

Mitigating surface temperature errors using approximate radiation updates

Robin J. Hogan and Alessio Bozzo

Research Department

February 2015

*This paper has not been published and should be regarded as an Internal Report from ECMWF.
Permission to quote from it should be obtained from the ECMWF.*



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

©Copyright 2015

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

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

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

Abstract

Due to computational expense, the current radiation scheme in the IFS is called infrequently in time (every hour in the high resolution forecast and every three hours in the ensemble system) and on a reduced spatial grid. This can lead to large surface temperature errors at coastal land points due to surface fluxes computed over the ocean being used where the skin temperature and surface albedo are very different. It can also lead to a lag in the diurnal cycle of surface temperature. This memorandum describes a computationally efficient solution to these problems, in which the surface longwave and shortwave fluxes are updated every timestep and gridpoint according to the local skin temperature and albedo. In order that energy is conserved, it is necessary to compute the change to the net flux profile consistent with the changed surface fluxes. The longwave radiation scheme has been modified to compute also the rate of change of the profile of upwelling longwave flux with respect to the value at the surface. Then at each gridpoint and timestep, the upwelling flux and heating-rate profiles are updated using the new value of skin temperature. The computational cost of performing approximate radiation updates is only 2% of the cost of the full radiation scheme, so increases the overall cost of the model by only of order 0.2%. Testing the new scheme by running daily 5-day forecasts over an eight-month period reveals significant improvement in 2-m temperature forecasts at coastal stations compared to observations.

1 Introduction

Given the complexity of gaseous absorption spectra, it is remarkable how rapidly modern radiation schemes for General Circulation Models (GCMs) are able to compute accurate broadband radiative fluxes and heating rates across the wide range of conditions found in the Earth's atmosphere. This is possible thanks to key approximations such as the two-stream approximation ([Schuster, 1905](#)) where the diffuse radiation field is treated by radiation travelling in just two discrete directions, and the correlated- k distribution method ([Lacis and Oinas, 1991](#)), where the effect of $O(10^6)$ absorption lines is captured by $O(10^2)$ quasi-monochromatic calculations. expensive parts of a weather or climate model and is generally too expensive to run every timestep and gridpoint. [Morcrette et al. \(2008b\)](#) provided a history of the reduced radiation resolution in time and space at ECMWF. In 2007, the "McRad" radiation scheme became operational at ECMWF ([Morcrette et al., 2008a](#)), which represented the full shortwave and longwave spectrum with 252 pseudo-monochromatic bands via its use of the Rapid Radiative Transfer Model for GCMs (RRTM-G) described by [Mlawer et al. \(1997\)](#), [Morcrette et al. \(2001\)](#) and [Clough et al. \(2005\)](#). This led to an increase in the computational cost of the radiation scheme by around a factor of 3.5, necessitating a further reduction of the spatial resolution of the radiation calculations relative to the model resolution. Operational practice at the time of writing is to run the high-resolution model at a spectral resolution of $T_L 1279$ with the radiation scheme run every 1-h at an effective resolution of $T_L 511$, and to run the ensemble prediction system at a resolution of $T_L 639$ with the radiation scheme run every 3-h at a resolution of $T_L 255$. Thus, in both cases the radiation scheme is run on 6.25 times fewer gridpoints than the rest of the model physics, and for only one in every 6 or 9 model timesteps.

[Morcrette \(2000\)](#) examined the impact of temporal and spatial sampling of radiation on ECMWF forecasts and analyses and found that (a) there was negligible degradation of forecast skill in terms of 500-hPa geopotential in at least the first 7 days of the forecast, but biases emerged in seasonal forecasts, (b) there were locally significant changes to skin temperature and cloudiness, which were more sensitive to temporal than spatial sampling of the radiation, and (c) the weakening of the coupling between rapidly varying cloud fields and the radiation field led to a change to the model's climate sensitivity. In recent years there have been complaints from forecast users, particularly in Norway, that forecasts of nighttime coastal 2-m temperature can sometimes be too cold by in excess of 10 K (personal communication Linus Magnusson and Tim Hewson, 2014). This is essentially due to the surface net longwave radiation being

taken from an adjacent sea point with a warmer surface temperature, and kept constant between calls to the radiation scheme.

