1 (shallow maritime convection)			7,976,470	9.7
Emissivity retrieval failed			756,986	0.9
Last model timestep			732,131	0.9
Excessive SI over land			210,088	0.3
Observed value missing			4,863	0.0
Negative humidities in model			278	0.0
Surface emissivity too variable			159	0.0
After basic QC	34,422,618	41.9		
First guess check			958,670	1.2
VarQC			336,197	0.4
After all QC	33,127,751	40.3		

important check is for inhomogeneous surfaces, which rejects locations where the model grid-point contains a mixture of water and land. This is applied to all channels but could be relaxed in the future for the higher peaking channels, given how many observations are lost to it (around 14%).

Less significant groups of rejections come from the failure to generate an emissivity (0.9%; land and sea-ice surfaces only; see Sec. 2.7) and when (over land only) we avoid assimilating strongly convective situations. Here we remove observations with a symmetric SI greater than 20 K; these situations exhibit mean FG departure biases around +20 K, suggesting a problem with the forecast model or the observation operator (0.3% rejected, see Baordo et al., 2013). There are also a number of essentially technical issues that cause small amounts of data rejection, the most significant (at 0.9%) coming from the last 15 minutes of the assimilation time window, where for performance reasons the model physics is switched off and hydrometeor inputs to the observation operator are not available.

Figure 2 (left column) examines the histograms of FG departure after 'basic' QC, after the first guess check and after VarQC. There are FG departures as large as 150 K in the humidity sounding channels; these are situations where convection is present in the model and absent in the observations, or vice-versa. Most of these situations are excluded by the first guess check which rejects normalised departures greater than 3 (MHS channels 3 and 4) or 2.5 (MHS channel 5 and all SSMIS channels - more consistency would be desirable). In MHS channel 3, where the highest possible assigned observation error is about 10 K (Fig. 1) that means all FG departures greater than around 30 K are rejected. Likely as a result of this tight first guess check, VarQC does not eliminate much additional data and the histograms before and after VarQC are nearly on top of each other. This is a mixed sign: it means that most observations that got past the basic QC and first guess check were in good agreement with other observations and with the final analysis. However, it also motivates further work on the first guess check and observation error model.

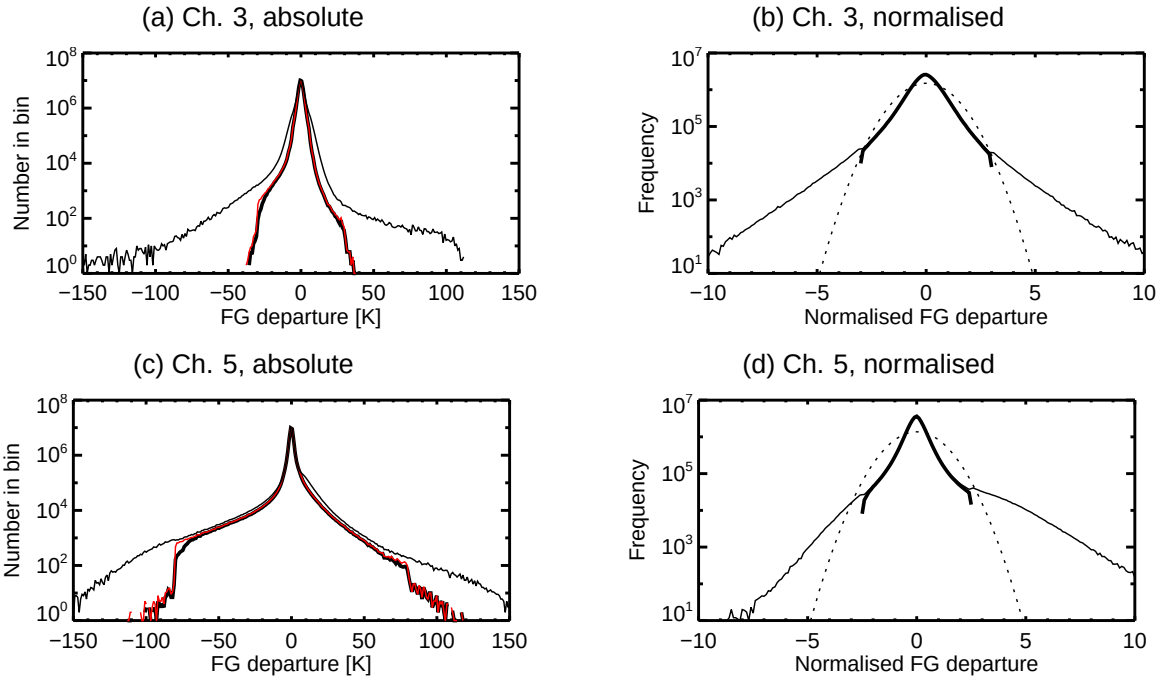

Figure 2: Histograms of MHS FG departure in terms of absolute brightness temperature (left column) or normalised by the observation error (right column). The thin black line is for all observations passing basic QC checks; thin red line after the first guess check (left column only) and thick black line after VarQC, i.e. the observations actually assimilated. On the right column only, the dotted line shows the Gaussian distribution with a standard deviation equivalent to that of the sample passing basic QC checks. The sample is 6 months of MHS observations.

The right column in Fig. 2 shows the FG departure normalised by the observation error, after basic QC and after first guess check and VarQC. This second histogram (the assimilated data) is truncated within the range imposed by the first guess check. The symmetric observation error model is broadly doing its job, which is to transform a highly non-Gaussian PDF (left column) into a more Gaussian PDF to satisfy the requirements of data assimilation algorithms for Gaussian error models. The PDF is not perfect, being overly peaked and also having large tails which are thrown away by the first guess check. The excessive peaking suggests our data usage is too cautious, as a lot of observations are probably getting less weight (i.e. they have a larger observation error) than they may deserve. SSMIS results look similar (Geer, 2013). Particularly in the upper-tropospheric channels, displacements in humidity features can cause brightness temperature signals larger than many cloud and precipitation signals. The SI-based error model is not perfect because it assumes convection is the main driver of error. We are capturing the dominant source of error, but clear-sky humidity errors are also important and are not yet treated in this error model. Possibly they can be addressed in the future as background error using estimates from the EDA (Bormann and Bonavita, 2013). In the lower channels it might also be useful to better model the errors coming in cloud and precipitation in non-convective frontal areas, which are not perfectly captured by a scattering predictor.

2.6 Bias correction

Biases are corrected adaptively using VarBC (Dee, 2004; Auligné et al., 2007), though there is no attempt to correct cloud and precipitation-related biases. Cloud and precipitation biases are best addressed through improvements to the radiative transfer or the forecast model and failing that, quality control should be applied to prevent the biased observations being assimilated. The bias for the humidity sounding channels is modelled as a constant offset plus a linear function of layer thicknesses (1000 - 300 hPa; 200 - 50 hPa and 10 - 1 hPa) and a fourth order polynomial in scan position. This is the configuration that has been used in clear-sky humidity sounding for many years. For MHS, bias corrections in the lower two sounding channels are very small (well within ± 0.5 K) and map onto air-mass and scan-dependent components (not shown). The upper humidity sounding channel has a positive bias around 1 K with strong pole-to-tropics variation, i.e. an air-mass dependence that might suggest instrument calibration issues or biases in the upper-tropospheric humidity in the forecast model. However, the biases are not consistent with SSMIS, which might suggest instrumental biases in either MHS or SSMIS. The 90 GHz and 150 GHz channels are bias corrected as microwave imager channels which means the bias is modelled as a function of skin temperature, total column water vapour and wind-speed, in addition to the scan-bias. Generally the MHS and SSMIS bias correction is stable in time and not much affected by changing from clear-sky to all-sky assimilation.

2.7 Surface emissivity from atlas and retrievals

Accurate surface emissivity (and/or skin temperature) is important in correctly simulating the lower-peaking 183 GHz channels. For example, over land, after quality control that includes the removal of high altitude and snow surfaces (Tab. 3) MHS channel 5 has surface-to-space transmittances a greater than 0.2 in roughly 5% of situations. These situations are caused by low relative humidity in the troposphere and they are just as prevalent in the tropics as the midlatitudes. With this level of surface-to-space transmittance, the emissivity needs to be known to within ± 0.1 to limit the errors in simulated brightness temperatures to below 1 K (following English, 2008). Ocean surface emissivity models are very accurate, but the same accuracy is not possible over land or sea-ice. To get a surface emissivity for land and sea-ice we attempt an emissivity retrieval (e.g. Jones and Vonder Haar, 1990; Prigent et al., 1997; Karbou et al.,

2005). If that fails we use emissivities from the TELSEM atlas (Aires et al., 2011), though that is not possible over sea-ice.

Over land, we retrieve emissivity in a 90 GHz channel and apply that to the 183 GHz channels, assuming that the spectral variability of emissivity is minimal (see Tab. 1). The retrieval and quality control of all-sky land surface emissivities will be fully documented in Baordo and Geer (2014) following initial work described in Karbou et al. (2008); Baordo et al. (2012, 2013). The all-sky emissivity retrievals use the first guess atmospheric profiles and skin temperature and take into account cloud absorption but not scattering. Hence, errors in the FG skin temperature and atmospheric profile lead to errors in the emissivity retrievals. The most obvious retrieval errors are associated with cloud and precipitation. The retrievals are subject to errors when either the model has much more cloud than the observations (leading to unphysical emissivities greater than 1) or when the observations have much more cloud than the model (leading to unphysically low land emissivities, e.g. less than 0.4). Hence, a quality control is applied to the retrievals, which are used only when the difference between the emissivity retrieval and the equivalent value from the TELSEM atlas is less than 0.09 (SSMIS) or 0.07 (MHS). These limits are roughly three times the standard deviation of differences between carefully screened clear-sky retrievals and the corresponding atlas values. A further check ensures the emissivity is less than or equal to 1 and greater than 0.55. If the retrieval is rejected, the TELSEM atlas value for 183 GHz is used.

Over sea-ice, the atmosphere is transparent enough to use a 150 GHz channel for emissivity retrieval (e.g. Di Tomaso et al., 2013). This is useful because the spectral variability of surface emissivity over snow and ice can be very large (e.g. Hewison and English, 1999). However, there is no emissivity atlas for backup or quality control. Any retrieval between 0.65 and 1.0 is used; otherwise the observation is rejected. Overall the emissivity retrieval is much less reliable than over land and there are obvious biases over sea-ice in the FG departures of the lowest peaking 183 GHz channel (e.g. MHS channel 5). For this reason, these channels are not used over sea-ice.

Generally the mean FG departures of assimilated observations show no sign of problems coming from inaccurate land or sea-ice surface emissivities. Aside from the inclusion of cloud effects in the emissivity retrieval, the surface emissivity treatment broadly replicates the framework used in the previous clear-sky assimilation of humidity sounding channels, which has been well-proven in operational clear-sky assimilation (Krzeminski et al., 2009; Di Tomaso et al., 2013). There are even indications (e.g. Baordo and Geer, 2014) that the all-sky approach slightly reduces FG departure biases over land compared to clear-sky assimilation.

2.8 Moist physics modelling: nonlinear and linearised

As explained in the introduction, moist physics modelling is crucial to using cloud and precipitation from all-sky radiances. There are three types of moist physics model in the ECMWF system. The first level runs in the high-resolution forecast model, using complex and highly nonlinear algorithms (e.g. Forbes et al., 2011; Bechtold et al., 2013). Prognostic variables are used for large-scale cloud and precipitation, but the convection scheme is diagnostic. This package can be termed the full non-linear physics (FNL). For data assimilation, a simplified non-linear (SNL) physics scheme is required, where minor processes are discarded and particularly nonlinear behaviour is smoothed (Tompkins and Janisková, 2004; Lopez and Moreau, 2005; Janisková and Lopez, 2013). The tangent linear and adjoint of the simplified non-linear scheme form the third level; these are used in the 4D-Var minimisation. Much care is taken to ensure the SNL and FNL schemes produce similar results, both in the forward and tangent-linear sense; generally this can be achieved remarkably well out to forecast ranges of about 48 h.

To understand all-sky data assimilation, the diagnostic nature of the SNL scheme is important. Essentially, cloud and precipitation is diagnosed at each timestep from the temperature and moisture profile (plus some surface quantities such as moisture and heat flux). Cloud and precipitation is not advected and it does not persist until the next model timestep. The 4D-Var control vector (surface pressure, wind, temperature and humidity only) is valid at the beginning of the first model timestep of the 4D-Var window. The diagnostic moist physics operators then generate cloud and precipitation fields valid at the same time. Hence it is possible to adjust cloud and precipitation to fit all-sky observations even at the beginning of the assimilation window, but only by adjusting the temperature and moisture profile at the observation location. Background error correlations can of course transfer these adjustments into other variables and other locations.

Cloud-tracing wind sensitivity develops in the second timestep after the dynamical model timestep has been run and the humidity field has been advected. In an adjoint (reverse) sense, the process works as follows: (a) cloud and precipitation gradients are generated from the adjoint of the observation operator; (b) these are passed (along with contributions from any other colocated cloud or precipitation sensitive observations) into the adjoint of the SNL physics, which generates gradients in the temperature and moisture field at the same time and location; (c) these temperature and humidity gradients are added to all other gradients (coming from later timesteps and from other observations) and passed into the adjoint of the dynamical model, which translates them into gradients in humidity, temperature, wind and surface pressure at the beginning of the first timestep; (d) the adjoint of the control variable transforms (e.g Bannister, 2008) converts these dynamical variables into the uncorrelated parameters that are the true control variable for 4D-Var.

Note that a cloud control variable is still an important goal of cloud assimilation efforts at ECMWF and is necessary when the moist physics in the data assimilation is fully prognostic, but the viability of the current all-sky assimilation shows it is not essential. As has been seen in previous single observation test cases (Bauer et al., 2010), 4D-Var is capable of modifying wind and mass fields at the beginning of the time window to better fit cloud and precipitation observations. The incremental formulation of 4D-Var is particularly important, meaning that the tangent-linear and adjoint model can be relinearised against a steadily improving model trajectory. Further single observation cases are shown in this study that also clearly demonstrate the adjustment of winds and synoptic structures to fit cloud and precipitation features in the observations.

3 Understanding the observations

Figure 3 shows the observations from the microwave humidity sounder (MHS) on the European Metop-B satellite. These are the observations available in the 00 UTC assimilation window on 15 August 2013, during the southern hemisphere winter where synoptic activity in the southern storm track is at its height. The top panel shows MHS channel 3 (at 183$\pm$1 GHz), sensitive to upper-tropospheric humidity (e.g. Buehler and John, 2005) and ice cloud and precipitation (e.g. Sreerekha et al., 2008; Hong et al., 2005). In the absence of cloud or precipitation, low relative humidity corresponds to a high observed brightness temperature and vice-versa. The climatological subtropical high pressure regions are marked by high brightness temperatures indicating a very dry free-troposphere. In the mid-latitudes and in the inter-tropical convergence zone (ITCZ) there is much higher relative humidity and the brightness temperatures are relatively low. Also, in the mid-latitudes there is great variability and a mixture of dry and moist air masses. The main intention of assimilating the humidity sounding channels is that the movement, position and water amounts of these air masses can be used to infer information on the wind fields and ultimately on the large-scale synoptic structures of the atmosphere. The all-sky brightness temperatures

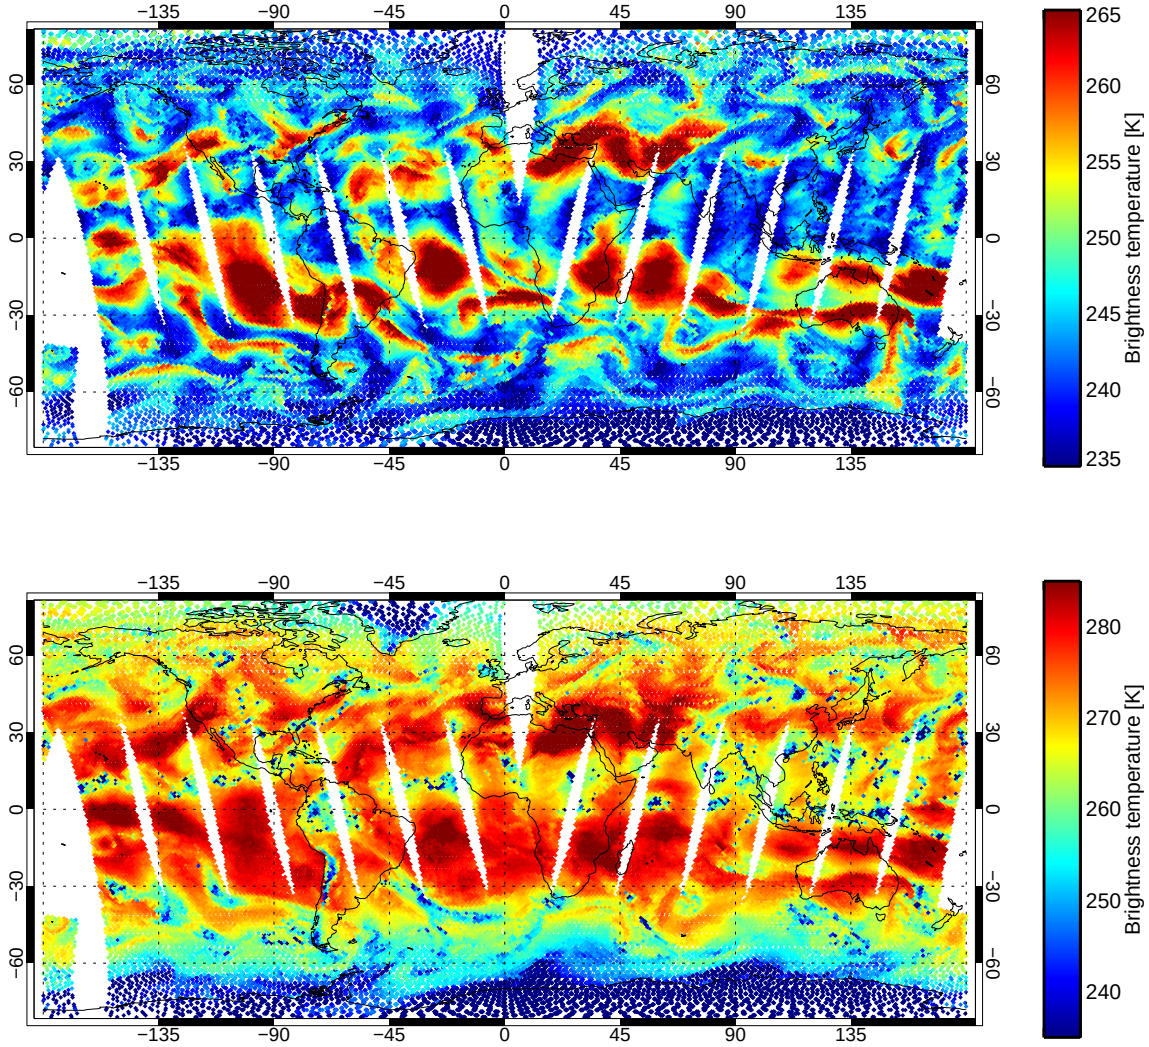

Figure 3: Observed brightness temperatures from Metop-B MHS in the 00 UTC assimilation window on 15 August 2013: (a) Channel 3, at 183±1 GHz, sensitive to upper-tropospheric humidity, cloud ice and frozen precipitation; (b) Channel 5, at 190 GHz, sensitive to mid-tropospheric humidity, cloud ice and frozen precipitation. High brightness temperatures correspond to low humidity; low brightness temperatures correspond to high humidity, cloud ice or frozen precipitation (typically deep convection).

include cloudy and precipitating scenes. These are hard to distinguish in the figure, at least in a global view. At the microwave frequencies used by MHS, cloud ice and frozen precipitation cause scattering, which reduces brightness temperatures. High relative humidity also reduces brightness temperatures: without additional information it is hard to distinguish cloudy or deep-convective areas from those which just have high relative humidity. In channel 3 the effect of humidity on brightness temperature is often larger than that of cloud and precipitation.

Figure 3b shows MHS channel 5 (at 190 GHz), which peaks lower in the atmosphere, sometimes low enough to observe the surface. Again humidity is the main factor controlling the observed brightness temperature, but cloud and precipitation are more easily distinguished against the warmer clear-sky background: spots of localised low brightness temperature (the dark blue colours) indicate deep-convective systems or broader areas of thick ice cloud, for example in the ITCZ and in mid-latitude frontal zones.

For the set of observations that are assimilated, Fig. 4 shows the normalised all-sky first-guess departures (first-guess departure divided by observation error) corresponding to the observations from MHS channels 3 and 5. This is a fundamental quantity in the data assimilation system because, when squared, it gives the contribution of that observation to the 4D-Var cost function. In other words, the normalised first-guess departure is a guide to the influence of that observation in the data assimilation, though its ultimate effect in the analysis is also controlled by the background error term. The largest normalised departures are in the mid-latitude storm tracks and in the ITCZ: hence these are the areas where we might expect to have greatest influence on the analysis.

Given that only active observations are shown in Fig. 4, and as already seen in Sec 2.5, in the lower-peaking channels the quality control removes a lot of data. Overall, we might expect the upper-tropospheric channel to have the greatest impact on the data assimilation because it has a much greater geographical coverage. Looking closely at the departures in the tropics, there are positive and negative regions corresponding to displacements in the position of large convective systems in the ITCZ. It is still very hard for the model to simulate these features in the right place and at the right time, even in the 12 h first-guess forecast. In the mid-latitudes, there are elongated regions of negative and positive departures with a width of the order 100 to 300 km and lengths up to 1000 km. These correspond to displacements and intensity variations in the mid and upper-tropospheric humidity and cloud fields. This is the information we hope will lead back to the wind fields and ultimately to improvements in the analysis of the large-scale synoptic situation.

One of the key advantages of the all-sky approach is likely to be its enhanced coverage in meteorologically interesting areas. Figure 5 shows the number of MHS channel 3 observations that can be assimilated using either a clear-sky or an all-sky approach. In either case, coverage is lowest in the northern hemisphere, mainly because high altitude and snow-covered land surfaces need to be discarded. But polewards of 50°S, across the southern hemisphere storm tracks, the all-sky technique provides around double the number of observations. In the clear-sky technique, cloud screening removes a majority of observations in the southern high latitudes. Comparing the number of all-sky observations assimilated to the number available also highlights a remaining challenge for microwave sounding observations: filling in the gaps over high-altitude and snow-covered land surfaces, which are now responsible for the majority of data rejections.

Figure 6 shows the mean FG departures of assimilated observations over the 6 month experimental period, presented either raw or normalised by the prescribed observation error. The forecast model lacks liquid water in boundary-layer maritime cloud and its diurnal cycle is not realistic (Kazumori et al., 2014). Hence there are systematic negative FG departures in the subtropical oceans in the 190 GHz channel (Fig. 6e and f), when the free troposphere is very dry, can sense deep enough into the troposphere

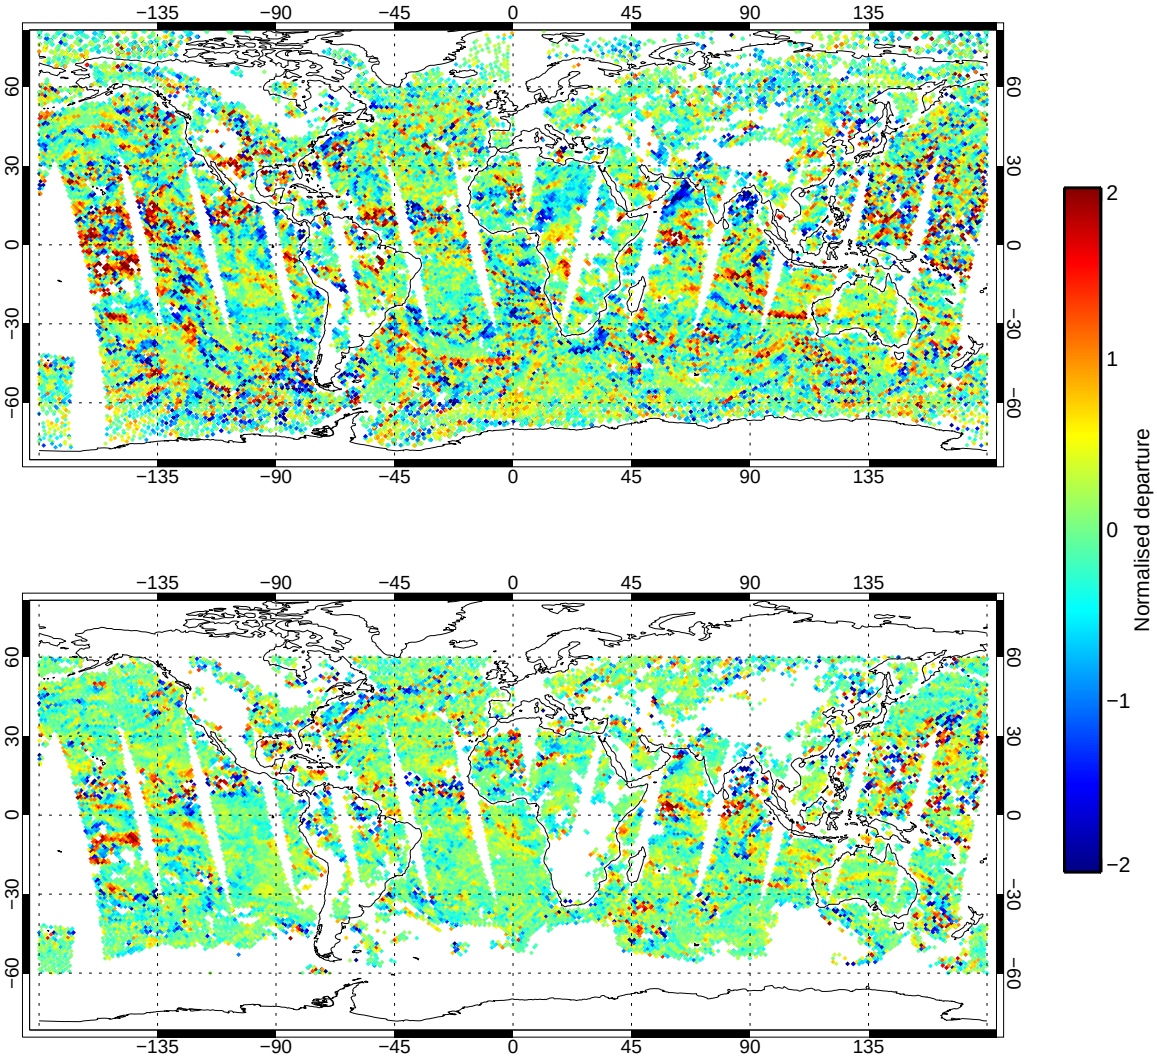

Figure 4: Normalised, bias corrected first-guess departures from Metop-B MHS in the 00 UTC assimilation window on 15 August 2013: (a) Channel 3, sensitive to the upper troposphere; (b) Channel 5, sensitive to the mid or lower troposphere. Normalised first-guess departures are computed as the observation minus the first guess divided by the observation error. Squared, these would give the observation's weight in the 4D-Var cost function; as they are, they retain information on the direction in which the analysis should move: either moistening or drying the model. The sample is all assimilated observations.

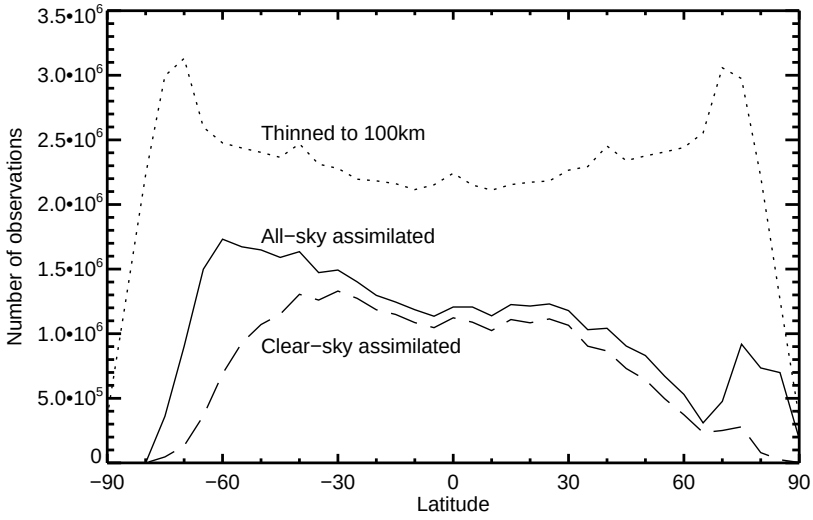

Figure 5: Number of MHS channel 3 observations actively assimilated, per 5° latitude bin, in 6 months of combined winter and summer experiments (solid, 37 million obs). Also shown are the total available observations after grid-based thinning to 100 km separation (dotted, 82 million obs) and if clear-sky quality control had been applied to remove cloudy scenes (dashed, 27 million obs).

to observe the maritime boundary layer. The same biases affect the microwave imager channels and drive the model to create more humidity and cloud. Experience with the microwave imager assimilation shows that these systematic changes in the subtropics have surprisingly little effect on tropical and midlatitude forecasts; in other words, they do not trigger growing errors in the forecast and they are not dangerous to forecast scores (Geer and Bauer, 2010).

Another obvious feature is the systematic positive departures in the lowest-peaking channel in the mid-latitudes (Fig. 6e) that are not visible in the normalised departures (Fig. 6f), which at least means they cannot affect the forecasts. Because of the way the observation errors are generated, this indicates that the bias is limited to areas with high assigned observation errors, which in the midlatitudes means frontal regions. The version of the forecast model used in these experiments systematically overestimates liquid water path in frontal areas. This problem has been corrected in the latest model cycle, 40r3, and the systematic departures in frontal areas in microwave imager and humidity sounder observations are much reduced.

A second area of positive departures is present in all channels around the Andes. Note that assimilation is not done in the highest parts of the Andes; quality control removes observations where the surface altitude is high (see Sec. 2.5). An obvious but incorrect inference is that the bias is associated with surface visibility. In the areas where this bias is present, even in the lowest peaking channel, a surface to space transmittance of 0.01 is extremely rare. Further evidence against this hypothesis is the fact that the bias is quite consistent in all channels. A second incorrect hypothesis is that the bias is associated with deep convection. If the deep-convective and heavy cloud observations are removed, the positive bias around the Andes remains (not shown, this is done by eliminating all scenes where either observed or modelled scattering index exceeds 5 K). Most likely there is a small moist bias in the upper troposphere in the forecast model in these areas.

Finally, there are negative biases in tropical convective regions, largest in the western equatorial pacific.

(a) Mean departure, 183±1 GHz

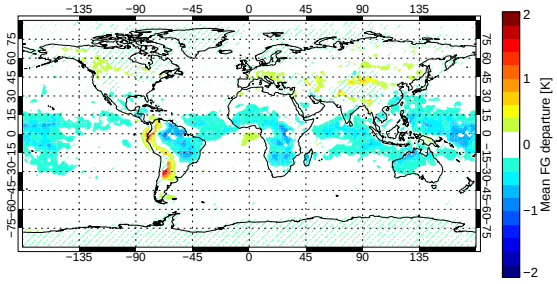

(b) Mean normalised, 183±1 GHz

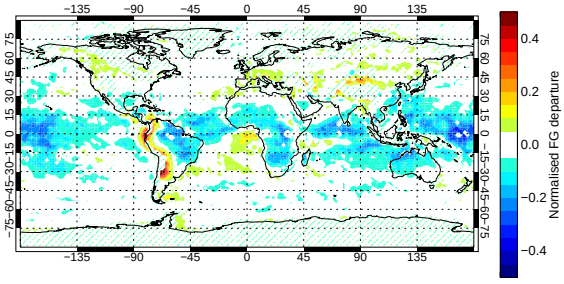

(c) Mean departure, 183±3 GHz

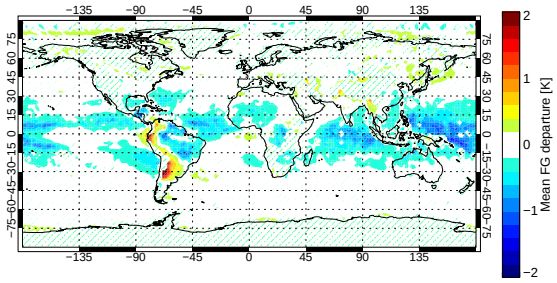

(d) Mean normalised, 183±3 GHz

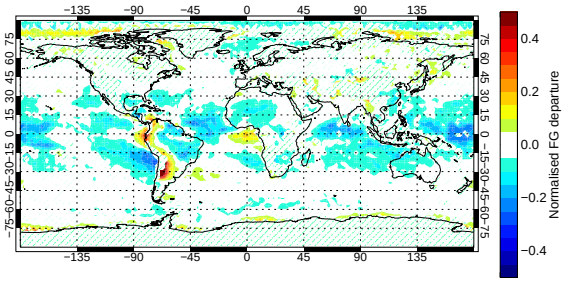

(e) Mean departure, 190 GHz

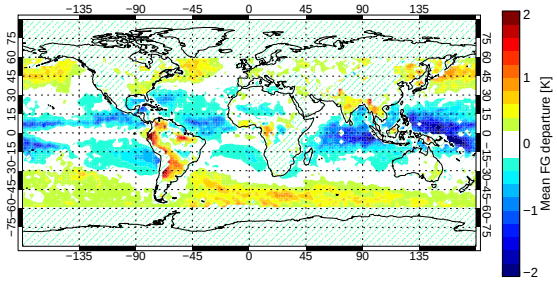

(f) Mean normalised, 190 GHz

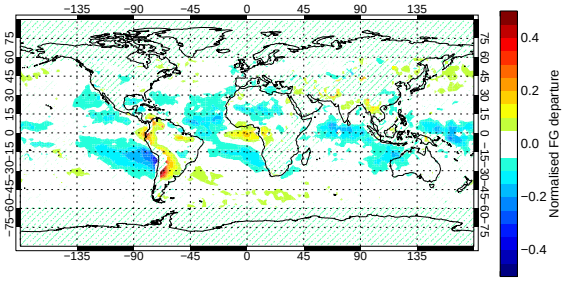

Figure 6: Mean of FG departures for actively assimilated MHS observations over the combined 6 month period, based on a total of 37 million observations in the highest-peaking channel but fewer in lower peaking channels. The blue diagonal lines indicate areas where no data has been assimilated. Left column: absolute FG departures in K; Right column: FG departures normalised by the assigned observation error.

One part of this bias probably comes from under-scattering in the observation operator: frozen particles in the upper parts of deep convective clouds should cause more scattering, and so they should ideally produce lower simulated brightness temperatures. In choosing a particle shape to represent scattering from frozen hydrometeors that produced good results at all frequencies from 10 GHz to 183 GHz, Geer and Baordo (2014) had to make compromises that they acknowledge resulted in a slight under-scattering at high frequencies in deep convective areas. However, most of those scattering areas are assigned high observation errors. The remaining scattering biases are small in respect of the observation error and are not too important in the data assimilation. So it is curious the negative biases around the equator become more obvious in the normalised departures, particularly in the upper-troposphere channel (panel b). The explanation is that negative biases also occur around the tropical convective systems in regions that are apparently free from significant cloud, in model or observations. Here, the assigned observation error is small. There are three hypotheses that could explain these biases: first, the model may be deficient in cirrus detrainment and there should be more high ice cloud near the convective systems; second, scattering from this ice cloud may not be strong enough in the observation operator; third, these truly are clear-sky areas and the model humidities surrounding convective areas may not well modelled. These negative biases (first guess too warm means too little moisture in the model) could be explained by too much dry detrainment in the tropical upper troposphere, but much more work would be needed to confirm this. As shown in later sections, these uncorrected tropical upper troposphere biases, when assimilated, do have an impact on the mean state of the model, but verified against other observations the impact is strongly positive. The impact is visible in all temperature and wind observing systems in the upper troposphere and lower stratosphere. That would perhaps favour the third hypothesis, on dry detrainment. However, a major caveat is that biases can change with model resolution and with every new version of the model convection scheme (e.g. Bechtold et al., 2013). Hence, it is necessary to monitor these biases closely and in collaboration with model physics developers. This is a strategy that has been applied for many years to the cold-sector bias in the microwave imager channels, and which may soon lead to an improved model representation of supercooled liquid water in high-latitude shallow cumulus (Richard Forbes, personal communication). Such improvements to the forecast model are an additional benefit of all-sky assimilation.

4 Improving the clear-sky assimilation of microwave humidity sounding channels

This study seeks to separate the clear-sky and all-sky contributions to the assimilation of microwave humidity sounders. The clear-sky assimilation produces useful forecast benefits but that has not always been the case. Before cycle 39r1, microwave humidity sounders produced very little impact on forecast scores at ECMWF. At the time, the only information at 183 GHz came from MHS (Metop-A,B NOAA-18,19) and assimilation was restricted to land and ocean surfaces. Recent developments have increased the amount of 183 GHz observations used in clear-sky conditions, or have given the potential to do so:

- Extension of MHS assimilation to sea-ice surfaces and cold oceans, i.e. sea surface temperature less than 278 K (Di Tomaso et al., 2013).
- Ability to assimilate SSMIS humidity channels over ocean (Geer, 2013), land (Baordo et al., 2013) and sea-ice.
- Allowing the assimilation of the 18 outermost scan positions of MHS. These were previously blacklisted because of fears for their quality that turned out to be unfounded.

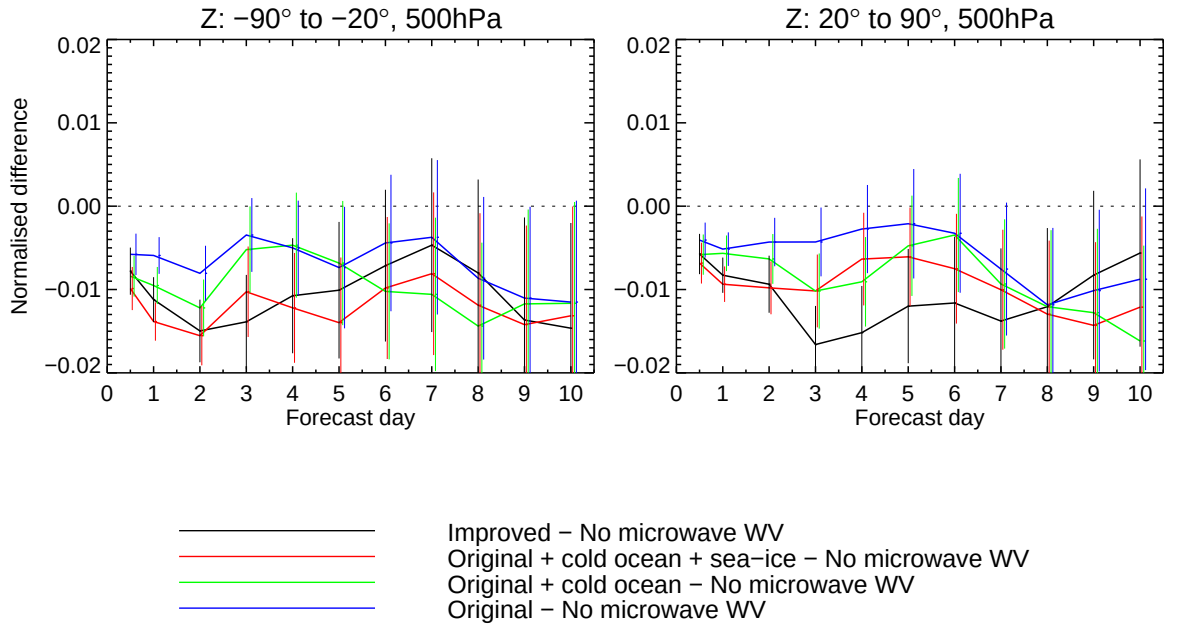

Figure 7: Normalised change in RMS of geopotential (Z) forecast errors when clear-sky microwave humidity assimilation is added to the otherwise full observing system. Error bars indicate statistical significance at the 95% confidence level. Results are based on a total of six months experimentation, winter and summer, and are based on 322 to 360 forecasts depending on the forecast range. Verification is against own analysis.

The latter two changes were made as part of all-sky developments described here. In other words, trying to implement the all-sky assimilation has shown us how to do a better job of clear-sky assimilation. The complete package of microwave humidity assimilation updates in cycle 40r3 includes changes that simply improve the use of clear-sky humidity information, but these are a significant part of the package. To test the all-sky contribution specifically, the appropriate reference is the best available clear-sky assimilation, not the less-capable version used in the previous operational system. Because it will not be examined elsewhere, this short section explores the forecast impact of recent improvements to clear-sky assimilation.

Figures 7 and 8 illustrate forecast impact compared to a control from which microwave humidity sounding channels have been excluded, called 'No Microwave WV'. The experiment titled 'Original' replicates the pre-39r1 version of 183 GHz assimilation, using four MHS instruments over land and ocean. This is enhanced by adding cold ocean surfaces ('Original + cold ocean') and sea-ice areas ('Original + cold ocean + sea-ice') as described in Di Tomaso et al. (2013). Finally, the 'Improved' clear-sky experiment adds the 18 outer scan positions of MHS (bringing about 50% extra observations) and the 183 GHz channels of SSMIS F-17 over ocean, land and sea-ice. This last experiment was created in the all-sky framework by turning off the cloud and precipitation radiative transfer and imposing the quality control checks necessary to remove cloud-affected scenes. There are many minor differences in data usage between clear-sky and all-sky frameworks (Bauer et al., 2010) so the comparison to 'Original + cold ocean + sea ice' is not completely clean. The most important of these minor differences is around 10% additional observations that are assimilated in 'Improved', because the grid-based thinning selects slightly more data than the more 'random' thinning algorithm used for clear-sky observations. The grid-based thinning gives around 100 km distance between neighbouring observations and there has been no attempt

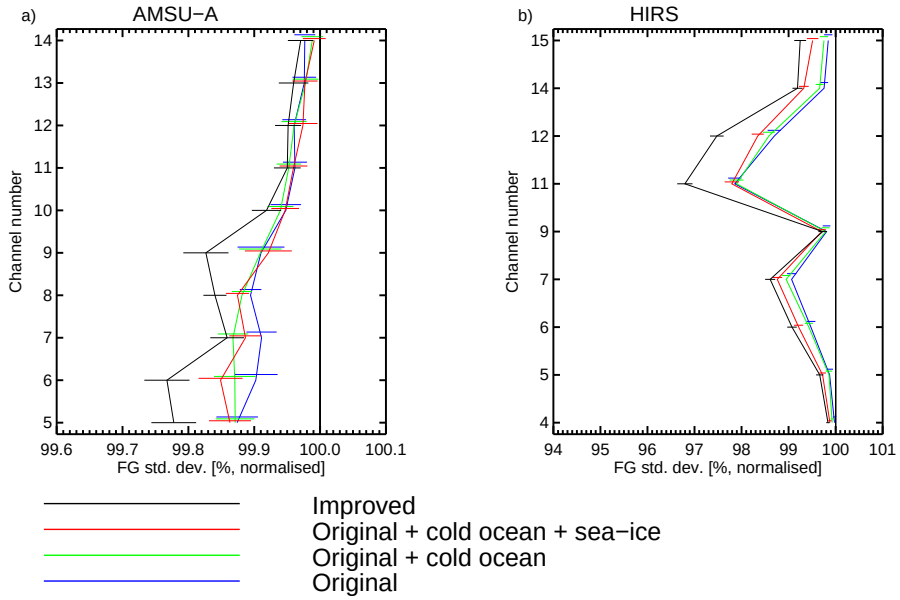

Figure 8: Standard deviation of (a) AMSU-A and (b) HIRS first guess departures, given as a percentage relative to the control, and based on a global sample of assimilated observations. Error bars indicate the 95% confidence range. Control is without microwave humidity sounders but otherwise contains the full observing system, and it sets the 100% line on this figure.

to make this exactly consistent with the old approach. Overall, 'Improved' clear-sky roughly doubles the number of assimilated 183 GHz observations compared to 'Original + cold ocean + sea-ice'.

The 'Improved' clear-sky configuration has roughly twice the impact of 'Original' on forecast scores (Fig. 7) and short-range tropospheric temperature fields (e.g. fits to the microwave temperature sounder AMSU-A in channels 5 to 9, Fig. 8a). However, there is only 50% more impact in the short-range moisture fields (the upper-tropospheric water vapour channels 11 and 12 of the infrared sensor HIRS, Fig. 8b). The forecast impacts are statistically significant out to day 4 (SH) and 6 (NH). It is harder to measure the forecast impacts of the incremental improvements in data usage, such as adding cold oceans and sea-ice. The six-month period of experimentation is not long enough to reduce the error bars to the point where these experiments can be reliably distinguished from one another. However, there is statistical significance in the observation fits, where the 'Improved' version is clearly best, suggesting that doubling the number of observations (without changing the geographical coverage) is probably the most significant development.