The forecast errors associated with intermittent radiation have prompted a number of attempts to make radiation schemes more efficient, for example by running only a randomly selected subset of the spectral intervals in each profile (Bozzo et al., 2015), or by running only the optically thin parts of the spectrum (where the effects of clouds and the surface are felt) at higher resolution (Manners et al., 2009). Another approach is to perform approximate updates of the radiation fields between calls to the radiation scheme. In the shortwave, the IFS already accounts partially for the changing solar zenith angle between calls to the radiation scheme by computing the shortwave flux profile for an incoming top-of-atmosphere (TOA) flux of unity, and then at every timestep and gridpoint multiplying it by the local value of incoming TOA flux, which is proportional to the cosine of the solar zenith angle μ_0 (Morcrette, 2000). Manners et al. (2009) proposed a more accurate scheme to account also for the μ_0 -dependence of the path length of the direct solar beam through the atmosphere.

In this memorandum we consolidate a number of new and existing methods for approximately updating the radiation fields between calls to the full radiation scheme, and examine their effects on forecasts, particularly at the surface. In the shortwave, we propose a scheme to account for large horizontal variations in surface albedo that also includes the effect of back-reflection from the atmosphere. In the longwave, not only are surface fluxes modified to respond immediately to changes to skin temperature, but also the profiles of upwelling and downwelling fluxes, in order to capture the strong coupling between surface temperature and the temperature of the lowest few hundred metres of the atmosphere.

Sections 2 and 3 describe the approximate updates applied in the longwave and shortwave, respectively. Section 4 presents a case study of a global forecast in which the impact of the scheme on both coastal errors and errors in the diurnal cycle of surface temperature is demonstrated. Then in section 5, a total of eight months of daily 5-day forecasts are run with different model configurations to assess the improvement to the forecasts. Coastal examples are presented from New York and Arabia in section 4 and from Norway in section 5.

2 Longwave method

The modifications to the model needed to provide an approximate update of the longwave net fluxes (surface and atmosphere) at every timestep and gridpoint are in two parts. Firstly, the radiation scheme is modified to output extra variables such as the profile of partial derivative of the upwelling flux profile as described in section 2.1. Offline radiation calculations are carried out in section 2.2 to illustrate the typical shape of these profiles and the radiative coupling between the surface and lowest layers of the atmosphere. Secondly, these extra variables are used to update the net fluxes every timestep and gridpoint as described in section 2.3. Section 2.4 then describes two minor additional changes to the code concerning the calculation of skin temperature.

2.1 Extra variables from the radiation scheme

The only output from the longwave radiation scheme currently returned and used by the rest of the model is the profile of net longwave flux at each model half-level including the surface, $L_{i-1/2}^n = L_{i-1/2}^\downarrow - L_{i-1/2}^\uparrow$, where i is the vertical layer index counting down from 1 at the top. We modify the scheme to return two additional variables: (1) the surface downwelling flux $L_{\text{surf}}^\downarrow$, and (2) the partial derivative of

upwelling longwave flux at all model half-levels with respect to the surface upwelling longwave flux, i.e. $\partial L_{i-1/2}^\uparrow / \partial L_{\text{surf}}^\uparrow$. This is a partial derivative in the sense that we are treating the atmospheric temperature and composition constant.