5 All-sky impact in the absence of other observations

5.1 Single-observation experiments

To illustrate how 4D-Var assimilates the all-sky observations, fourteen single-observation test cases have been chosen (Fig. 9). The sample was picked by hand, looking for unusual situations with large first guess departures. These are good examples because the analysis has to make noticeable changes like shifting a precipitation event or adjusting the position of an airmass boundary. All examples have been taken from

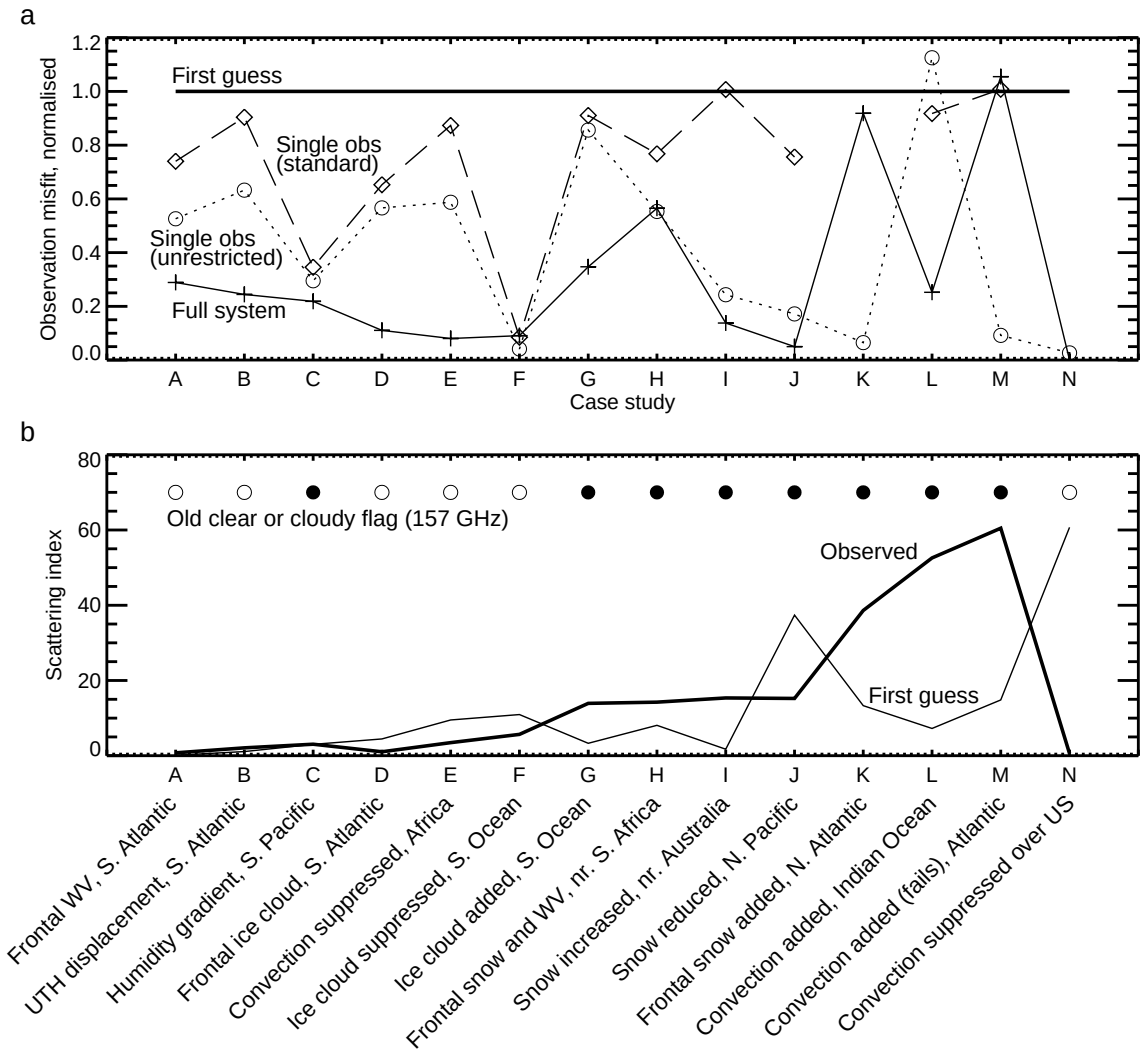

Figure 9: Summary of single observation test cases: (a) Observation misfit in 183 GHz channels normalised by the misfit at first guess; (b) Scattering index in observations and first guess, plus the results that of the cloud detection in the old clear-sky approach (open circle indicates 'clear' and filled circle 'cloudy').

the Metop-B observations assimilated in the 00Z analysis on 15th August 2013. The experiments have been ordered by scattering index (the maximum of the observation and the first guess, Fig. 9b) so that the first are typical of clear-sky and smaller cloud and precipitation amounts, whereas the last involve strong convection in the model or observations. There is a mix of land and ocean cases across the tropics and midlatitudes.

Globally, the features with the largest 183 GHz FG departures fall into these categories:

- Displacements in upper-tropospheric airmasses in the extratropics, where large departures are found along sharp upper-tropospheric humidity (UTH) boundaries (cases A-D);
- Displacement and intensity errors in ice cloud and large-scale precipitation associated with fronts in the extratropics (cases F-K);
- Displacement and intensity errors in snowfall in convective areas, most prevalent in the tropics (cases E and L-N).

Three configurations are available for each test case. The reference configuration is the analysis generated by the full observing system. Then single observation experiments have been performed either with the usual observation errors and QC ('standard'), or with relaxed QC and no inflation of observation errors in cloudy situations ('unrestricted'). These latter experiments use the following adaptations:

- The first guess check has been relaxed to 5 times observation errors and VarQC has been turned off;
- Screening of some difficult situations has been deactivated: high scattering index situations over land and cold-air outbreaks over ocean are now assimilated.
- Observation errors are always set to clear-sky values, even in cloud and precipitation (over ocean, 2 K in channels 3 and 4 and 2.2 K in channel 5; over land, 3 K);

Figure 9a quantifies the fit to the MHS observations in terms of a norm of 183 GHz channel departures,

$$J = \sum_{i=3,4,5} (\mathrm{Obs}_i - \mathrm{FG}_i)^2. \tag{7}$$

where i is the channel number. What is actually plotted is $J_{\mathrm{Analysis}}/J_{\mathrm{FG}}$, i.e. the fractional reduction in error between FG and analysis. The norm is reminiscent of the observation cost-function in 4D-Var but it ignores the observation error. In most of the examples, the full observing system improves the fit to observations by 70% to 90%. This is not representative and it results from deliberately choosing cases where the full observing system was successful in fitting the MHS observations. However, two convective examples (K and N) have been chosen to represent more difficult situations where the analysis from the full observing system did not improve the FG, and other examples also illustrate the limitations of the system as well as its benefits.

The examples span clear-sky and all-sky cases. Figure 9b shows the results of the old cloud screening, based on a 5 K threshold for clear-sky FG departures in channel 2 (157 GHz). This can be compared to the scattering index, which indicates the presence of cloud ice or snow. In all cases where the observed scattering index is greater than 10 K (G-M), cloud-screening would also have identified cloud and would have thrown out the observation. However, there are cases where the cloud-screening would have failed to identify significant cloud contamination (cases E and F, with observed scattering indexes of 3.5 K and

Table 4: Single observation case C (Metop-B, 36.9°S, 112.6°W, 05:54 UTC 15 Aug 2013.) Channels 1 and 2 are only monitored, not assimilated. TCWV, CWP, IWP, RWP and SWP refer to total column amounts of water vapour, cloud water, cloud ice, rain and snow.

		Analysis	
	FG	Single-obs (normal)	Full system
Departures [K]			
1: 89 GHz	2.43	0.39	1.28
2: 157 GHz	2.39	1.38	1.63
3: 183±1 GHz	2.96	1.99	1.54
4: 183±3 GHz	1.38	0.25	0.40
5: 190 GHz	1.03	0.16	0.14
Total columns [kg m^{-2}]			
TCWV	13.1	12.8	13.0
CWP	0.03	0.05	0.03
IWP	0.04	0.03	0.02
RWP	0.00	0.00	0.01
SWP	0.07	0.06	0.04

5.7 K respectively). In retrospect, the scattering index could have been used as an additional, more precise test of ice cloud and snow contamination in the mid- and upper-tropospheric WV channels (channels 3 and 4). The 157 GHz clear-sky departures (channel 2) are (a) sensitive to low-level cloud and precipitation features invisible in the upper channels and (b) affected by compensating errors, such as a cloudy observations masked by excessive relative humidity in the FG. Significant effort would be required to try to improve MHS cloud detection algorithms (e.g. Zou et al., 2013) but one of the advantages of all-sky assimilation is that we can abandon cloud-screening and along with it any problems of residual cloud contamination. A second advantage is demonstrated in cases E, F and N, where there are clear sky observations but significant cloud or precipitation in the model. The all-sky approach gives a mechanism to remove the unwanted cloud or precipitation from the model.

Cases A - D illustrate the capability of 4D-Var to adjust dynamical fields to fit humidity features in the upper troposphere. Cases A and B feature clear-skies in model and observation but they are interesting because the 'normal' single observation tests do not produce so much impact. In case A, channel 3 was downweighted by VarQC and channel 5 was lost to the cold-air-outbreak screening. In case B, all channels (3-5) were lost to the first guess check or downweighted by VarQC. Removing QC in the 'unrestricted' experiments (note that observation errors were not affected, these being clear-sky examples) the single observation results are much closer to those of the full system. Cases A and B are typical examples of big FG departures (order 5 K to 10 K in TB) in channels 3 and 4 where distinct humidity boundaries are displaced in the first guess compared to observations. There is no cloud or precipitation and no gross error in the observations: the observations are good and should not have been rejected. For the future it would be good to investigate improved specification of observation and background errors to enable the observations to pass QC and remain in the analysis.

Cases C and D illustrate humidity features where small amounts of cloud or snow are present in model or observations, indicated for example by the scattering index. In both cases the single observation produces a reasonable analysis with the normal settings and the relaxed settings are not needed. Case C gets the closest to the full observing system solution; detailed information is given in Tab. 4 and Figs. 9 and 10. The hydrometeors are well forecast and the scattering index is similar in FG and observations. The model has an ice water path (IWP) of 0.04 kg m^{-2} and a snow water path (SWP) of 0.07 kg m^{-2}

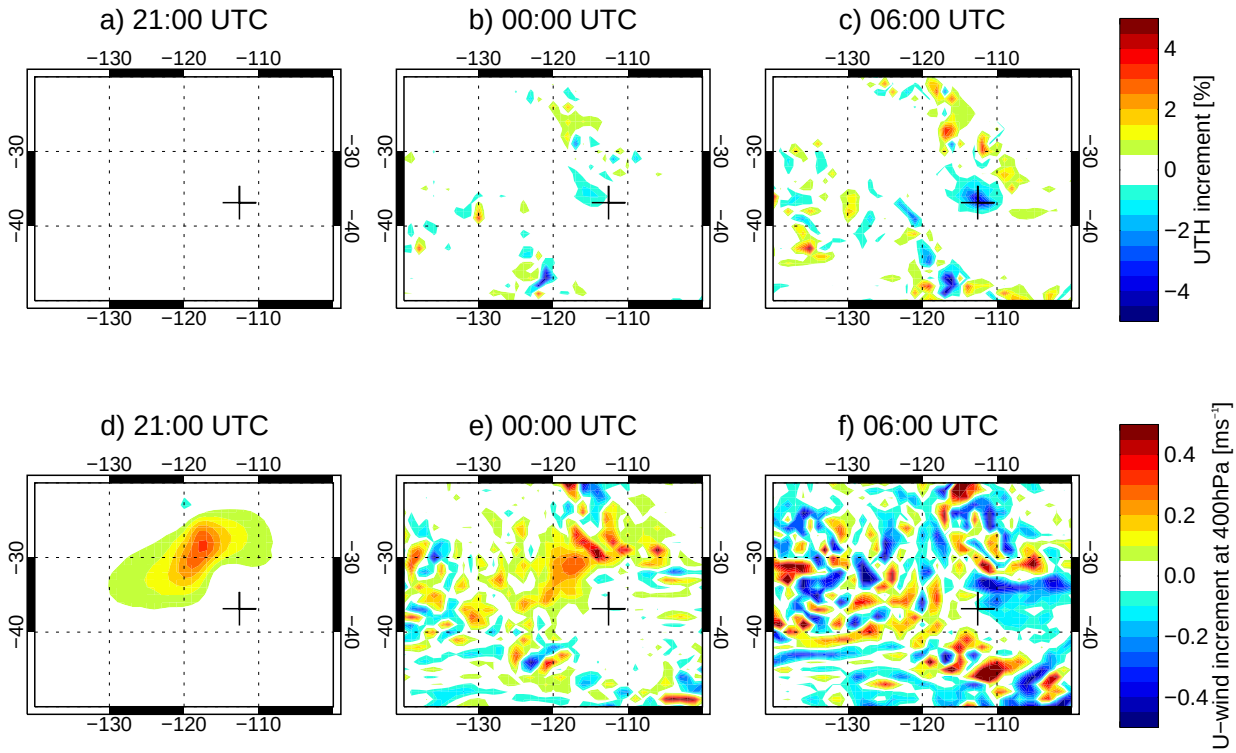

Figure 10: Increments in single observation test case C through the assimilation window; (a-c) UTH; (d-f) u-component wind at 400hPa. 00 UTC is the start of the assimilation window so at 12 UTC and 18 UTC we are seeing the 'evolved' increments. The cross shows the location of the observation, which was made at 05:54 UTC.

and these are slightly reduced from the first guess to analysis. There is also a reduction of 3% in upper tropospheric humidity at the observation location (from 43% to 40%; Fig. 10c). However, these changes in moist quantities have been generated in 4D-Var through a modification to the dynamical fields, not the moist fields, at the start of the assimilation window (compare Figs. 10a and d). This is a robust pattern in all the midlatitude frontal or UTH boundary test cases: at the beginning of the window, changes to wind and geopotential, not humidity, generate the required changes in water vapour, cloud and precipitation at the observation time. This is clear evidence of the model tracing effect in 4D-Var.

Figure 10 also demonstrates an important (and for some surprising) feature of single-observation test cases in the ECMWF system: the increments at the beginning of the window are local to the observation, but the evolved increments (e.g. at 00 UTC and 06 UTC) are not at all local. The background term keeps the increments reasonably local at the beginning of the window (there is some small noise globally, not shown). However, the evolved increments are not constrained except at the time and location of the single observation, so there is nothing to stop them varying freely. In contrast, with a full observing system, the evolved increments are constrained globally and throughout the time window. To create the final analysis, the increments computed using a relatively low resolution linear model are added into the higher-resolution nonlinear model (T511 in this case) and evolved forward in time, so natural chaotic error growth will also start to add features into the evolved increments.

Table 5: Single observation case J (Metop-B, 47.41°N, 159.3°W, 08:00 UTC 15 Aug 2013.) Channels 1 and 2 are only monitored, not assimilated.

		Analysis	
	FG	Single-obs (unrestricted)	Full system
Departures [K]			
1: 89 GHz	-5.71	2.00	-6.36
2: 157 GHz	16.45	4.63	2.11
3: 183±1 GHz	-0.22	-0.63	-0.86
4: 183±3 GHz	3.67	1.09	0.00
5: 190 GHz	15.72	6.57	3.5
Total columns [kg m^{-2}]			
TCWV	30.5	29.3	30.2
CWP	0.11	0.04	0.13
IWP	0.19	0.08	0.07
RWP	0.06	0.03	0.03
SWP	0.97	0.48	0.42

Cases F to I are midlatitude frontal situations with a significant amount of cloud or precipitation in model or observation (Fig. 9b). In case G, the full observing system has fitted the observations well by creating ice cloud and snow, mainly in the lower- and mid-troposphere, but the single observations were not successful. In the normal case, all channels were lost to VarQC or first guess check, but even in the 'unrestricted' case channels 4 and 5 were still eliminated by the more relaxed first guess check. These channels have sensitivity to lower-level cloud and snow; channel 3 on its own does not, and it cannot move the analysis. This illustrates the importance of the lower-peaking channels in sensing snow and cloud ice, and it again gives an incentive to find ways to get the screening to let more good data through. In cases H and I the observations did not move the analysis too much in the normal approach because of inflated observation errors in cloudy areas. In the 'unrestricted' approach they caused the model to create ice cloud and snow to fit the observations better. This is one of the continuing difficulties in all-sky assimilation: with sufficiently low observation errors, we can usually fit the observations quite well, but this can degrade forecast quality (Geer and Bauer, 2011). This is a reason for inflated observation errors in cloudy situations. Normally the observations provide only a push in the right direction, not a firm constraint. For example in case H, the normal approach improves the observation fit by 20% but with smaller observation errors the improvement is increased to 40%. As models get better at representing and forecasting cloud and precipitation, we will slowly be able to increase the weight given to cloud and precipitation in the analysis. However, it should be noted that even with the normal constraints, and in some quite extreme situations, 10 of the 14 single observations improve the analysis by at least 9%. Even with high observation errors and tight quality control, the observations move the analysis in the right direction; used in bulk the observations have much greater impact, as shown in the next section.

Case J is another frontal case, where snowfall occurs ahead of a trough in the N. Pacific (Fig. 11). The channel 5 (190 GHz) departures are +16 K, indicating that the first guess snow amount at the observation time and location is too high (Tab. 5). Interestingly the channel 3 departures (183±1 GHz) are very small and are not significantly changed between FG and analysis: as mentioned earlier this channel is high enough that it is often insensitive to midlatitude precipitation features. With normal settings, neither the first guess check nor VarQC are triggered and the single observation provides a useful push in the right direction, improving the fit to observations by 25% (Fig. 9). With reduced observation errors, the analysis

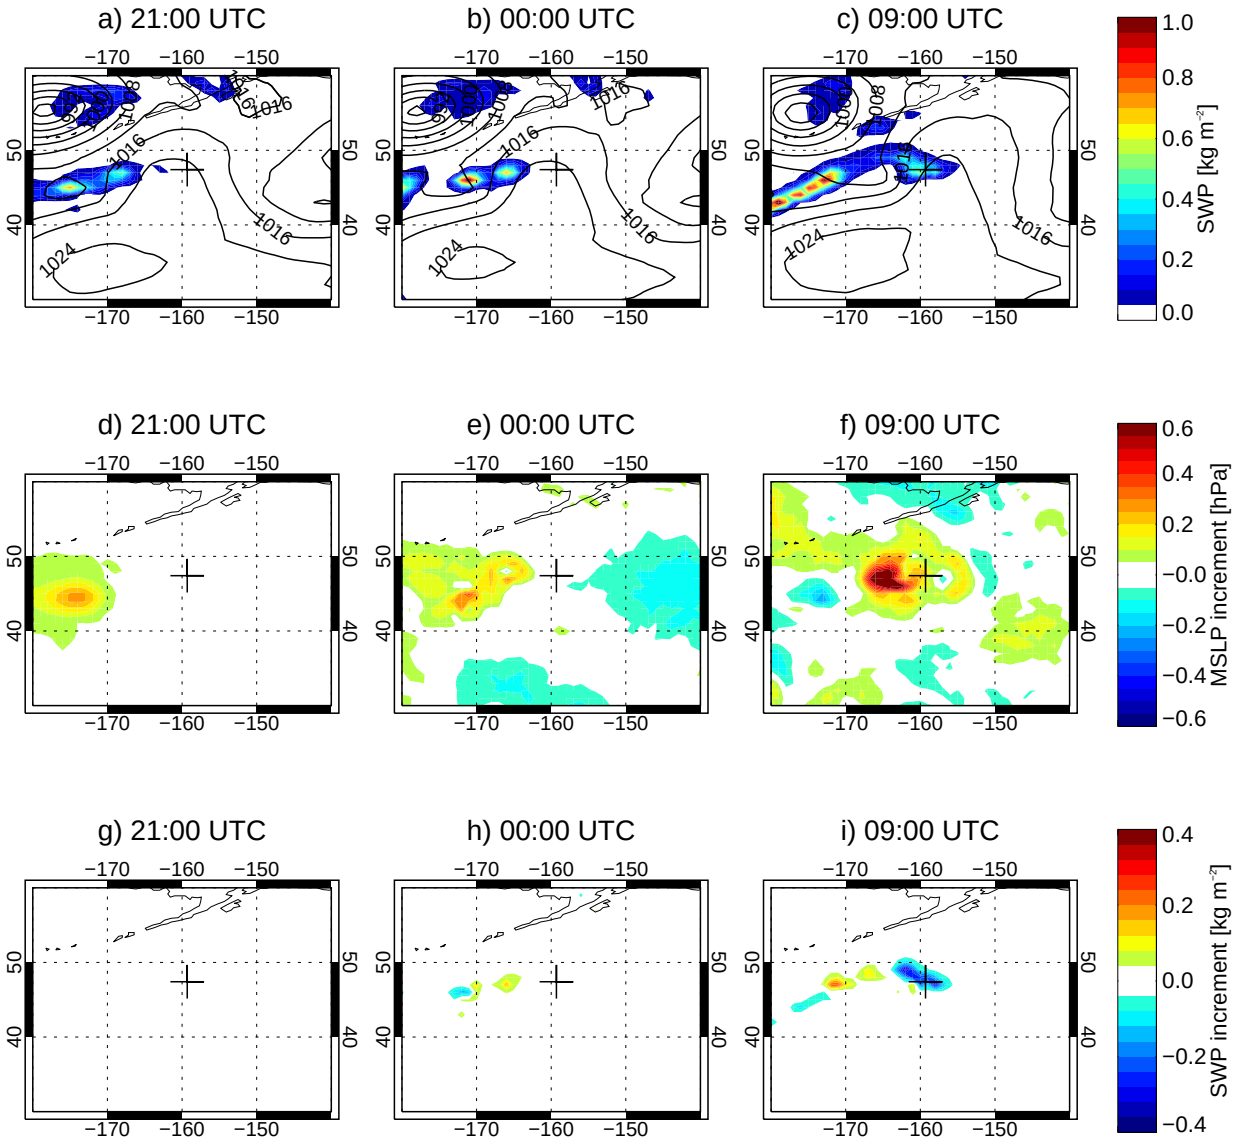

Figure 11: Case J through the assimilation window (21 UTC - 09 UTC) with the 'unrestricted' approach, i.e. with relaxed QC and smaller observation errors than usual: (a-c) First guess mean sea-level pressure (line contours, in hPa) and snow water path (filled contours); (d-f) Increments in mean sea-level pressure; (g-i) Increments in SWP. The cross shows the location of the observation, which was made at 08:00 UTC.

improves the fit to observations by 80%, principally by reducing SWP from 1.0 kg m^{-2} to 0.5 kg m^{-2}, which reduces the scattering index from 37 K to 18 K, close to the observed value of 15 K. Figures. 11c and 11i show that the SWP was reduced along a 700 km length of the front. As usual, the adjustments at the beginning of the time window were mainly dynamical and they are best illustrated by the 0.5 hPa increase in the pressure (Fig. 11d) at the centre of a small area of low pressure (Fig. 11a) that goes on to become a trough 12 h later (Fig. 11c). The intensity of the trough is reduced and hence the intensity of precipitation in the front that develops 12 h later is also reduced. This is further evidence of the ability of a 4D-Var system without a cloud-control variable to generate cloud and precipitation increments in the trajectory.

Case K is a frontal case with similarities to J, only it demonstrates the ability of the system to create cloud and precipitation, as well as to remove it. However, in the normal configuration the observation is rejected. With relaxed QC and smaller observation errors, the observation is used successfully in the 'unrestricted' test with a 95% improvement in the fit to observations (Fig. 9a). This is one of the cases where the full observing system analysis fails to fit the all-sky observation but the observation can be fitted when all other data are excluded.

Finally, cases E, L, M and N illustrate convection over land and ocean. There are a number of issues here: The observation in case N is rejected due to the check on excess scattering index over land. We could try relaxing this check in the future, given that the 'unrestricted' assimilation works quite well in this case. More worrying, in cases L and M the single observations don't make much impact when assimilated normally. In case L, giving extra weight to the observations makes the fit to observations worse, rather than better. These results are reminiscent of those of Bauer et al. (2010), who also tested all-sky assimilation in convective areas, with variable results. Convection is difficult to assimilate because in some cases it can be unpredictable over the 12 h timescales of the assimilation window. In other words it may not be possible to determine any change in the control variables that will successfully generate convection in a precise location later in the model trajectory. Even when that is possible, the increments generated in a low-resolution, simplified, linear model may not work when applied to a high-resolution nonlinear model. Increments which affect convection at one location in the minimisation may not affect it in the same place in the outer-loop. This can come from the changes in resolution, in the differences in the precise details of the forecast model, or simply from the invalidity of the tangent-linear approximation. Nevertheless, as demonstrated in Bauer et al. (2010) and in this work, incremental 4D-Var is usually able to adjust convection in order to fit a single all-sky observation when the observation errors are small.

5.2 Single-cycle experiments

When assimilating a single observation, there is no observational constraint except at the chosen time and location. Even if the analysis is improved at the observation location, 4D-Var could be degrading it elsewhere. It is important to demonstrate that the increments are meteorologically realistic globally and through the whole assimilation window. However, we cannot expect this in a single observation test case. Instead we can assimilate the observations globally from the five available microwave humidity sounders (i.e. MHS on Metop-A, Metop-B, NOAA-18, NOAA-19 and SSMIS on DMSP-F17) but in the absence of any other observations. As a reference, we will use the analysis from the full observing system, being by definition the best possible estimate of the atmospheric state. The experiment is performed for the 00 UTC assimilation window on 15 August 2013, the same date as the observations examined in previous sections.

On its own, the all-sky humidity sounder assimilation can generate mid-tropospheric increments that replicate a substantial portion of those from the full observing system. Figures 12 and 13 show the

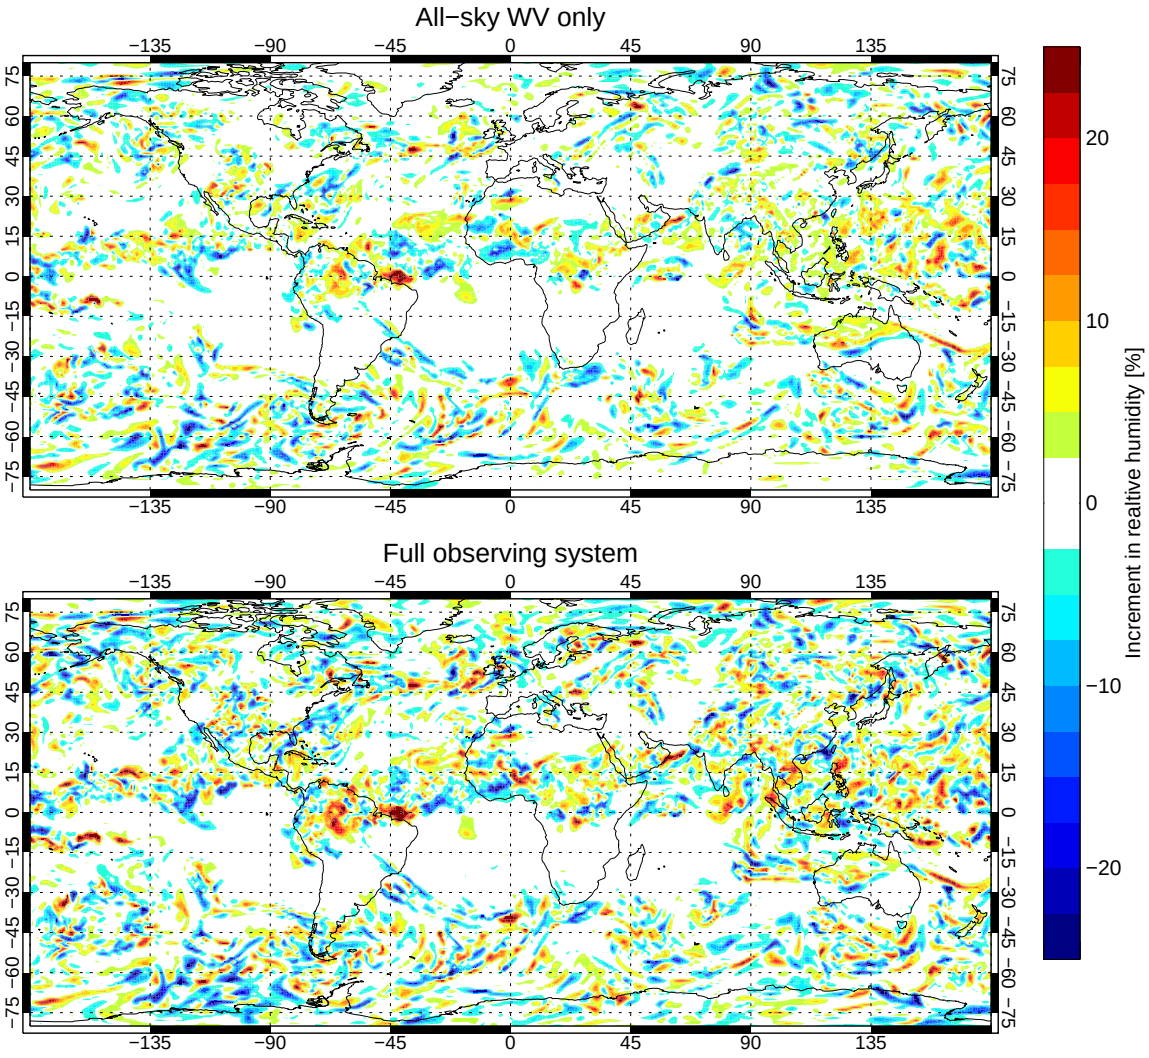

Figure 12: Increments in relative humidity on model level 95 (the closest level to 500 hPa) at 06 UTC on 15 August 2013: (a) assimilating only the all-sky microwave humidity observations; (b) assimilating the full observing system, including all-sky humidity channels. The first guess is identical in both cases and is created using the full observing system. Correlation between panels (a) and (b) is 0.72.

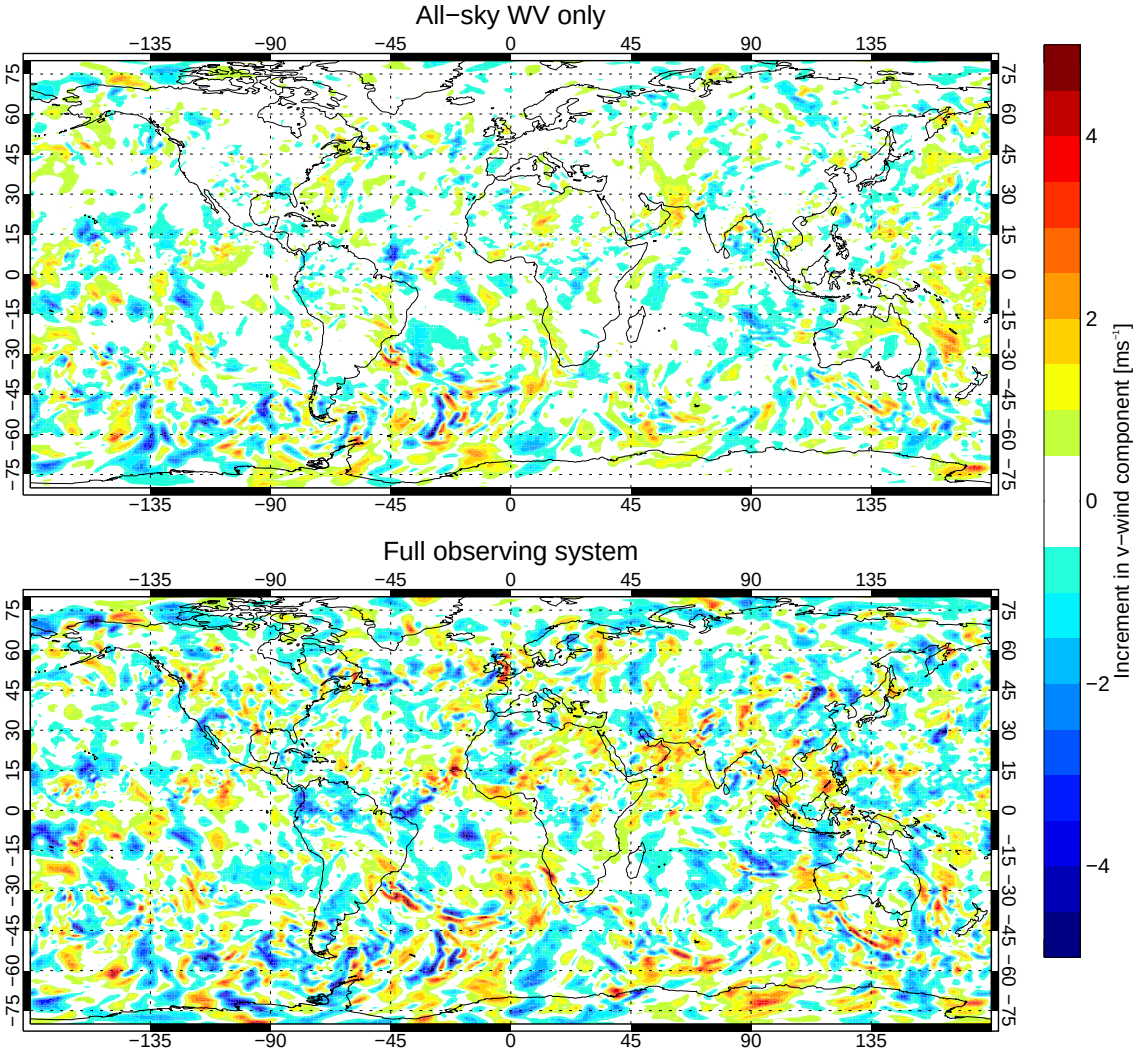

Figure 13: Increments in meridional wind component on model level 95 (the closest level to 500 hPa) at 06 UTC on 15 August 2013: (a) assimilating only the all-sky microwave humidity observations; (b) assimilating the full observing system, including all-sky humidity channels. The first guess is identical in both cases and is created using the full observing system. Correlation between panels (a) and (b) is 0.58. Increments in the zonal wind component show qualitatively similar patterns.

humidity and meridional wind increments 9 hours into the assimilation window at 500 hPa. The humidity increments resemble the normalised departures from Figure 4 in terms of their size and morphology, though they cannot be compared quantitatively because the satellite observations are valid at different times through the assimilation window and are sensitive to a broad layer of humidity, not just a single model level. These humidity increments are being generated specifically to fit the all-sky observations. The morphology of the wind increments is qualitatively consistent with those in the full system, on similar 100–400 km scales, and in similar areas of the globe: these wind increments will be associated with displacements, stretching and compression of the humidity field to make it fit the observations.

It is surprising how well the all-sky-only increments resemble those from the full observing system. The correlation between the two is 0.72 and 0.58 for the wind and humidity increments shown in Figs. 12 and 13. Figure 14 examines the correlations at the beginning of the time window for a range of levels and parameters. Correlations are generally a little smaller than those later in the window (not shown), suggesting that some of the correlation 9 hours into the window is being generated by the patterns of large-scale advection, which are similar in both experiments. However, correlations at the beginning of the window still reach 0.5 in the upper troposphere for the dynamical fields and 0.7 in relative humidity. The all-sky-only assimilation does not replicate the full system very well in the stratosphere and lower troposphere, but between about 250 hPa and 750 hPa, correlations are quite high. All-sky assimilation on its own can generate mid and upper-tropospheric wind, temperature and humidity increments that are reasonably consistent with those from the full observing system.

5.3 Cycled, reinitialised experiments

To gain a more representative sample, the single-cycle, single-observing system experiments were repeated twice-daily for short periods in summer and winter (15 to 31 August 2013 and 15 to 31 January 2014) and were used to generate a total of 68 forecasts. In this framework the first guess was always taken from the full-observing system experiment (the assimilation is 'reinitialised'). This avoids the problem of cycling the assimilation system with a restricted observing system, where the quality of the first guess would drop over a number of days. Cycling an assimilation system with a limited observing system is often referred to as a 'low-baseline' experiment (e.g. Kelly et al., 2008). However, low-baseline experiments are problematic because they do not represent a realistic scenario outside of special applications like early-20th century re-analysis (e.g. Compo et al., 2006). Further, in a 4D-Var context, and without a costly re-run of the ensemble data assimilation to correctly specify the background errors, the background errors in a low-baseline experiment will be representative of a much higher quality first guess, leading to a sub-optimal analysis. Hence we prefer to use 'reinitialisation' experiments with a high quality first guess. One way to think of these experiments is that they show the average change in forecast quality if, in just one 12 h assimilation cycle, we were to lose all observations except the all-sky microwave humidity channels.

The lower bound in this set of experiments is the quality of forecasts when no observations are assimilated for one cycle. The upper bound is the quality of forecasts where the full observing system is assimilated. We can create a metric that puts these points at 0% and 100% impact respectively. We can define $\mathrm{RMS}_{\mathrm{EXP}}$ as the root mean square of forecast error in experiment EXP, where the RMS error is computed across all forecasts and a chosen set of model points, for example midlatitude points at 500 hPa. As a verification reference we use the operational analyses for these experiments. Then the impact can be defined as:

$$I = 100 \times \frac{\mathrm{RMS}_{\mathrm{EXP}} - \mathrm{RMS}_{\mathrm{NO\_OBS}}}{\mathrm{RMS}_{\mathrm{FULL}} - \mathrm{RMS}_{\mathrm{NO\_OBS}}}. \quad (8)$$

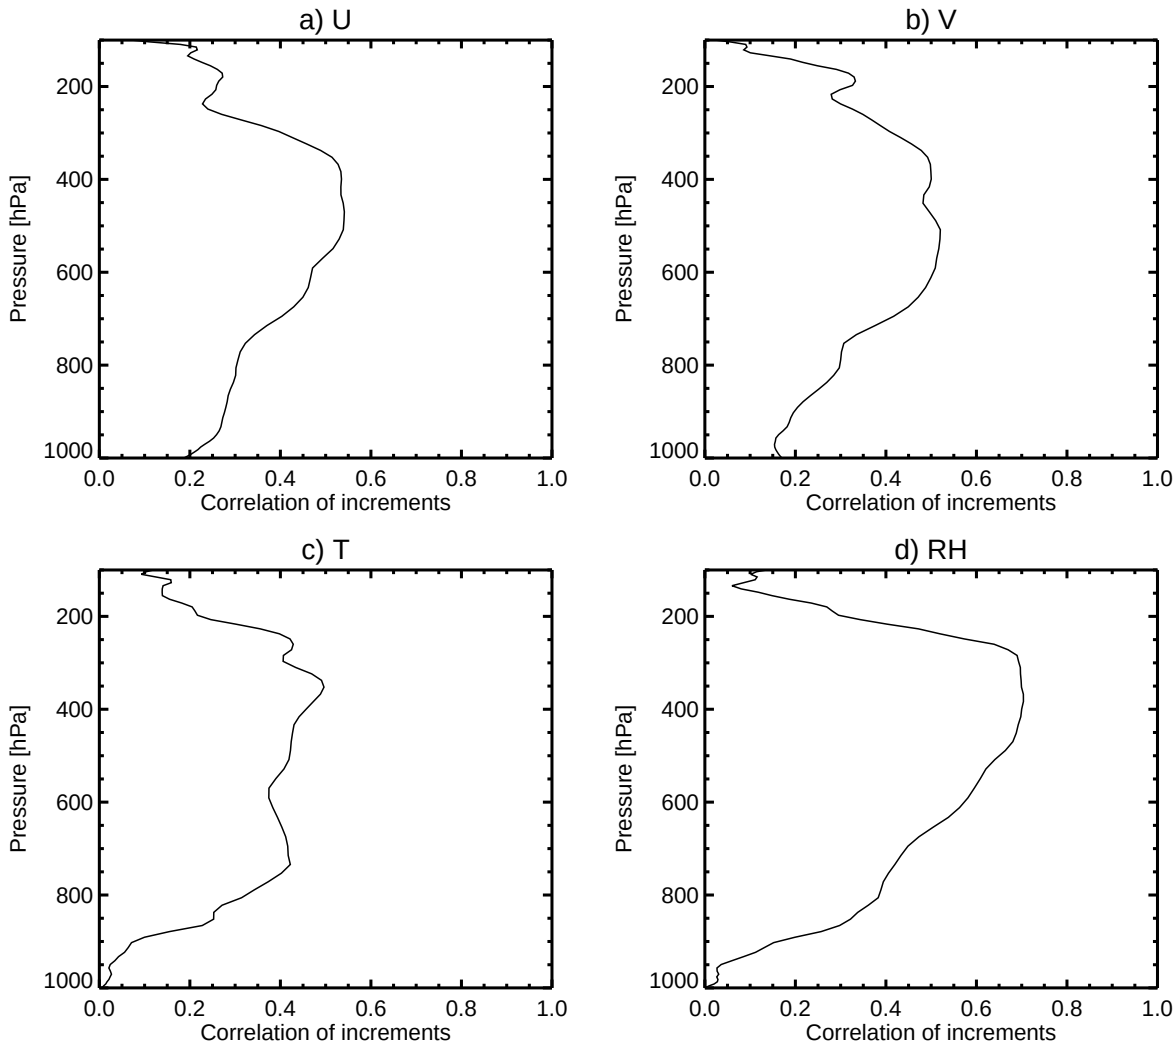

Figure 14: Global correlation of increments between a system assimilating only all-sky microwave humidity observations and a system with all observations assimilated. Correlations are shown at 21 UTC 14 August 2013, at the beginning of the assimilation window: a) zonal wind (U); b) meridional wind (V); c) temperature; d) relative humidity (RH).

Here, EXP is an experiment assimilating a single observing system, FULL is the full observing system and NO_OBS is no observations. If forecast quality were degraded compared to the no observation case, this 'impact' would be a negative number. This metric could be criticised for treating forecast errors as if they combine linearly when random errors usually sum in quadrature. However, the differences in RMS error between experiments (outside of the low-baseline scenario) is usually much smaller than the absolute RMS error, so an impact factor computed from squared errors is much the same as the one defined above. A more important feature in practice is that forecast impacts (i.e. changes in forecast error) combine neither linearly or in quadrature because impacts are often correlated. For example, different observing systems often bring similar information on the large-scale structure of the atmosphere. With these caveats, Fig. 15 shows the impact of assimilating only microwave humidity channels, through the all-sky approach.

Humidity errors are reduced in the mid and upper troposphere by 40% to 80% of the error reduction generated by the full observing system. Temperature errors are reduced by up to 40% in the SH, though outside of the mid and upper troposphere the impact is negative in places. In the Antarctic, these areas are below the true surface which can be around 600 hPa and hence this is just an artefact of using fixed pressure levels in this comparison. Elsewhere, a possible explanation is that microwave humidity observations cannot constrain the low-level temperature field and in the absence of other observations, so the assimilation system is free to make spurious changes to the temperatures. However, vector wind errors are improved throughout the troposphere, by 50% in the SH midlatitudes where the 4D-Var tracing effect is expected to be largest. Impact is smaller in the NH where snow-covered land surfaces and high orography reduce the coverage of the satellite data. However, in contrast to the Antarctic, there is some impact at very high northern latitudes because observations can be assimilated over the Arctic ocean.

Figure 16 compares the impact of clear-sky and all-sky assimilation of microwave humidity channels on forecasts of relative humidity and as a reference, also includes the impact of clear-sky AMSU-A and ATMS temperature-sounding observations, which are perhaps the most important single observing system for weather forecasting. In the midlatitudes, in the short-range (T+12) clear-sky and all-sky assimilation have similar impact on humidity forecasts and substantially more impact than AMSU-A in the mid and upper troposphere (70% versus 30% in the SH). This is a clear sign of the direct sensitivity of 183 GHz channels to mid and upper-tropospheric relative humidity. However, in the three-day forecasts (T+72) the humidity channels have a smaller impact on relative humidity than do the microwave temperature channels on AMSU-A and ATMS (40% versus 60% in the SH). At longer ranges, improvements in humidity forecasts come from better prediction of the large-scale dynamics and ultimately from the direct temperature sensitivity of AMSU-A and ATMS.