The flux $L_{\text{surf}}^\downarrow$ is straightforward since it is already computed by the radiation scheme but is not passed to the rest of the model. We calculate $\partial L_{i-1/2}^\uparrow / \partial L_{\text{surf}}^\uparrow$ as follows. The upwelling and downwelling longwave fluxes are currently computed without scattering via n_g independent pseudo-monochromatic calculations (known as g points) representing the full longwave spectrum. In RRTM-G, $n_g = 140$. Denoting g as the index to g points, the longwave upwelling flux at any half-level may be written as

$$L_{i-1/2}^\uparrow = \varepsilon \sum_{g=1}^{n_g} B_g(T_{\text{surf}}) \tau_{i\dots n,g} + f_g, \quad (1)$$

where ε is the surface emissivity (here assumed constant across the longwave spectrum), $B_g(T_{\text{surf}})$ is the Planck function (as a flux in W m^{-2}) at surface temperature T_{surf} integrated across the parts of the spectrum corresponding to one g point, f_g is the contribution to the upwelling flux from emission by the atmosphere (including downward emission that is reflected back up by the surface) and $\tau_{i\dots n,g}$ is the transmittance of the atmosphere between layers i and n inclusive, which may be written as the product of the transmittances of individual layers:

$$\tau_{i\dots n,g} = \prod_{k=i}^n \tau_{k,g}. \quad (2)$$

Taking the derivative of (1) with respect to T_{surf} we obtain

$$\frac{\partial L_{i-1/2}^\uparrow}{\partial T_{\text{surf}}} = \varepsilon \sum_{g=1}^{n_g} \frac{\partial B_g(T_{\text{surf}})}{\partial T_{\text{surf}}} \tau_{i\dots n,g}. \quad (3)$$

Under the assumption that emissivity is constant across the longwave spectrum, we may write the surface upwelling flux using the Stefan-Boltzmann law as

$$L_{\text{surf}}^\uparrow = \varepsilon \sigma T_{\text{surf}}^4 + (1 - \varepsilon) L_{\text{surf}}^\downarrow, \quad (4)$$

where σ is the Stefan-Boltzmann constant. The derivative of this with respect to T_{surf} is

$$\frac{\partial L_{\text{surf}}^\uparrow}{\partial T_{\text{surf}}} = 4\varepsilon \sigma T_{\text{surf}}^3. \quad (5)$$

Combining with (3) yields

$$\frac{\partial L_{i-1/2}^\uparrow}{\partial L_{\text{surf}}^\uparrow} = \frac{1}{4\sigma T_{\text{surf}}^3} \sum_{g=1}^{n_g} \frac{\partial B_g(T_{\text{surf}})}{\partial T_{\text{surf}}} \tau_{i\dots n,g}, \quad (6)$$

Since B_g is held as a look-up table versus temperature for each g point, computing its derivative numerically in the code is trivial. Note that when half-level $i - 1/2$ corresponds to the surface layer, i.e. $i = n + 1$, the partial derivative in (6) becomes unity.

The current longwave radiation scheme is particularly well suited to adding the computation of $\partial L_{i-1/2}^\uparrow / \partial L_{\text{surf}}^\uparrow$ alongside the existing calculation of fluxes for two reasons. Firstly, the individual g points in the McICA scheme treat the atmosphere as plane-parallel, i.e. there is no partial cloudiness treated within each single monochromatic calculation. Secondly, the neglect of scattering reduces the calculation to a first absorption-emission pass down through the atmosphere to compute the downwelling fluxes followed by a second absorption-emission pass back up through the atmosphere to compute the upwelling fluxes. In both passes the layer transmittances $\tau_{i,g}$ are used, and so we add the computation of the partial derivatives to the upward pass.