Figure 17 shows impacts on vector wind errors. The all-sky approach gives a stronger impact on wind scores than the clear-sky approach. The SH is where the all-sky approach brings the greatest increase in observational coverage (Fig 5). At 500 hPa in the SH, the all-sky assimilation of humidity channels brings almost the same amount of wind information as AMSU-A and ATMS, with an impact of about 45%. The wind impact is more uniform across the vertical extent of the troposphere than is the humidity impact. As seen in the single-observation tests, broad changes in the dynamical structures of the atmosphere are required to improve humidity and cloud forecasts in the mid and upper-troposphere. In the day-3 wind forecasts in the SH, the all-sky humidity observations reach a 50% impact on forecasts compared to 40% from clear-sky humidity and 65% from temperature-sounding observations. In the SH, the dynamical impact of humidity sounding observations is beginning to approach that of the temperature sounding observations.

An important question is whether the benefits of all-sky assimilation come through constraining the water vapour fields in the presence of clouds or through the cloud fields themselves. A crude way to test

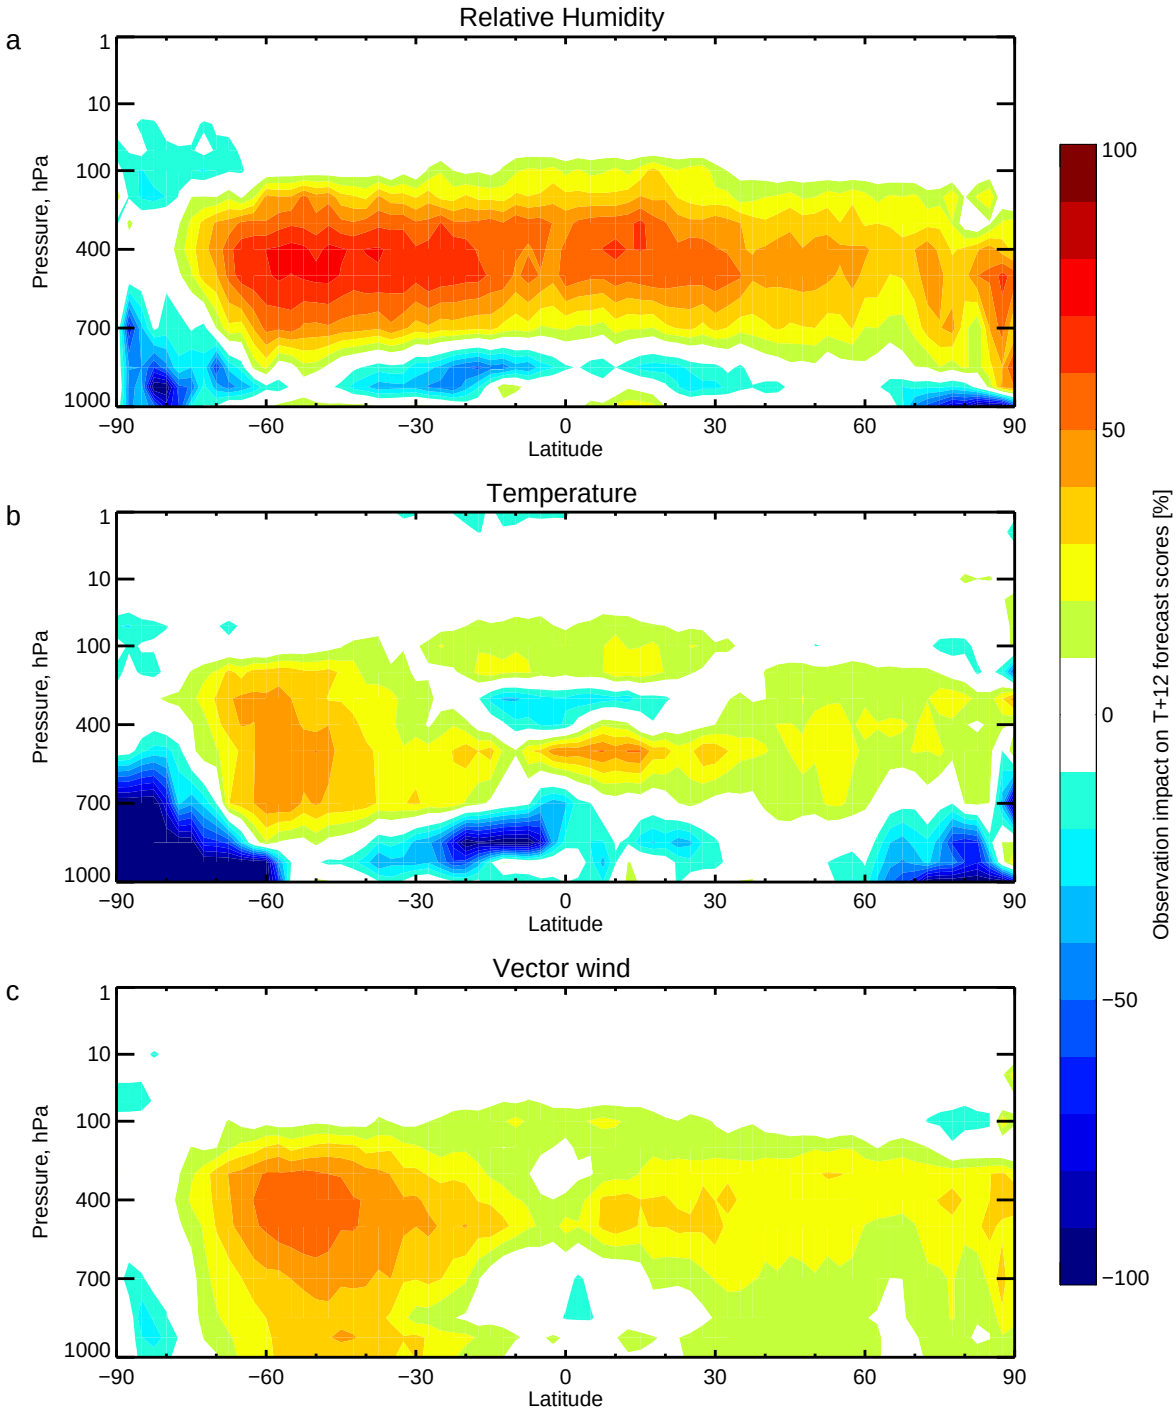

Figure 15: Impact on forecast errors at T+12 when all-sky microwave humidity sounders are assimilated in a framework where all analyses start from a high quality first guess. Here, 0% corresponds to no observations being assimilated and 100% corresponds to the full observing system: a) relative humidity; b) temperature; c) wind.

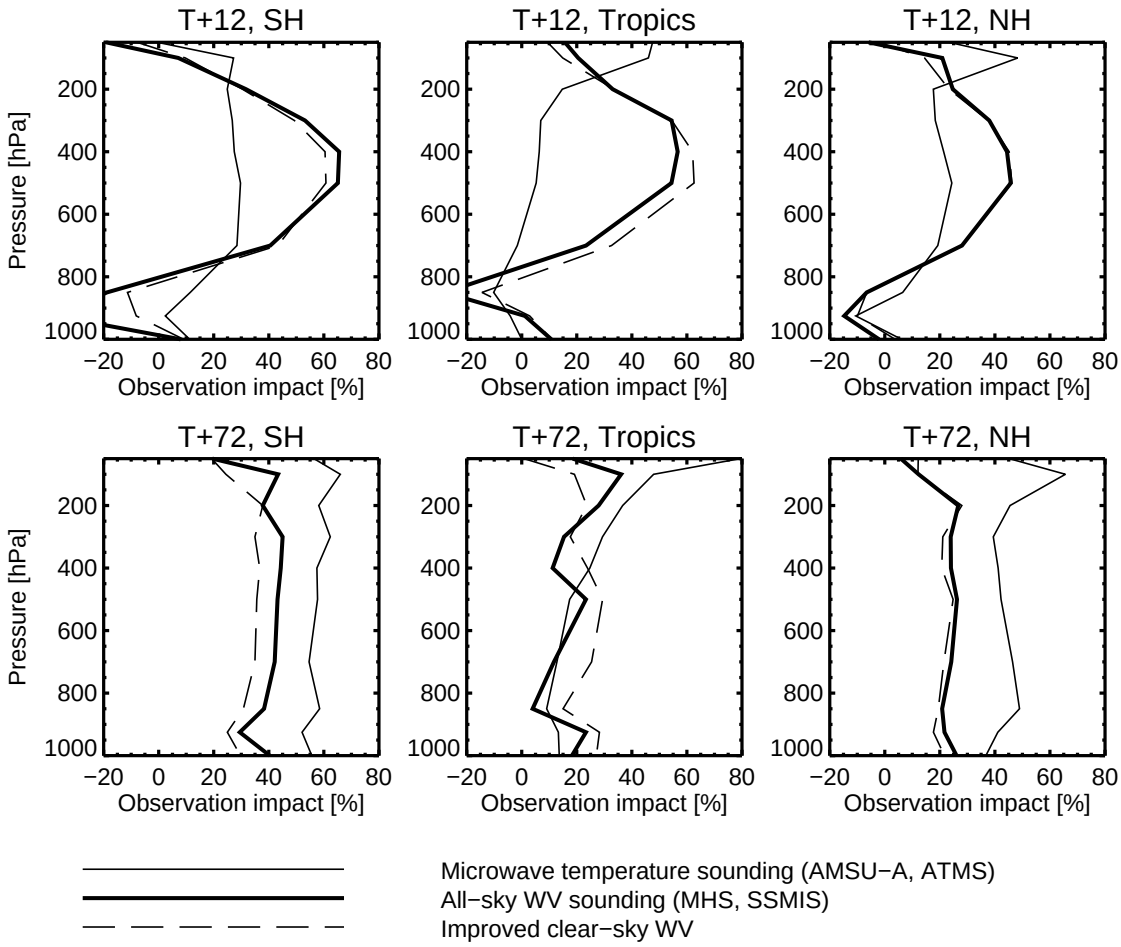

Figure 16: Impact on relative humidity forecast errors of single observing systems in a framework where all analyses start from a high quality first guess. Here, 0% corresponds to no observations being assimilated and 100% corresponds to the full observing system.

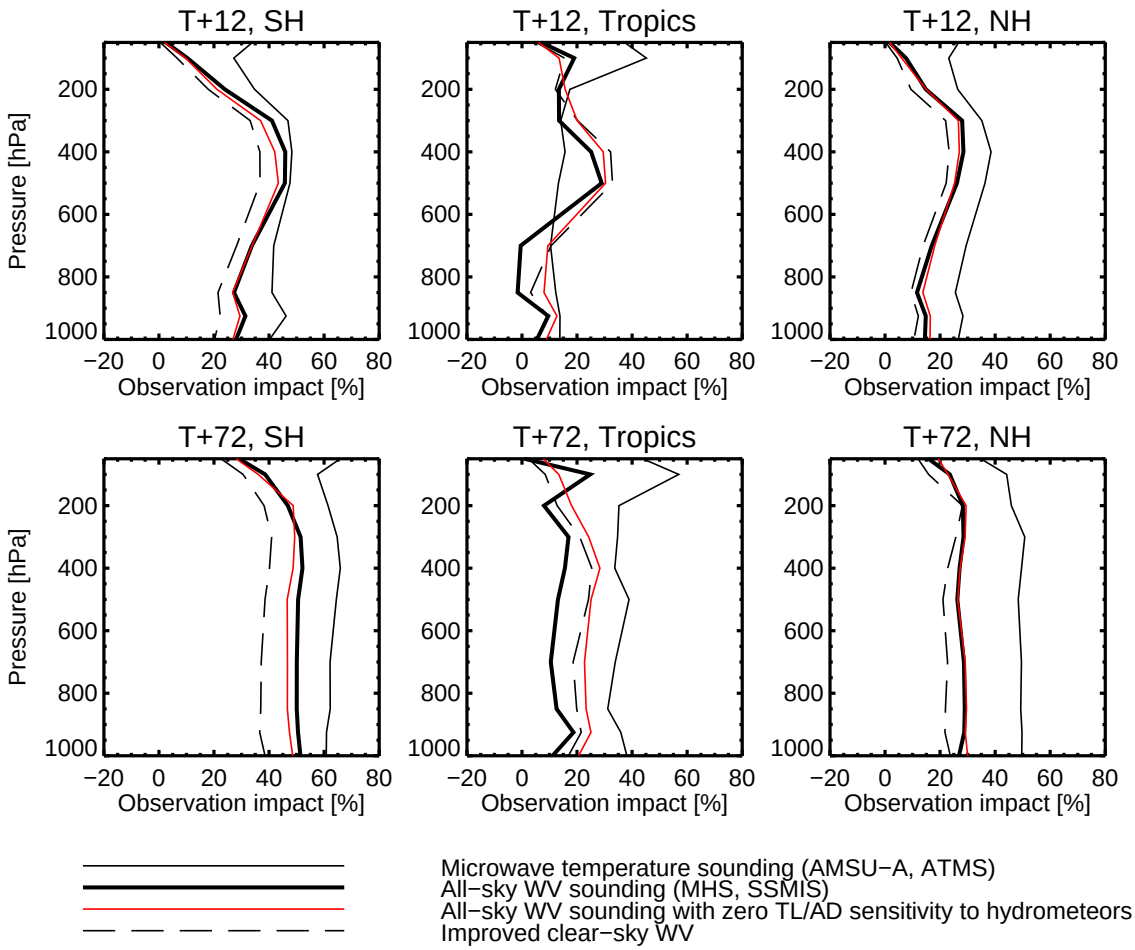

Figure 17: Impact on vector wind forecast errors; details otherwise for Fig. 16

this is to turn off the adjoint and tangent linear sensitivity to cloud and precipitation in the observation operator used in 4D-Var. Having done this, the assimilation system can only directly improve the fit to all-sky observations by changing the moisture fields at the observation location. The assimilation system is prevented from directly adjusting clouds and precipitation to fit the observations, though this may be a natural secondary result of improving the humidity. The impact of this is included in the wind figure (Fig. 17) but not in the relative humidity one (Fig. 16) where it is hard to distinguish it from other lines. The wind figure shows that turning off the cloud and precipitation sensitivity in the minimisation has relatively little impact in the midlatitudes. It looks like the majority of impact in going from clear-sky to all-sky assimilation comes from a better constraint on relative humidity in the presence of cloud, rather than directly through the sensitivity to cloud and precipitation. Nevertheless, the cloud and precipitation sensitivity is clearly beneficial, bringing up to an additional 4% forecast skill.

So far we have not mentioned the tropics, where in Figs. 16 and 17 the microwave humidity assimilation has a smaller impact in the all-sky approach than in the clear-sky approach. Further, turning off the tangent-linear and adjoint sensitivity to cloud and precipitation actually improves the forecast quality, bringing it closer to the clear-sky approach. In the context of the full observing system the picture is quite the opposite (see the results in the next section). We can speculate that difficulties in the tropics come from trying to assimilate observations of deep convection. As shown in the single observation tests, this is a much harder problem than assimilating mid-latitude cloud and precipitation. To fit the observed deep-convection, the analysis may be forced to make adjustments to humidity and dynamical fields that may not be consistent with the true atmosphere. The full observing system probably helps better constrain the humidity and dynamical fields in the tropics to prevent these spurious adjustments. In the tropics, it is only in the context of the full observing system that the all-sky approach can have a positive impact.

6 Impact in the full observing system

6.1 Impact on dynamical fields

Here we examine the impact of SSMIS and MHS humidity sounding channels in the context of the full observing system. All-sky or clear-sky microwave humidity observations are added to a control which contains the full observing system minus humidity sounding observations from these instruments. Figures 18 and 19 show the change in forecast error in vector wind. The impact of all-sky is bigger than clear-sky assimilation and there is a maximum impact around forecast day 2 (T+48). Here, the impact is significant and greater than 1% in most of the troposphere and stratosphere. The impact reaches 4% in the SH upper-troposphere, where we expect the greatest benefits from wind-tracing, but it is greater than 2% in most of the NH extratropics, too. Impacts are significant at least out to day 4 (T+96) but they are starting to tail off at day 6 (T+144).

The only area where microwave humidity sounders apparently increase wind forecast errors is in the tropics at around 400 hPa, early in the forecast range. However, tropical wind errors are significantly reduced in the medium range (Fig. 18) and first-guess fits to other wind observations are improved at all levels in the tropics (see later) so this is a curious feature. The feature is more obvious in the clear-sky approach and at T+12 where it extends into the mid-latitudes (Fig. 19). By construction, the change in T+12 own-analysis scores is equal to the change in the RMS of the analysis increments. Assimilation of microwave humidity observations makes analysis increments slightly larger in the tropical mid-troposphere. Looking at maps of the impact in either framework (not shown) the effect is largest over parts of central Africa

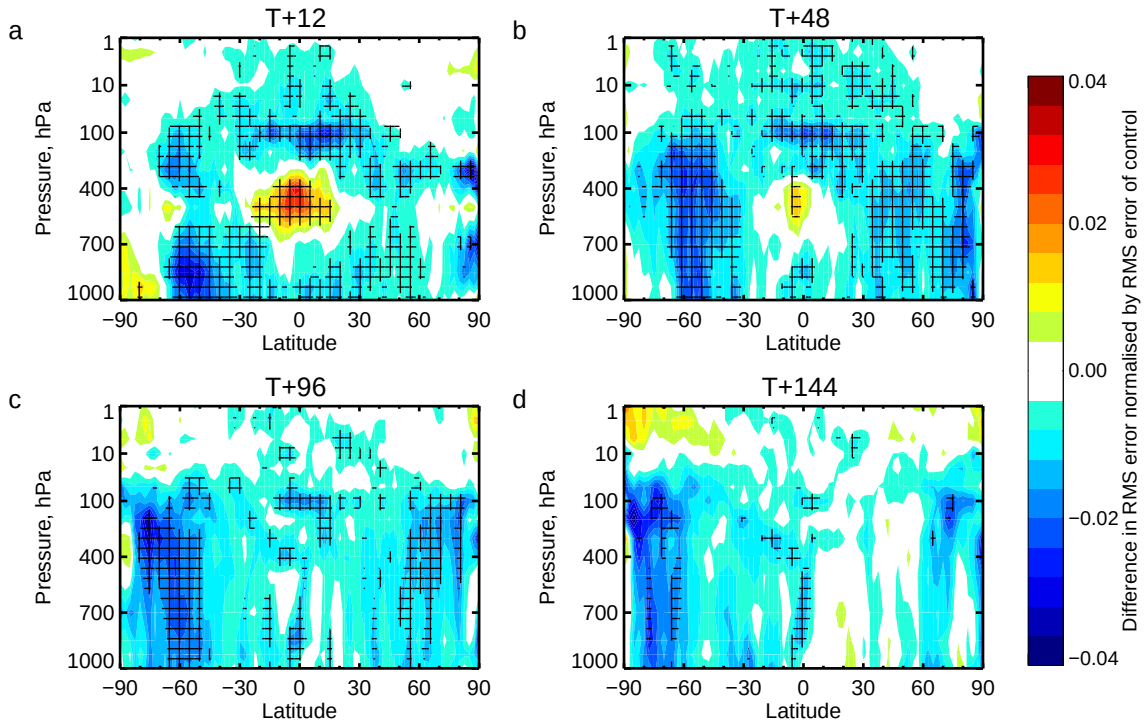

Figure 18: Normalised change in RMS of vector wind forecast errors when all-sky microwave humidity assimilation (MHS and SSMIS) is added to the otherwise full observing system. Reductions in errors are desirable and are shown in blue. Cross-hatching indicates statistical significance at the 95% confidence level. Results are based on a total of six months experimentation, winter and summer, and are based on 322 to 360 forecasts depending on the forecast range. Verification is against own analysis.

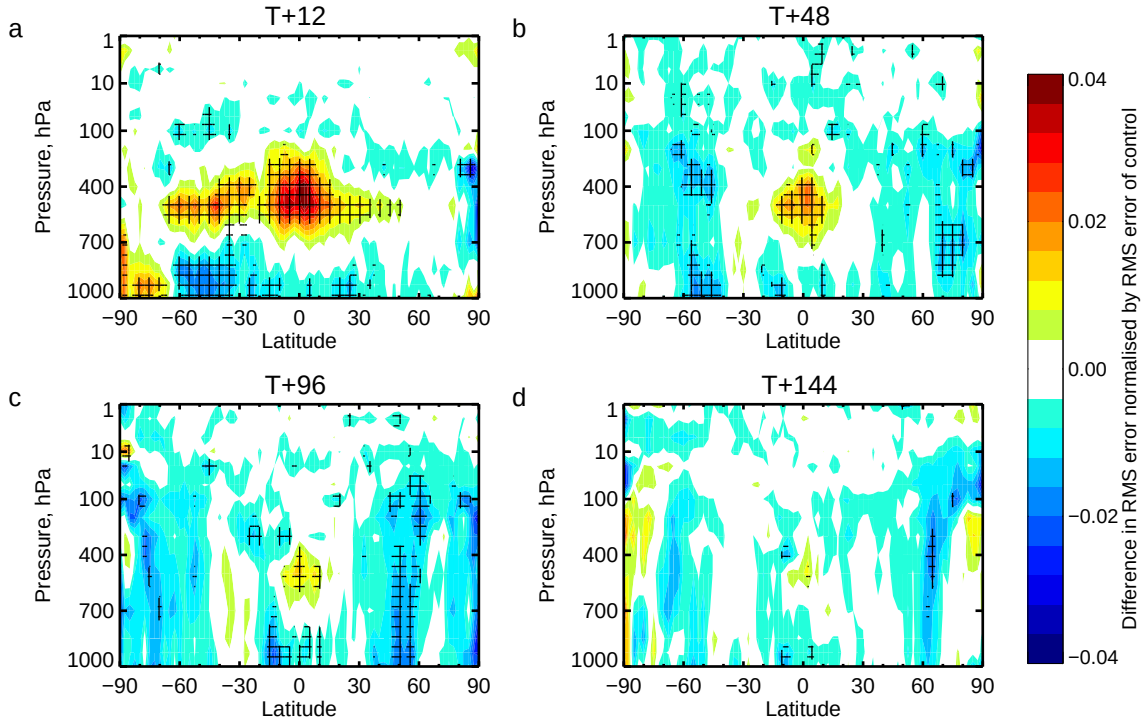

Figure 19: As Fig. 18 but for improved clear-sky humidity sounding versus no microwave humidity.

and the equatorial Atlantic. It is a feature most clearly of the ITCZ or its immediate vicinity. There are few signs of mean changes in forecast winds fields aside from a reduction in the equatorial easterly flow of around 0.1 ms^{-1} and even this is restricted very much to the equator and to the 400hPa and 500 hPa levels (see later). Previous sections have shown the difficulty of inferring wind and moisture increments in the vicinity of deep convection. With the assimilation of the microwave humidity sounders, we may be making additional increments to fit transient convective features. In this context, it is a good sign that the issue is reduced in the all-sky approach compared to the clear-sky approach.