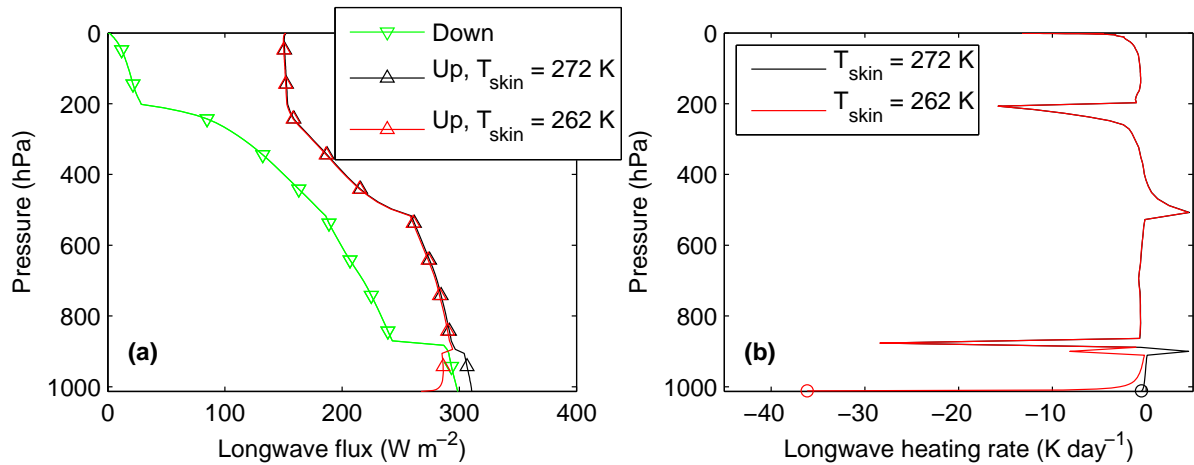


Figure 1: Illustration of the impact of skin temperature (T_{skin}) on (a) longwave flux and (b) heating rate profiles by applying RRTM-G to the Mid-Latitude Winter standard atmosphere with a liquid cloud (mixing ratio 0.2 g kg^{-1} , effective radius $10 \text{ }\mu\text{m}$, cloud fraction 0.75) between 860 and 900 hPa and an ice cloud (mixing ratio 0.05 g kg^{-1} , effective radius $50 \text{ }\mu\text{m}$, cloud fraction 0.75) between 200 and 500 hPa. The black lines depict the control scenario in which T_{skin} is equal to the air temperature at the lowest half-level, while the red lines depict the case when T_{skin} is reduced by 10 K but the air temperature is kept constant. The red lines could be reproduced exactly by taking the control profile and updating the upwelling fluxes using (6) and (7). The downwelling flux profile depicted by the green line in panel a is the same for any value of T_{skin} since when scattering is neglected it does not depend on surface temperature. The heating rates of the lowest model layer are shown by the circles in panel b.

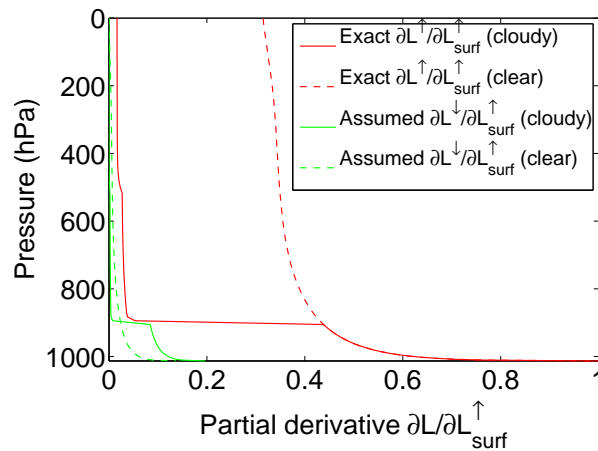


Figure 2: Partial derivative of upwelling and downwelling fluxes with respect to the surface upwelling flux for clear and cloudy versions of the profile shown in Fig. 1: (red) the exact derivative of upwelling flux, and (green) the assumed derivative of downwelling flux accounting approximately for the warming of the atmosphere.