Figure 20 examines the changes in RMS errors in dynamical fields as a function of forecast time, giving a summary for the main regions of the globe. In the extratropics, microwave humidity assimilation improves wind, temperature and geopotential scores by 1% – 2% with statistical significance out to at least day 6 in both hemispheres. All-sky assimilation has roughly double the impact of clear-sky assimilation in the shorter forecast range, and the improvement over clear-sky is significant out to around day 5 in the SH and day 3 in the NH (the hemispheric differences and the exact day that the significance is lost should not be taken too literally because in previous runs with no differences apart from an older compiler version the situation was opposite: day 3 in the SH and day 6 in the NH). For rough comparability to other studies, these figures can be converted to a gain in forecast lead time for the same level of skill. Taking the SH scores at day 5, there is around ¾ hour improvement in wind and geopotential forecast skill for the clear sky approach and about 1½ hours for the all-sky approach.

In the tropics, all-sky assimilation has a particularly strong impact at 100 hPa, reducing errors by 2% in wind and 3% in temperature. However, the strong tropical improvement at 500 hPa in geopotential and smaller degradation in temperature at 850 hPa are not important. These features appear neither in fits to temperature-sensitive observations nor when looking at changes in the standard deviation of forecast

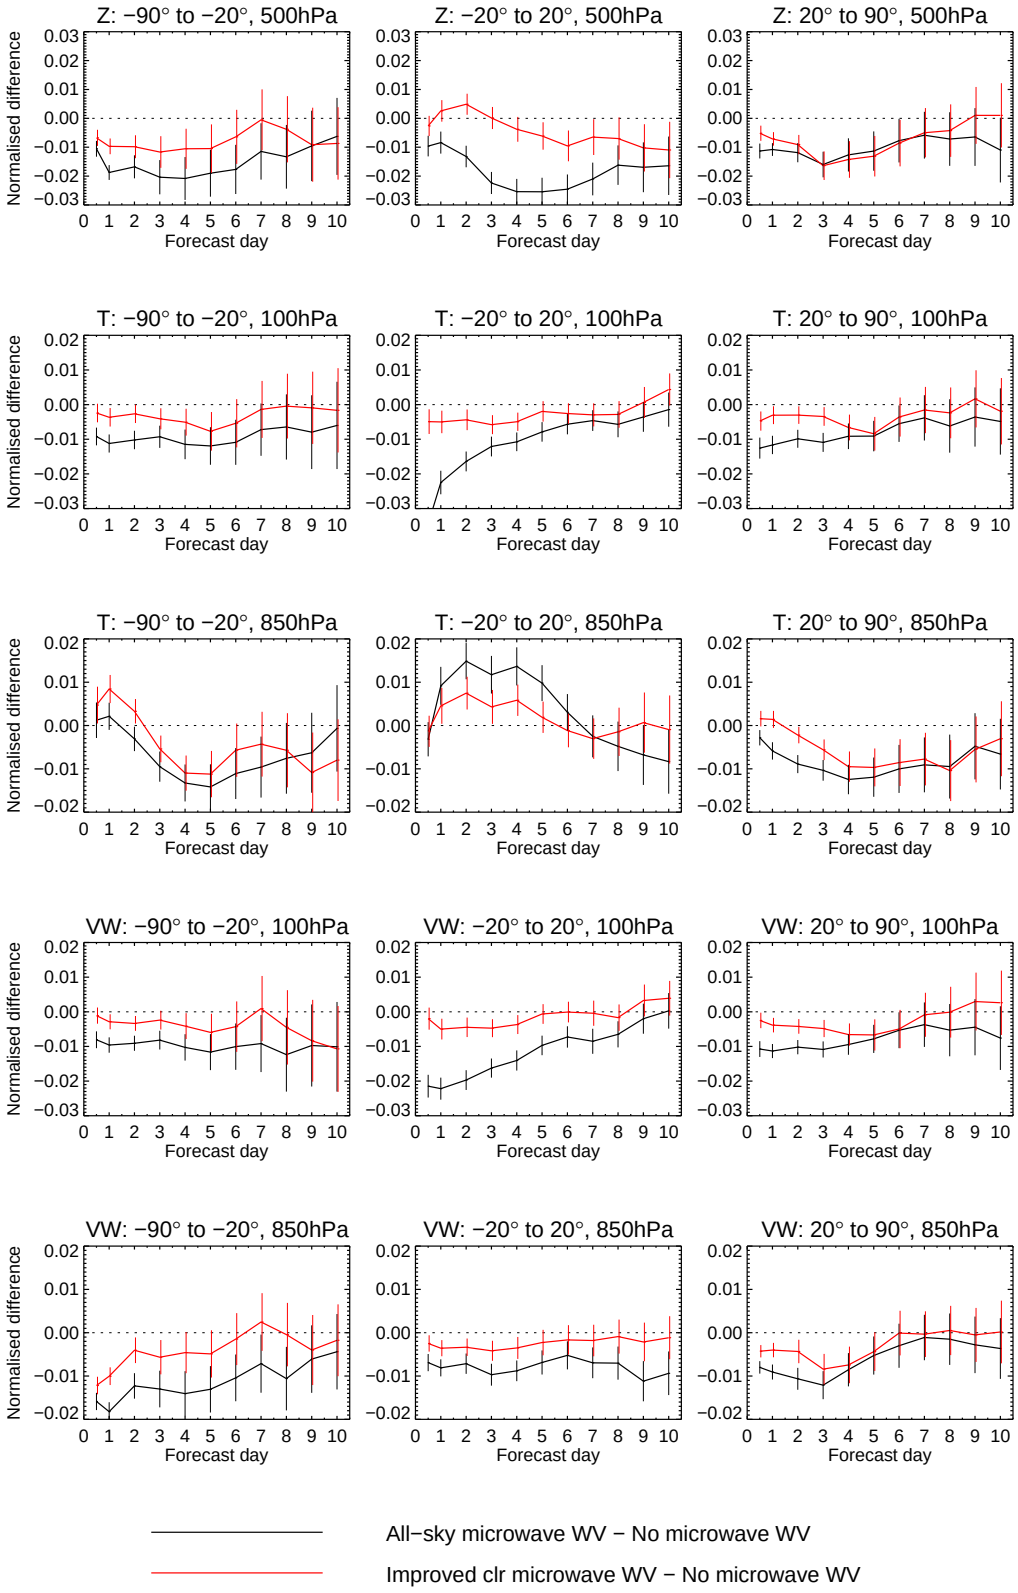

Figure 20: Changes in RMS error as a function of forecast time: geopotential at 500 hPa (top row), temperature at 100 hPa and 850 hPa (middle rows), and vector wind at 100 hPa and 850 hPa (bottom rows).

error instead of the RMS. These features are caused by small changes in the mean analysed temperature in the tropics. In the tropics the absence of fast growing random errors in the temperature fields makes changes in RMS temperature and geopotential errors tricky to interpret.

First-guess fits to observations give an alternative measure of the quality of T+12 forecasts. Figure 21 shows that fits to wind observations are improved by the assimilation of microwave humidity data, whether in clear-sky or all-sky. The additional benefit of all-sky assimilation is seen in the tropics in the improved fit to conventional wind observations around the tropopause and into the stratosphere. The impact peaks at 150 hPa, where the FG fit is 1.5% better. The SATOB winds (atmospheric motion vectors) have negligible coverage at 150 hPa so they do not contradict the conventional observations. In the midlatitudes, the SATOB winds show all-sky bringing twice the impact of clear-sky assimilation, whereas the conventional winds show all-sky and clear-sky having similar impact. Certainly in the maps of forecast impact (not shown) the additional benefit of all-sky is greater in ocean areas: these are better covered by the SATOB winds than the conventional data.

Fits to temperature-sensitive observations show a similar picture (Fig. 22). In the tropics, all-sky has a pronounced impact around and above the tropopause. The greatest impact is seen in channel 9 of AMSU-A, which has a weighting function peaking at around 80 hPa, in the radiosonde temperatures at 100 hPa to 200 hPa, and in the radio-occultation (GPSRO) bending angles at 16 km to 23 km (40 hPa to 100 hPa). In the midlatitudes, all-sky assimilation is significantly better than clear-sky assimilation, though the most pronounced impact is in AMSU-A channel 6, which has a weighting function peaking at roughly 8 km or 300 hPa. Given the greater impact on satellite fits compared to conventional observations, this again suggests the additional benefit of all-sky assimilation is greatest in the remote ocean areas that are only observed by the satellites. This is consistent with the spatial coverage of the microwave observations, which is severely restricted over cold or high land surfaces.

From the forecast scores and observation fits, we can see that microwave humidity observations help to infer large-scale, dynamical information that brings substantial benefit to medium-range forecast scores. Previous sections have illustrated how this can be generated through the 4D-Var tracer effect. Going to all-sky assimilation roughly doubles the observational coverage at higher latitudes and it leads to a rough doubling of the impact on forecasts, at least in the shorter ranges. By comparing mean, RMS and standard deviations of the forecast scores (not shown) we can eliminate any possibility that forecast improvements in the midlatitudes have been generated by systematic changes. The reduction in forecast errors is a genuine improvement in day-to-day weather patterns.

In the tropics, the all-sky approach does not greatly increase the number of observations but it leads to a disproportionate impact on the quality of the dynamical forecasts of the tropical tropopause and lower stratosphere. This impact may come through 4D-Var tracing if the upper-level winds in the vicinity of deep convective systems are particularly important. However, as discussed in section 3, there are remaining systematic biases between model and observations (up to half the size of the observation error) in and around the convective areas. Figure 23 shows the normalised change in forecast errors in temperature and vector wind at 100 hPa. The impact of microwave humidity assimilation is broadly spread between 30°N and 30°S, though the greatest impacts are in the Indian Ocean and Western Pacific, where RMS errors are locally improved by over 10%. Looking at changes in zonal mean fields, Fig. 24 shows that all-sky humidity assimilation cools the tropical upper troposphere by around 0.01 K (panel a) and slightly retards the Hadley circulation (panel e) and the equatorial easterly flow (panel c). This systematic impact is relatively small. A question is whether the 10% impacts on RMS errors at 100 hPa come from systematic changes like this, or from improved forecasts of day-to-day weather variability. However, maps of the difference in error standard deviations (not shown) are very similar to those for RMS, so day-to-day variability is definitely being improved. However, convection moves around and

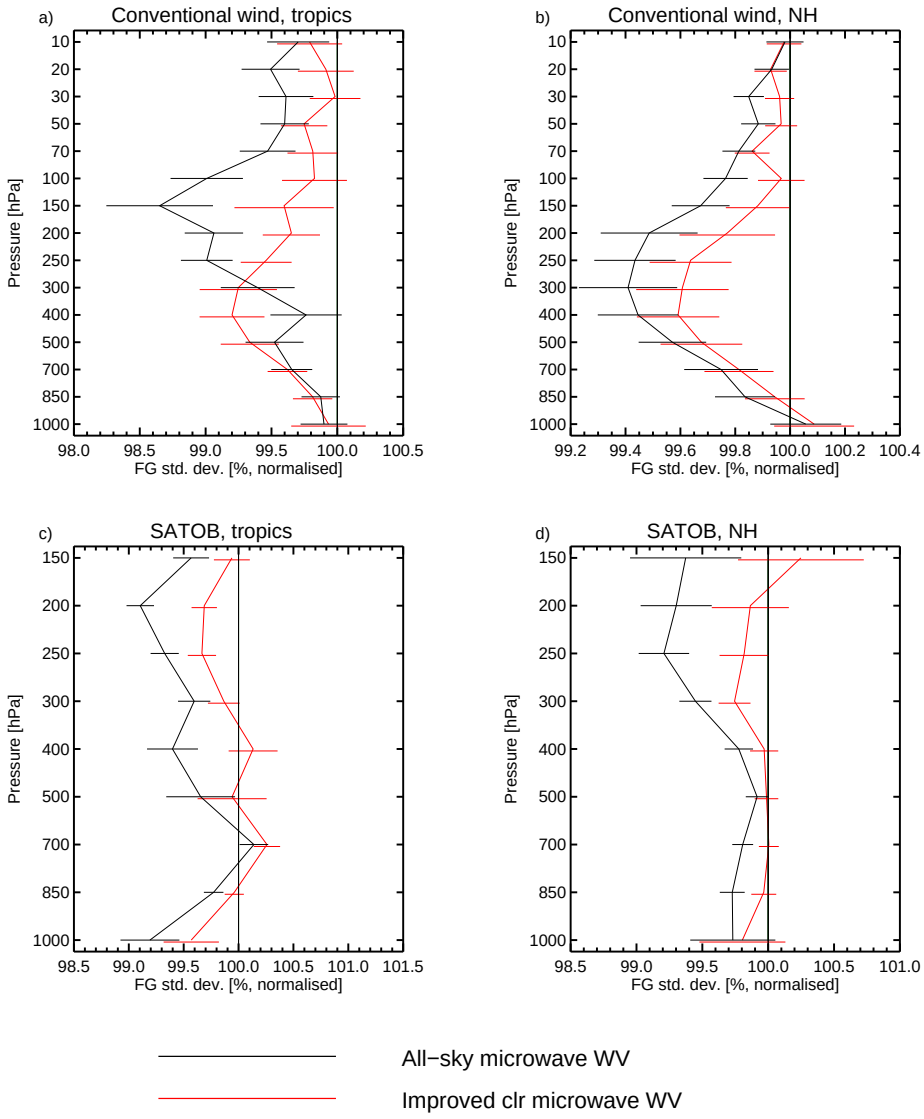

Figure 21: Standard deviation of FG departures for wind observations as a percentage relative to the control. Error bars indicate the 95% confidence range. Control is without microwave humidity sounders but otherwise contains the full observing system. Conventional observations include u and v component winds from profilers, radiosondes and aircraft. SATOB winds are atmospheric motion vectors derived from satellite images. Only tropical and NH fits are shown; results in the SH are similar.

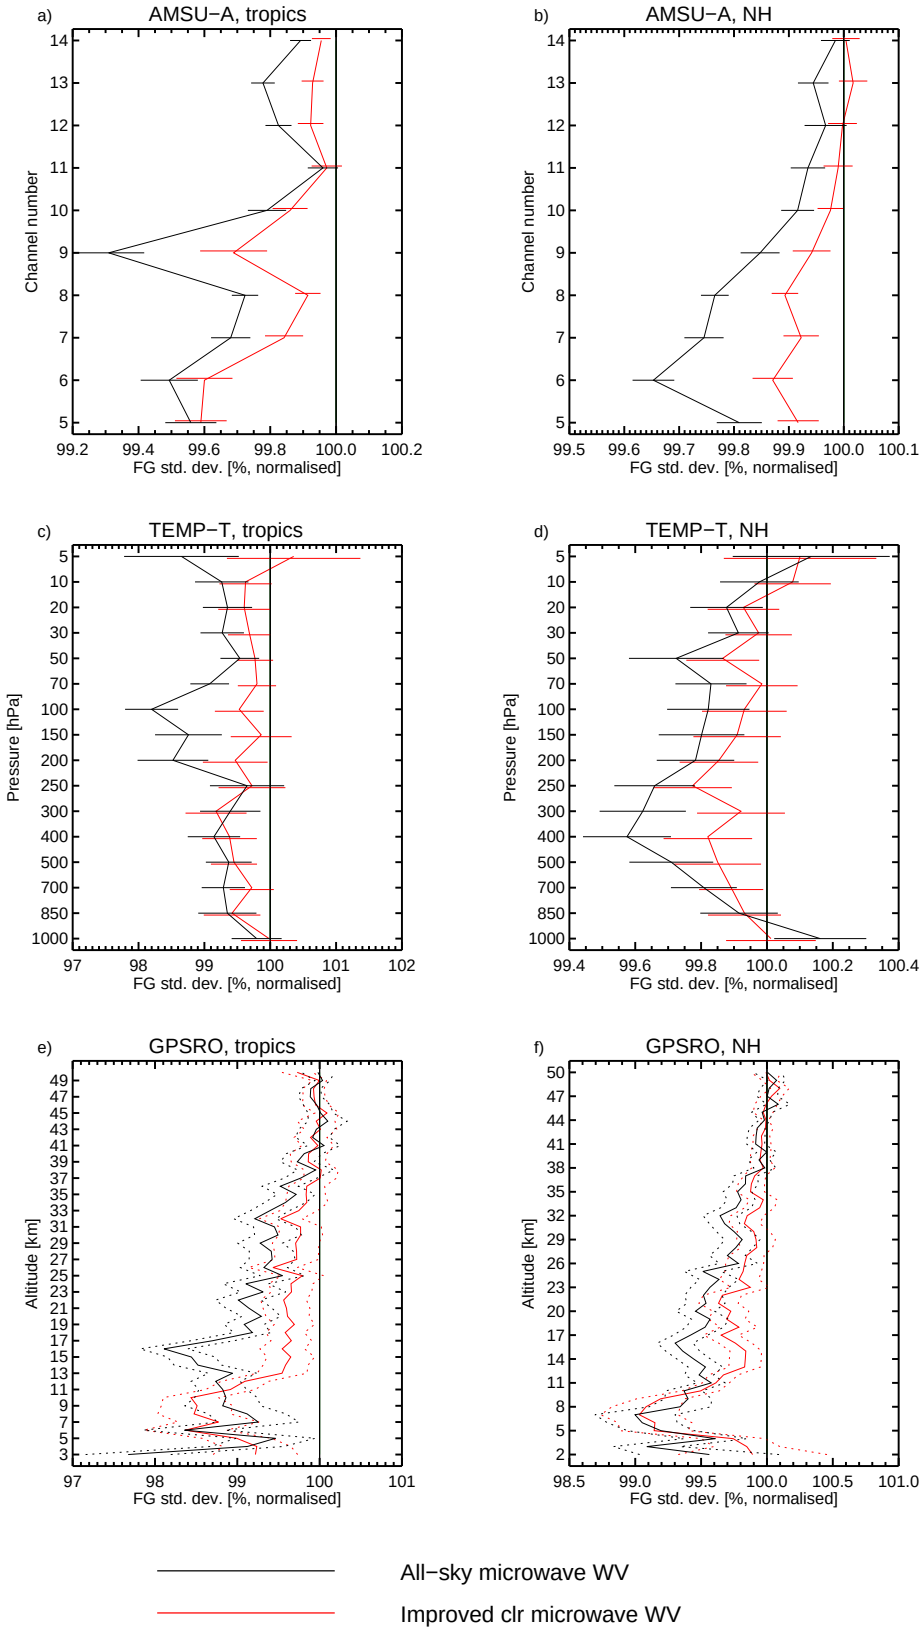

Figure 22: Standard deviation of FG departures for temperature-sensitive observations (TEMP-T is radiosonde). Other details as for Fig. 21

if we were to follow the convection, it is quite possible that the changes in the model might look more systematic.

The microwave humidity assimilation does cause two systematic changes in dynamical fields, both in the lower troposphere and visible in the temperature fields but not the wind fields. Figure 25 shows the change in the temperature field at 850 hPa. There is a systematic warming of around 0.2 K over the sea-ice areas that is slightly smaller in the all-sky approach, compared to the clear-sky approach. The issue was already present after the extension of clear-sky assimilation to sea-ice areas but it was tolerated because of improved fits to Antarctic radiosonde humidity and temperature profiles (mostly on the edge of the sea-ice Di Tomaso et al., 2013). The channels assimilated over sea-ice, e.g. middle and upper 183 GHz channels, show no systematic departures in these areas (Fig. 6a to d) but it may be that the whole system (bias-corrected observations and model) has moved to a new, stable temperature.

The other systematic change is a cooling of the subtropics peaking at around 0.4 K in regions of maritime stratocumulus and trade cumulus, in the western parts of the subtropical oceans. This cooling goes with a moistening in similar areas (next section) and it is similar to the impact of microwave imager assimilation, which come from model biases in cloud liquid water (including a poor diurnal cycle) in boundary-layer maritime cloud (Kazumori et al., 2014). However, the presence of cooling even in the clear-sky assimilation suggests that either the observations contained undetected cloud or that the lack of cloud liquid water is associated with too-low relative humidities in the boundary layer.

6.2 Impact on humidity and precipitation

Humidity forecast scores are shown in Fig. 26. In both clear-sky and all-sky assimilation of humidity sounders, there are increases in RMS forecast error, particularly at short forecast ranges. As explained previously, T+12 own-analysis scores are equivalent to a measurement of the change in the RMS of the humidity increments. The biggest early-range impacts are over the sea-ice regions and they do not seem to persist into the medium-range forecast. These effects are linked to the warming over sea-ice examined in the previous section. In mean terms, relative humidity is increased by around 1% in the mid and lower troposphere and decreased by around 1% in the upper troposphere in both clear-sky and all-sky assimilation (Fig. 27). The regions of moistening are clearly bounded within the sea-ice regions (maps of mean changes, not shown) and go with the warming described in the last section. As mentioned, this feature is tolerated because it produces improved fits to observations.

Outside of the sea-ice regions, the relative humidity forecast impact is dominated by increased RMS errors persisting many days into the forecasts. But in the all-sky assimilation we start to see RMS errors being reduced significantly, particularly later in the forecast range. The reduction in all-sky RMS error is associated with a general moistening of around 0.5% relative humidity in the mid and lower troposphere (Fig. 27) that makes the analysis and forecast more consistent and hence makes the own-analysis RMS errors smaller. Clear-sky observations generally dry the troposphere whereas the all-sky impact is more balanced. Over the same levels, FG biases towards humidity radiosondes are also reduced by going to all-sky assimilation (not shown), so this may be a genuine improvement. This pattern is consistent with one hypothesised drawback of clear-sky assimilation: the data selection will be skewed towards dry areas. This selection bias is worth further study when we apply the all-sky approach to the infrared, where the current ‘hole hunting’ approach is likely to cause a greater dry bias.