2.2 Impact of different surface and near-surface temperatures

To illustrate the typical shape of the partial-derivative profile and the importance of the opacity of the lowest few hundred metres of the atmosphere, some offline radiation calculations have been performed using the two-stream radiation scheme of Pincus and Stevens (2009), which is very similar to that in the IFS. Figure 1 compares the flux and heating rate profiles for a cloudy Mid-Latitude Winter standard atmosphere (McClatchey et al., 1972), using the same 137 pressure levels as the IFS, with two different skin temperatures: one the same as the lowest atmospheric temperature and the other 10 K colder. This is intended to represent a profile over sea being applied over a neighbouring land point with a colder

surface. The latter leads to a very strong atmospheric cooling in the lowest 100 hPa of the atmosphere, with a peak value of -36 K day^{-1} in the lowest model layer. This is because the atmosphere is largely opaque to longwave radiative transfer and there is a strong imbalance between the energy emitted by these layers and the energy absorbed from the colder underlying surface. The sign of the heating rate at cloud base is also reversed. It should be noted that since the near-surface atmospheric cooling is occurring in the opaque parts of the longwave spectrum, the peak cooling is not much affected by the presence of cloud: if the cloud is removed then the peak cooling rate is reduced in magnitude by only 1 K day^{-1} (not shown).

While these cooling rates are large, it should be pointed out that the lowest model level is only 20 m thick, so the total energy involved is modest. We would expect both the radiative tendencies and the turbulent mixing scheme in the model to keep the near-surface atmospheric temperatures coupled to the temperature of the surface (which should itself be much improved by the use of an approximate radiation update). The rapid response of the near-surface atmospheric temperatures suggests that it may be necessary to update the downwelling fluxes as well.

The red lines in Fig. 2 depict the partial-derivative profile $\partial L^\uparrow / \partial L_{\text{surf}}^\uparrow$ for clear and cloudy versions of the same profile. It is striking how rapidly this curve decreases with height above the surface, indicating that around half of the emitted radiation from the surface is absorbed in the lowest 500 m of the atmosphere. In the cloudy case, most of the remainder is then absorbed at cloud base. Currently the Met Office use (4) to update L_{surf}^\uparrow each timestep and then assume any excess is lost to space, thereby neglecting any change to atmospheric heating. While this is better than not updating L_{surf}^\uparrow at all (as in the IFS), the shape of these curves show that a much better approximation would be to allow much of this radiation to be absorbed in the atmosphere. The green lines are discussed in the next section.

2.3 Updating the net longwave flux profile

Section 2.1 described how the longwave radiation scheme was modified to provide the surface downwelling flux $L_{\text{ref},i-1/2}^\downarrow$ and a profile of the partial derivatives $\partial L_{i-1/2}^\uparrow / \partial L_{\text{surf}}^\uparrow$, in addition to the profile of net longwave flux $L_{\text{ref},i-1/2}^n$ (where ‘‘ref’’ indicates reference values output by the radiation scheme, which will subsequently be modified to respond to local conditions). Since the radiation scheme is run on a lower resolution grid than the rest of the model, these variables need to be interpolated back on to the native model grid where they are available for several timesteps until the radiation scheme is called again.

The current version of the model uses the net longwave flux profile to compute the profile of atmospheric heating rate in each of these intervening timesteps, naturally predicting the same heating rate each time. The net flux at the surface is used in the surface energy budget equation and is also held fixed between calls to the radiation scheme, even if the surface temperature changes. This is the principal cause of the forecast errors near coastlines discussed in the introduction.

To allow the net longwave flux profile, including the surface value, to respond to any change in surface temperature, we may compute a new upwelling longwave flux from (4) and

$$L_{i-1/2}^\uparrow = L_{\text{ref},i-1/2}^\uparrow + \left(L_{\text{surf}}^\uparrow - L_{\text{ref},\text{surf}}^\uparrow \right) \frac{\partial L_{i-1/2}^\uparrow}{\partial L_{\text{surf}}^\uparrow}. \quad (7)$$