The standard deviations of FG departures for humidity observations show more uniformly favourable results than the forecast scores and help put the apparent degradations into perspective (Fig. 28). Both all-sky and clear-sky humidity assimilation improve fits to radiosonde humidities by 1%–2% in the upper

(a) Temperature

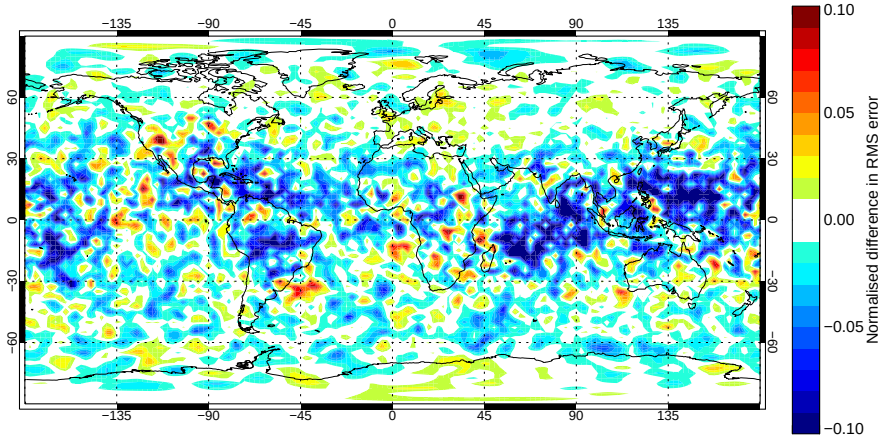

(b) Vector wind

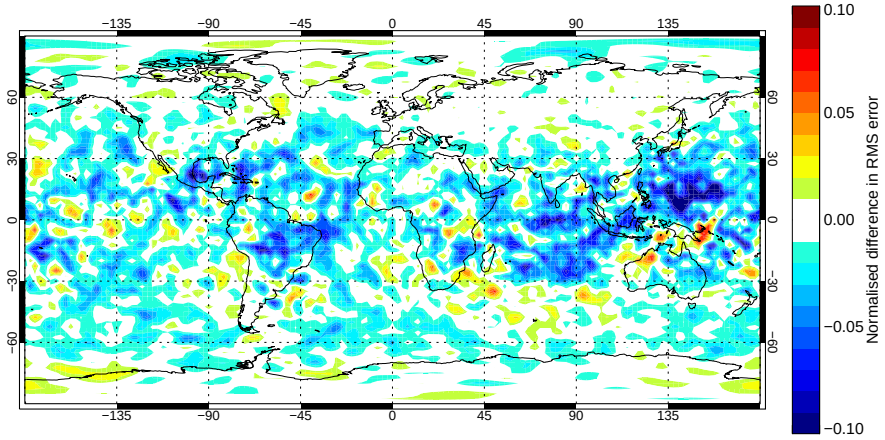

Figure 23: Normalised change in T+12 forecast errors at 100 hPa, all-sky microwave humidity assimilation minus no microwave humidity assimilation. No significance testing has been applied; however, the noisiness of the field should be an indication that small-scale features are not significant.

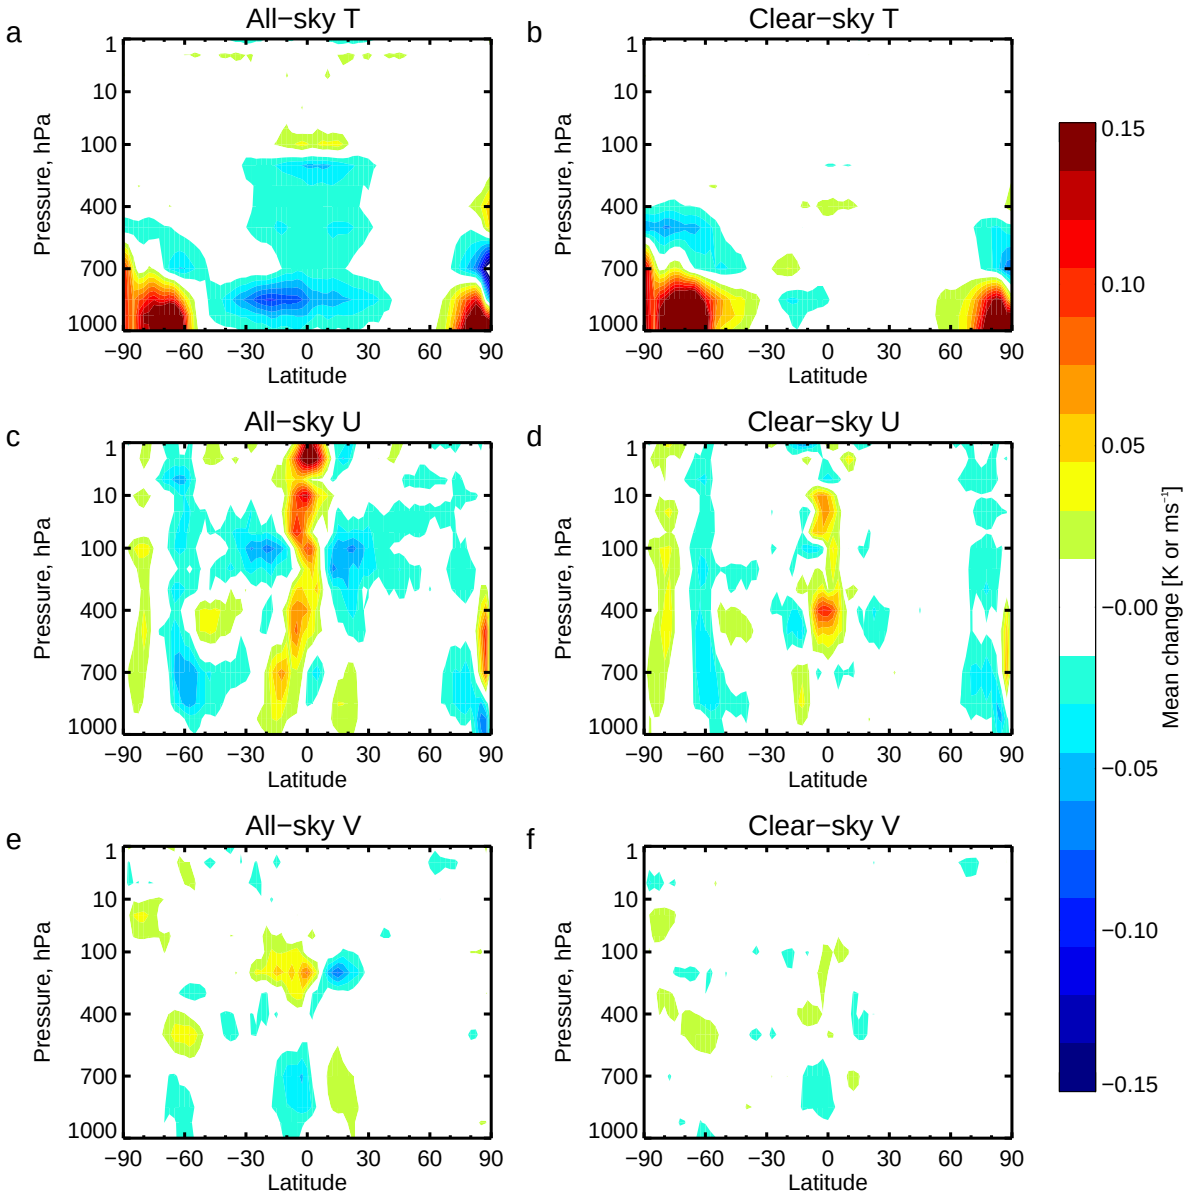

Figure 24: Mean change in T+12 forecast fields: with minus without microwave humidity assimilation. Left panels (a,c,e) for all-sky assimilation; right panels (b,d,f) for clear-sky assimilation.

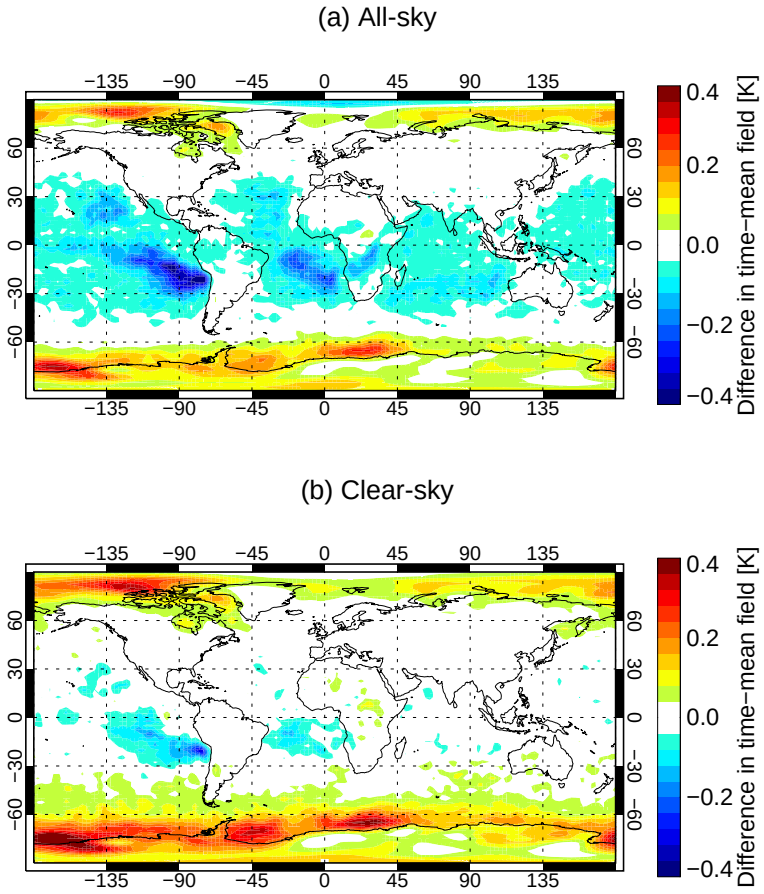

Figure 25: Mean change in T+12 temperatures at 850 hPa. This is the difference between experiments with and without microwave humidity assimilation, in the context of the otherwise full observing system.

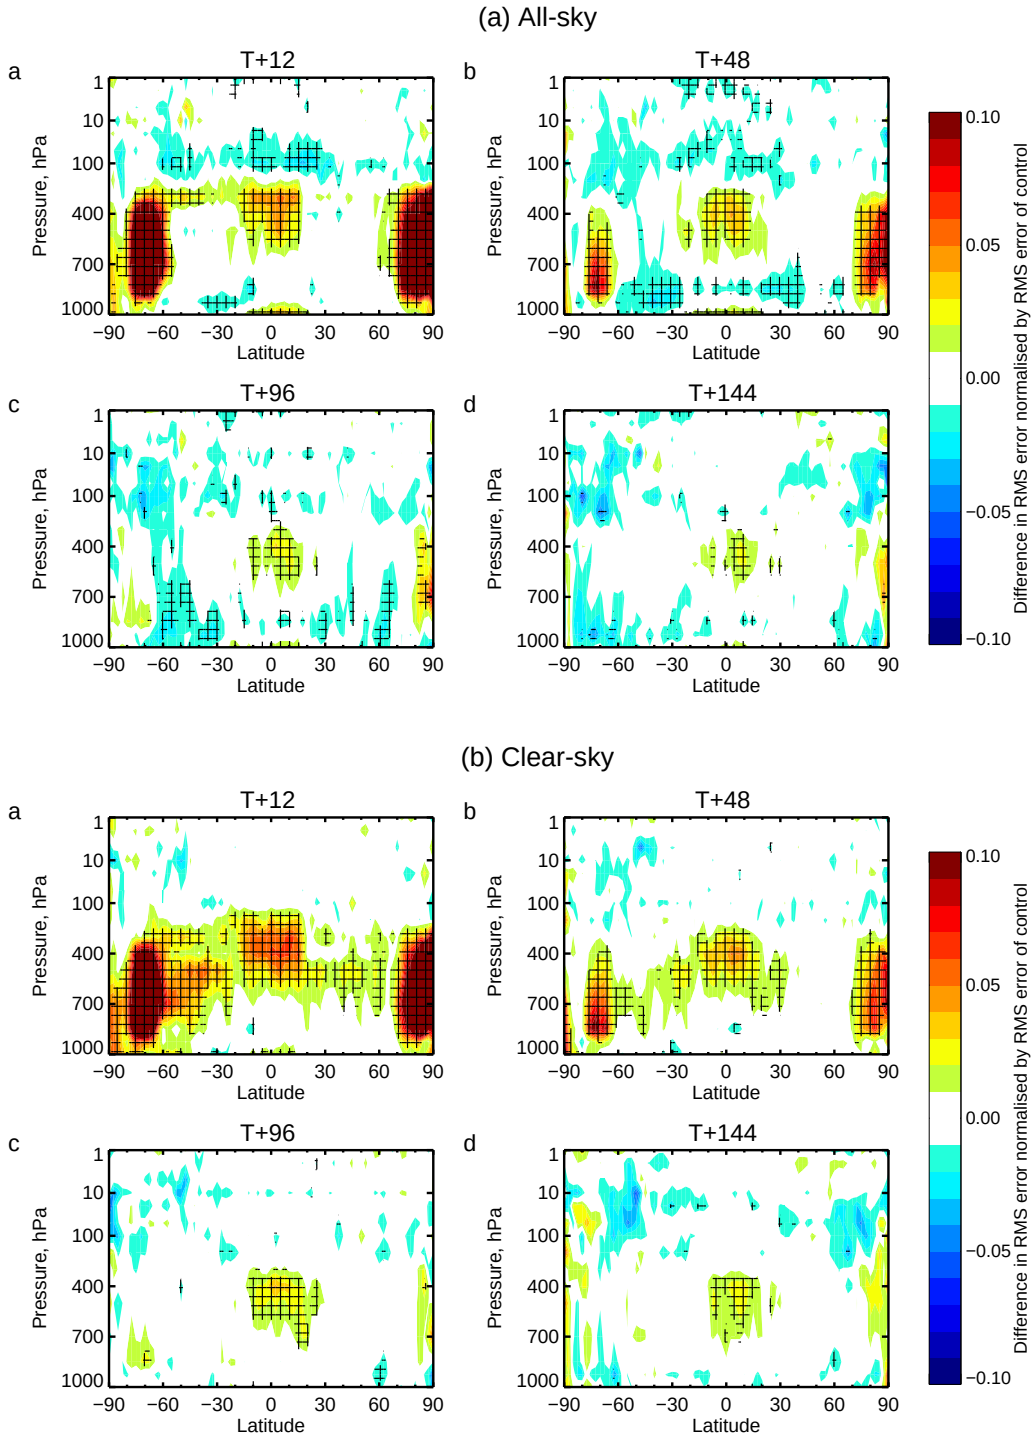

Figure 26: Normalised change in RMS of relative humidity forecast errors when all-sky (top) or clear-sky (bottom) microwave humidity assimilation is added to the otherwise full observing system. Reductions in errors are desirable and are shown in blue. Cross-hatching indicates statistical significance at the 95% confidence level. Results are based on a total of six months experimentation, winter and summer, and are based on 322 to 360 forecasts depending on the forecast range. Verification is against own analysis.

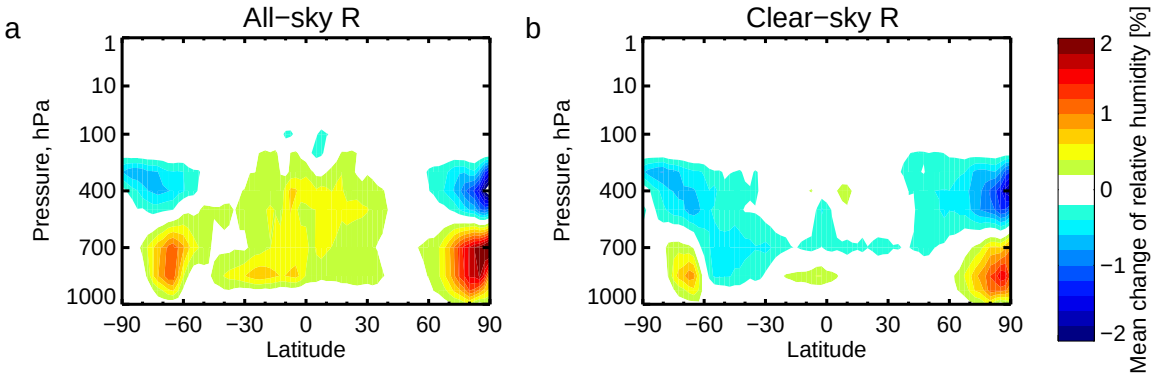

Figure 27: Mean change in T+12 relative humidity: with minus without microwave humidity assimilation for (a) for all-sky assimilation; (b) clear-sky assimilation.

troposphere and (with a more oceanic coverage) by up to 6% in HIRS channels 11 and 12, also sensitive to mid and upper-tropospheric humidity. Lower levels (channels 6 and 7 of HIRS and 500 hPa–1000 hPa in the radiosondes) show smaller improvements in humidities. Going from clear-sky to all-sky generates a proportionately lower impact in the humidity fields than in the dynamic fields (contrast AMSU-A channel 6, Fig. 22) which might suggest the all-sky observations have proportionately more dynamical impact.

Since ERA-40 the impact of satellite humidity assimilation on the ECMWF precipitation spin-down in the ECMWF system has been of concern, so Fig. 29 checks the situation with and without all-sky microwave humidity sounding channels. Even without these observations, the model gives 13% more precipitation in hours 3–6 and 6–9 than at 0–3. All-sky humidity sounding observations boost this slightly to 16%. Geer et al. (2010) show that all-sky microwave imagers also have little impact on the spin-down. The precipitation spin-down is a broader problem in the ECMWF system. Still, the situation is much better than in the days of ERA-40, when short-range precipitation was over-estimated by at least 50%.

7 Conclusion

All-sky assimilation of 183 GHz microwave water vapour sounding channels from five instruments (four MHS and one SSMIS) will be part of the next ECMWF operational system, to be implemented in early 2015. These humidity sounding channels have 1% – 2% impact on dynamical forecasts out until at least day 6 across the tropics and midlatitudes. At day 5 in the southern hemisphere, the total benefit of microwave humidity sounding channels is now equivalent to 1½ hours of forecast skill. Hours of forecast skill do not sound much but developments across observations, assimilation and models have gained ECMWF a total of about 1 day of forecast skill over the last decade, so this is still a useful contribution. It is very hard to improve the skill of the full system, with all observations included. However, in a set of ‘reinitialisation’ experiments, all-sky microwave water vapour sounding on its own can replicate 50% of the impact of the full global observing system on day-3 wind forecasts in the southern hemisphere. In comparison, eight microwave temperature sounding instruments (seven AMSU-A and one ATMS) replicate 65% of the global observing system impact. All-sky humidity sounding contributes real medium-range dynamical forecast skill and it is catching up with the impact of microwave temperature sounders, currently the most important single observing system in NWP.

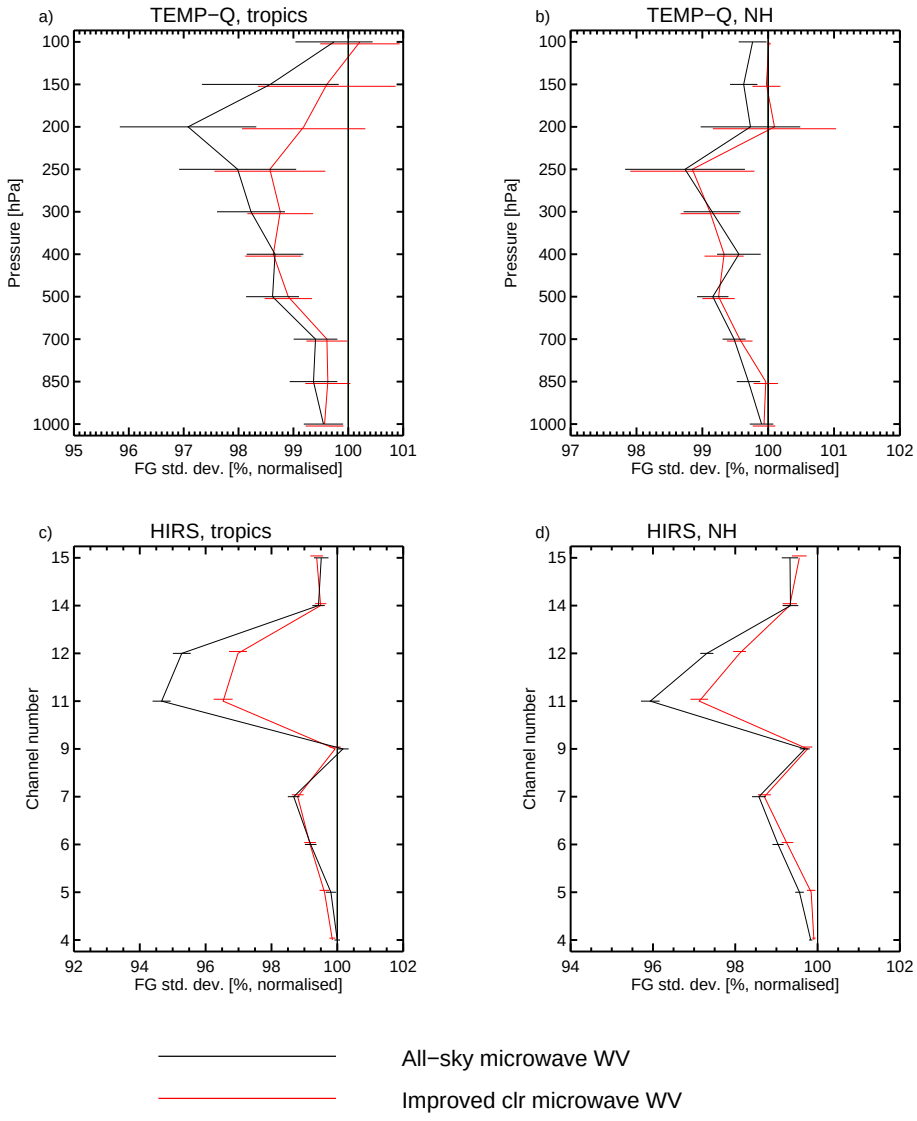

Figure 28: Standard deviation of FG departures for humidity-related observations as a percentage relative to the control. Error bars indicate the 95% confidence range. Control is without microwave humidity sounders but otherwise contains the full observing system; control is equivalent to 100% on these figures. Only tropical and NH fits are shown; results in the SH are similar. TEMP-Q refers to radiosonde humidities; HIRS channels 7, 11 and 12 are the main channels sensitive to humidity (in the lower, mid-upper and upper-troposphere respectively) while others are more sensitive to temperature or ozone.

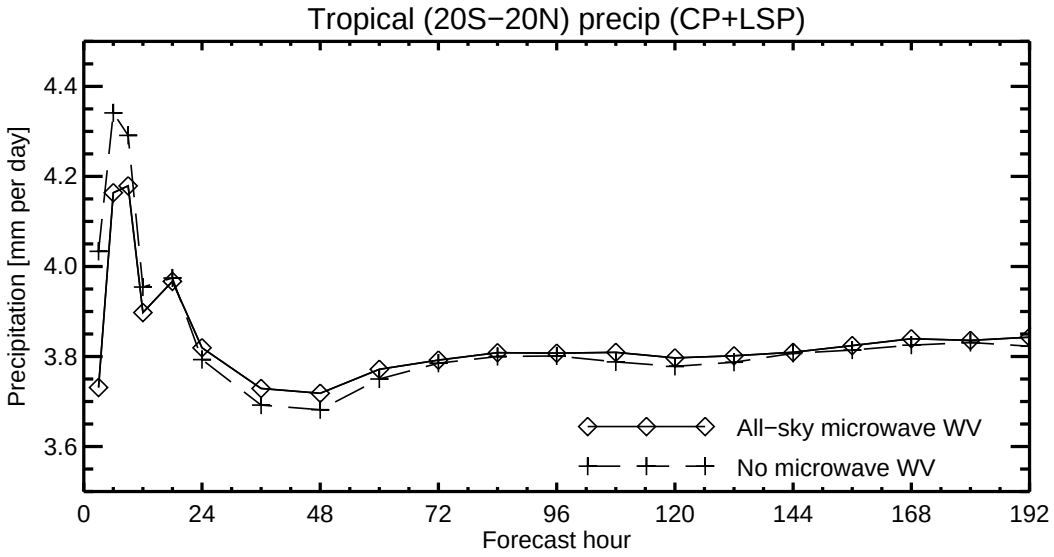

Figure 29: Mean accumulated precipitation in the tropics (20°N - 20°S) as a function of forecast time, for the summer experiments (August to October 2013) converted to an equivalent in mm per day. In the first 12 h, precipitation is accumulated in 3 h periods; in the next 12 h it is in 6 h periods. For the rest of the forecast time, precipitation is accumulated every 12 h.

The impact of of the humidity sounders comes through the model tracing of water vapour, cloud and precipitation. Single observation test cases demonstrate that:

- In the midlatitudes, 4D-Var adjusts dynamical initial conditions to better fit the observed water vapour, cloud and precipitation. Generally, adjustments to humidity initial conditions are a less significant part of the response;
- Even without a cloud or precipitation control variable, 4D-Var is capable of creating cloud and precipitation in the 12 h assimilation window to fit the observed cloud and precipitation;
- Although quality control is stringent and observation errors in cloudy errors are inflated, there is useful impact from cloud- and precipitation-affected observations.

We wanted to critically investigate the benefit of the all-sky approach, because a clear-sky approach (where clear-sky observation operators are used and cloudy scenes are discarded) provides forecast benefits with less complexity. Recent developments in clear-sky assimilation have more than doubled the amount of clear-sky data available to assimilate from MHS and SSMIS. In turn this has doubled the impact of clear-sky humidity sounding observations, now roughly a 1% impact on medium range forecast scores. Why not just continue to refine the clear-sky approach or to add humidity sounding data from new instruments in clear skies? The new clear-sky developments are important, but all-sky assimilation of humidity sounding channels doubles the impact again, to around 2%. Moreover, cloud and precipitation-affected data observe areas that clear-sky cannot, which brings a disproportionate impact:

- All-sky brings 10% more observations in the tropics, mostly around and inside convective precipitation. These make improvements of 1% to 3% in forecasts and in the fit to independent observations at the tropical tropopause. However, this may be related to systematic errors associated with

modelling of tropical convection (either in the forecast model or the observation operator) so it is important to monitor and better understand these issues.

- In the high midlatitudes, cloud-screening removes a high proportion of observations in the most dynamically active areas. All-sky assimilation doubles the number of observations available where model-tracing is likely to be most effective and most useful, in the midlatitude storm tracks. This directly benefits midlatitude synoptic forecasts through improved dynamical initial conditions.

Although the all-sky approach requires a more complex radiative transfer and careful attention to the observation error model, complexity is reduced in other areas. All-sky assimilation allows us to move away from unreliable and hard-to-maintain cloud-screening algorithms; further we do not have to worry about the sampling biases towards dry conditions that clear-sky screening likely introduces.

It is important to understand the mechanism behind the impact of cloudy and precipitation assimilation. This has been investigated by removing the tangent-linear and adjoint sensitivity to cloud and precipitation at the input to the observation operator. This reduces the impact on southern hemisphere day 3 forecasts from 50% to 46% in reinitialisation experiments, suggesting that cloud and precipitation sensitivity in the assimilation is not the dominant factor. Instead, based also on the single observation tests, the primary impact comes from being able to do water vapour tracing in areas where cloud or precipitation is present. Expectations about the current impact of cloud and precipitation assimilation in operational systems need to be realistic. Given the size of cloud and precipitation representivity errors in forecast models it is impossible for a global, full observing system analysis to fit every detail of cloud and precipitation in the observations. However, cloud and precipitation sensitivity is still an important secondary effect and it is likely to become more important as the quality of cloud and precipitation forecasts improves.

There are areas where all-sky assimilation of microwave humidity channels could be improved:

- Quality control is over-active, removing good observations that could have a useful impact.
- Using a scattering index predictor, the observation error model represents the dominant source of random error, which is the representivity of heavy precipitation in the forecast model. However, we need to deal with large background errors in clear-sky water vapour; our observation errors might currently be too large in clear areas. Also there might be benefit in representing frontal cloud contributions to the error budget in the lower peaking channels.
- The most significant remaining causes of microwave humidity sounder data rejection are high altitude and snow-covered land surfaces. Hence the greatest potential for adding more data is now in these areas, though it must be noted some of these regions can be dominated by high pressure and may not be as useful for model-tracing of synoptic features as the maritime storm tracks.

There are biases between model and microwave humidity observations that, although they do not have a negative impact on the current all-sky assimilation, need to be addressed with improvements both to forecast models and observation operators:

- There is insufficient scattering from frozen particles in deep convection (Geer and Baordo, 2014).
- Areas surrounding the tropical convection appear to be too dry, possibly through excessive dry detrainment, though there is a possibility of missing ice cloud effects - either through ice cloud missing in the model or if the observation operator is not sensitive enough to cloud ice.

- Over sea-ice, there is a 0.3 K warming and 1% moistening of the lower troposphere caused by assimilating humidity sounding channels. This improves fit to observations and consistency between analysis and forecasts, but it must indicate a limitation in data assimilation or in the model itself.

Based on the success of all-sky humidity sounding with MHS and SSMIS, it is hoped to further expand the all-sky assimilation approach to ATMS humidity sounding channels and any new microwave humidity sounding instruments that are added to the ECMWF system. Further, as pointed out by Chevallier et al. (2004), the easiest starting point for all-sky assimilation of infrared observations is the upper-tropospheric humidity sounding channels. These share many similarities with their equivalents in the microwave. Applying the all-sky technique to the infrared should bring additional benefits to forecasts through model-tracing of the humidity, cloud and precipitation fields.

Acknowledgements

Colleagues across ECMWF are thanked for their technical and scientific contributions; Erland Källén is thanked for reviewing the manuscript. Fabrizio Baordo was funded through a EUMETSAT fellowship.

References

Aires, F., C. Prigent, F. Bernardo, C. Jimnez, R. Saunders, and P. Brunel (2011). A tool to estimate land-surface emissivities at microwave frequencies (TELSEM) for use in numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society 137*(656), 690–699.

Allen, D., K. Hoppel, and D. Kuhl (2014). Wind extraction potential from 4D-Var assimilation of stratospheric O_3, N_2O, and H_2O using a global shallow water model. *Atmos. Chem. Phys. 14*(7), 3347–3360.

Allen, D. R., K. W. Hoppel, G. E. Nedoluha, D. D. Kuhl, N. L. Baker, L. Xu, and T. E. Rosmond (2013). Limitations of wind extraction from 4D-Var assimilation of ozone. *Atmos. Chem. Phys. 13*, 3501–3515.

Andersson, E., E. Hólm, P. Bauer, A. Beljaars, G. A. Kelly, A. P. McNally, A. J. Simmons, J.-N. Thépaut, and A. Tompkins (2007). Analysis and forecast impact of the main humidity observing systems. *Quart. J. Roy. Meteorol. Soc. 133*, 1473–1485.

Andersson, E. and H. Järvinen (1998). Variational quality control. *ECMWF Tech. Memo., 250, available from http://www.ecmwf.int*.

Andersson, E., J. Pailleux, J. N. Thépaut, J. R. Eyre, A. P. McNally, G. A. Kelly, and P. Courtier (1994). Use of cloud-cleared radiances in three/four-dimensional variational data assimilation. *Quart. J. Roy. Meteorol. Soc. 120*, 627–653.

Auligné, T., A. P. McNally, and D. P. Dee (2007). Adaptive bias correction for satellite data in a numerical weather prediction system. *Quart. J. Roy. Meteorol. Soc. 133*, 631–642.

Bannister, R. N. (2008). A review of forecast error covariance statistics in atmospheric variational data assimilation. II: Modelling the forecast error covariance statistics. *Quarterly Journal of the Royal Meteorological Society 134*(637), 1971–1996.

Baordo, F. and A. J. Geer (2014). All-sky assimilation of SSMI/S humidity sounding channels over land within the ECMWF system. *Quart. J. Roy. Meteorol. Soc.*, in preparation.

Baordo, F., A. J. Geer, and S. English (2012). SSMI/S radiances over land in the all-sky framework: one year report. *EUMETSAT/ECMWF Fellowship Programme Research Report No. 27, available from http://www.ecmwf.int*.

Baordo, F., A. J. Geer, and S. English (2013). All-sky assimilation of SSMI/S humidity sounding channels over land: second year report. *EUMETSAT/ECMWF Fellowship Programme Research Report No. 30, available from http://www.ecmwf.int*.

Bauer, P. (2001). Including a melting layer in microwave radiative transfer simulation for cloud. *Atmos. Res. 57*, 9–30.

Bauer, P., A. J. Geer, P. Lopez, and D. Salmond (2010). Direct 4D-Var assimilation of all-sky radiances: Part I. Implementation. *Quart. J. Roy. Meteorol. Soc. 136*, 1868–1885.

Bauer, P., E. Moreau, F. Chevallier, and U. O'Keeffe (2006). Multiple-scattering microwave radiative transfer for data assimilation applications. *Quart. J. Roy. Meteorol. Soc. 132*, 1259–1281.

Bechtold, P., N. Semane, P. Lopez, J.-P. Chaboureau, A. Beljaars, and N. Bormann (2013). Representing equilibrium and non-equilibrium convection in large-scale models. *ECMWF Tech. Memo., 705, available from http://www.ecmwf.int*.

Bengtsson, L. and K. Hodges (2005). On the impact of humidity observations in numerical weather prediction. *Tellus 57A*, 701–708.

Bennartz, R. (2000). Optimal convuolution of AMSU-B to AMSU-A. *J. Atmos. Ocean. Tech. 17*, 1215–1225.

Bonavita, M., L. Isaksen, and E. Holm (2012). On the use of EDA background error variances in the ECMWF 4D-Var. *ECMWF Research Department Memorandum 664*.

Bormann, N. and M. Bonavita (2013). Spread of the ensemble of data assimilations in radiance space. *ECMWF Technical Memorandum 708*.

Buehler, S. A. and V. O. John (2005). A simple method to relate microwave radiances to upper tropospheric humidity. *Journal of Geophysical Research: Atmospheres 110*.

Chevallier, F., P. Lopez, M. A. Tompkins, M. Janisková, and E. Moreau (2004). The capability of 4D-Var systems to assimilate cloud-affected satellite infrared radiances. *Quart. J. Roy. Meteorol. Soc. 130*, 917–932.

Compo, G. P., J. S. Whitaker, and P. D. Sardeshmukh (2006). Feasibility of a 100-year reanalysis using only surface pressure data. *Bull. Amer. Meteor. Soc. 87*, 175–190.

Daley, R. (1995). Estimating the wind field from chemical constituent observations: Experiments with a one-dimensional extended Kalman filter. *Mon. Weath. Rev. 123*(1), 181–198.

Dee, D. (2004). Variational bias correction of radiance data in the ECMWF system. In *ECMWF workshop proceedings: Assimilation of high spectral resolution sounders in NWP, 28 June – 1 July, 2004*, pp. 97–112. Eur. Cent. for Med. Range Weather Forecasts, Reading, UK, available from http://www.ecmwf.int.

Di Tomaso, E., N. Bormann, and S. English (2013). Assimilation of ATOVS radiances at ECMWF: third year EUMETSAT fellowship report. *EUMETSAT/ECMWF Fellowship Programme Research Report No. 29, available from http://www.ecmwf.int*.

English, S., T. McNally, N. Bormann, K. Salonen, M. Matricardi, A. Horanyi, M. Rennie, M. Janisková, S. Di Michele, A. Geer, E. Di Tomaso, C. Cardinali, P. de Rosnay, J. Muñoz Sabater, M. Bonavita, C. Albergel, R. Engelen, and J.-N. Thépaut (2013). Impact of satellite data. *ECMWF Tech. Memo., 711, available from http://www.ecmwf.int*.

English, S. J. (2008). The importance of accurate skin temperature in assimilating radiances from satellite sounding instruments. *Quart. J. Roy. Meteorol. Soc. 126*, 2911–2931.

English, S. J. and T. J. Hewison (1998). A fast generic microwave emissivity model. In T. Hayasaka, D. L. Wu, Y. Jin, and J. Jiang (Eds.), *Proceedings of SPIE, 3503: Microwave remote sensing of the environment*, pp. 288–300.

English, S. J., R. J. Renshaw, P. C. Dibben, A. J. Smith, P. J. Rayer, C. Poulsen, F. W. Saunders, and J. R. Eyre (2000). A comparison of the impact of TOVS and ATOVS satellite sounding data on the accuracy of numerical weather forecasts. *IEEE Trans. Geosci. Remote Sens. 46*, 403–408.

Eyre, J. R. (1991). A fast radiative transfer model for satellite sounding systems. *ECMWF Tech. Memo., 176, available from http://www.ecmwf.int*.

Forbes, R. M., A. M. Tompkins, and A. Untch (2011). A new prognostic bulk microphysics scheme for the IFS. *ECMWF Technical Memorandum 649*.

Fujita, T. T., E. W. Fearl, and W. E. Shenk (1975). Satellite-tracked cumulus velocities. *J. App. Met. 14*(3), 407–413.

Geer, A. J. (2013). All-sky assimilation: better snow-scattering radiative transfer and addition of SSMIS humidity sounding channels. *ECMWF Tech. Memo., 706, available from http://www.ecmwf.int*.

Geer, A. J. and F. Baordo (2014). Improved scattering radiative transfer for frozen hydrometeors at microwave frequencies. *Atmos. Meas. Tech. 7*, 1839–1860.

Geer, A. J. and P. Bauer (2010). Enhanced use of all-sky microwave observations sensitive to water vapour, cloud and precipitation. *Published simultaneously as ECMWF Technical Memoranda 620 and ECMWF/EUMETSAT fellowship reports 20*.

Geer, A. J. and P. Bauer (2011). Observation errors in all-sky data assimilation. *Quart. J. Roy. Meteorol. Soc. 137*, 2024–2037.

Geer, A. J., P. Bauer, and P. Lopez (2010). Direct 4D-Var assimilation of all-sky radiances: Part II. Assessment. *Quart. J. Roy. Meteorol. Soc. 136*, 1886–1905.

Geer, A. J., P. Bauer, and C. W. O'Dell (2009). A revised cloud overlap scheme for fast microwave radiative transfer. *J. App. Meteor. Clim. 48*, 2257–2270.

Gérard, E. and R. Saunders (1999). Four-dimensional variational assimilation of Special Sensor Microwave / Imager total column water vapour in the ECMWF model. *Quart. J. Roy. Meteorol. Soc. 125*, 3077–3102.

Hewison, T. J. and S. J. English (1999). Airborne retrievals of snow and ice surface emissivity at millimeter wavelengths. *IEEE Trans. Geosci. Rem. Sens. 37*(4), 1871–1879.

Hólm, E., E. Andersson, A. Beljaars, P. Lopez, J.-F. Mahfouf, A. Simmons, and J.-N. Thepaut (2002). Assimilation and modelling of the hydrological cycle: ECMWF's status and plans. *ECMWF Tech. Memo., 383, available from http://www.ecmwf.int*.

Hong, G., G. Heygster, J. Miao, and K. Kunzi (2005). Detection of tropical deep convective clouds from AMSU-B water vapor channels measurements. *J. Geophys. Res. 110*, D05205.

Janisková, M. and P. Lopez (2013). Linearized physics for data assimilation at ECMWF. In *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications (Vol. II)*, pp. 251–286. Springer.

Järvinen, H. and P. Unden (1997). Observation screening and background quality control in the ECMWF 3D-Var data assimilation system. *ECMWF Tech. Memo., 236, available from http://www.ecmwf.int*.

Jones, A. S. and T. H. Vonder Haar (1990). Passive microwave remote sensing of cloud liquid water over land regions. *J. Geophys. Res. 95*, 16673–16683.

Joseph, J., W. J. Wiscombe, and J. A. Weinman (1976). The delta-Eddington approximation for radiative flux transfer. *J. Atmos. Sci. 33*, 2452–2459.

Karbou, F., P. Bauer, A. Geer, and W. Bell (2008). Exploitation of microwave sounder/imager data over land surfaces in the presence of clouds and precipitation. *EUMETSAT Hydrology SAF visiting scientist report*, available http://hsaf.meteoam.it/documents/reference/VS–37–ECMWF–Karbou–final–report.pdf.

Karbou, F., C. Prigent, L. Eymard, and J. R. Pardo (2005). Microwave land emissivity calculations using AMSU measurements. *IEEE Trans. Geosci. Remote Sensing 43*(5), 948–959.

Kazumori, M. and S. J. English (2014). Use of the ocean surface wind direction signal in microwave radiance assimilation. *Quart. J. Roy. Meteorol. Soc.*, in press.

Kazumori, M., A. J. Geer, and S. J. English (2014). Effects of all-sky assimilation of GCOM-W1/AMSR2 radiances in the ECMWF system. *Quart. J. Roy. Meteorol. Soc.*, in preparation, available as ECMWF technical memorandum 732.

Kelly, G. A., P. Bauer, A. J. Geer, P. Lopez, and J.-N. Thépaut (2008). Impact of SSM/I observations related to moisture, clouds and precipitation on global NWP forecast skill. *Mon. Weather Rev. 136*, 2713–2726.

Krzeminski, B., N. Bormann, F. Karbou, and P. Bauer (2009). Improved use of surface-sensitive microwave radiances at ECMWF. In *Proc. EUMETSAT Meteorol. Satell. Conf*, pp. 21–25.

Kunkee, D., G. Poe, D. Boucher, S. Swadley, Y. Hong, J. Wessel, and E. Uliana (2008). Design and evaluation of the first Special Sensor Microwave Imager/Sounder. *IEEE Trans. Geosci. Remote Sensing 46*, 863–883.

Liu, G. (2008). A database of microwave single-scattering properties for nonspherical ice particles. *Bull. Am. Met. Soc. 111*, 1563–1570.

Liu, Q., F. Weng, and S. English (2011). An improved fast microwave water emissivity model. *IEEE Trans. Geosci. Remote Sensing 49*, 1238–1250.

Lopez, P. and E. Moreau (2005). A convection scheme for data assimilation: Description and initial tests. *Quart. J. Roy. Meteorol. Soc. 131*, 409–436.

Peubey, C. and A. P. McNally (2009). Characterization of the impact of geostationary clear-sky radiances on wind analyses in a 4D-Var context. *Quart. J. Roy. Meteorol. Soc. 135*, 1863 – 1876.

Peuch, A., J. N. Thepaut, and J. Pailleux (2000). Dynamical impact of total-ozone observations in a four-dimensional variational assimilation. *Quart. J. Roy. Meteorol. Soc. 126*, 1641 – 1659.

Prigent, C., W. B. Rossow, and E. Matthews (1997). Microwave land surface emissivities estimated from SSM/I observations. *J. Geophys. Res. 102*(D18), 21867–21890.

Rabier, F., H. Järvinen, E. Klinker, J.-F. Mahfouf, and A. Simmons (2000). The ECMWF operational implementation of four-dimensional variational assimilation. I: Experimental results with simplified physics. *Quart. J. Roy. Meteorol. Soc. 126*, 1148–1170.

Riishøjgaard, L. P. (1996). On four-dimensional variational assimilation of ozone data in weather-prediction models. *Quart. J. Roy. Meteorol. Soc. 122*, 1545–1571.

Robel, J. (2009). NOAA KLM user's guide (February 2009 revision). *Available from http://www.ncdc.noaa.gov/oa/pod-guide/ncdc/docs/intro.htm*.

Roberts, N. M. and H. W. Lean (2008). Scale-selective verification of rainfall accumulations from high-resolution forecasts of convective events. *Mon. Wea. Rev. 136*, 78–96.

Saunders, R., J. Hocking, P. Rayer, M. Matricardi, A. Geer, N. Bormann, P. Brunel, F. Karbou, and F. Aires (2012). RTTOV-10 science and validation report. NWPSAF-MO-TV-023 v1.11, EUMETSAT NWP-SAF.

Smagorinsky, J., K. Miyakoda, and R. Strickler (1970). The relative importance of variables in initial conditions for dynamical weather prediction. *Tellus 22*(2), 141–157.

Sreerekha, T., S. Buehler, U. O'Keeffe, A. Doherty, C. Emde, and V. John (2008). A strong ice cloud event as seen by a microwave satellite sensor: Simulations and observations. *J. Quant. Spectrosc. and Radiat. Transfer 109*(9), 1705 – 1718.

Tompkins, A. M. and M. Janisková (2004). A cloud scheme for data assimilation: Description and initial tests. *Quart. J. Roy. Meteorol. Soc. 130*, 2495–2517.

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

Velden, C. S., C. M. Hayden, S. J. Nieman, W. P. Menzel, S. Wanzong, and J. S. Goerss (1997). Upper-tropospheric winds derived from geostationary satellite water vapor observations. *Bull. Am. Met. Soc. 78*(2), 173–195.

Zou, X., Z. Qin, and F. Weng (2013). Improved quantitative precipitation forecasts by MHS radiance data assimilation with a newly added cloud detection algorithm. *Mon. Weath. Rev. 141*(9), 3203–3221.