The reference upwelling longwave flux from the surface is computed from the stored variables simply as $L_{\text{ref},\text{surf}}^\uparrow = L_{\text{ref},\text{surf}}^\downarrow - L_{\text{ref},\text{surf}}^n$. If surface temperature were the only thing to change then (7) would exactly match what would be output from the radiation scheme if it were run at high temporal and spatial

resolution. In reality, however, the atmospheric temperature and composition will also change in time and space. The most important change is that atmospheric temperature in the lowest few hundred metres of the atmosphere is strongly coupled to surface temperature, due both to turbulent heat fluxes and to longwave radiative exchange (see Fig. 1). This means that increased upwelling longwave radiation tends to be coupled to increased downwelling. A first-order representation of this effect is to assume that the change in surface downwelling is a fixed fraction γ of the change to the surface upwelling:

$$L_{\text{surf}}^{\downarrow} - L_{\text{ref,surf}}^{\downarrow} = \gamma \left(L_{\text{surf}}^{\uparrow} - L_{\text{ref,surf}}^{\uparrow} \right), \quad (8)$$

where a value of $\gamma = 0.2$ is justified a-posteriori in section 5, where it is found to provide the best match in global model simulations when compared to runs with radiation called every timestep.

Accompanying this assumption it is necessary to assume a profile for the partial derivatives of the downwelling fluxes. We assume that the profile has the same shape as the profile of partial derivatives of upwelling fluxes, but scaled and offset so that the surface value is $\partial L_{\text{surf}}^{\downarrow} / \partial L_{\text{surf}}^{\uparrow} = \gamma$ and the top-of-atmosphere value is $\partial L_{\text{TOA}}^{\downarrow} / \partial L_{\text{surf}}^{\uparrow} = 0$. This is achieved by

$$\frac{\partial L_{i-1/2}^{\downarrow}}{\partial L_{\text{surf}}^{\uparrow}} = \gamma \frac{\partial L_{i-1/2}^{\uparrow} / \partial L_{\text{surf}}^{\uparrow} - \partial L_{\text{TOA}}^{\uparrow} / \partial L_{\text{surf}}^{\uparrow}}{1 - \partial L_{\text{TOA}}^{\uparrow} / \partial L_{\text{surf}}^{\uparrow}}. \quad (9)$$

The green lines in Fig. 2 depict the partial derivative of downwelling flux under this assumption. The model deals with net fluxes, so the updated net flux is actually computed from (4), (9) and

$$L_{i-1/2}^n = L_{\text{ref},i-1/2}^n + \left(L_{\text{surf}}^{\uparrow} - L_{\text{ref,surf}}^{\uparrow} \right) \left(\frac{\partial L_{i-1/2}^{\downarrow}}{\partial L_{\text{surf}}^{\uparrow}} - \frac{\partial L_{i-1/2}^{\uparrow}}{\partial L_{\text{surf}}^{\uparrow}} \right). \quad (10)$$

2.4 Computation of skin temperature

The change in T_{skin} with time is currently computed using the surface energy balance equation, but since several terms in this equation depend on T_{skin} (and in the case of upwelling longwave radiation the dependence is nonlinear), this is done by linearizing the surface energy budget equation with respect to T_{skin} . The current version of the IFS linearizes about the T_{skin} at the time the radiation scheme was last called. Since the scheme described in this document updates the radiative fluxes at each timestep, it is necessary to change this to linearize about the T_{skin} at the last timestep.

A further more minor change concerns the spatial interpolation and averaging of T_{skin} when averaging the model fields for use in the lower-resolution radiation scheme. We require that if the surface properties of several adjacent gridboxes are merged then the surface upwelling longwave flux computed should be equal to the mean of the individual fluxes from the individual gridboxes. This can be achieved only if the averaging is performed on T_{skin}^4 , rather than T_{skin} as currently in the IFS.

3 Shortwave method

Unlike surface temperature, albedo is an almost static field between calls to the radiation scheme, so the modifications needed to update the shortwave net flux profile to respond to the local value of surface albedo need to be applied only once per radiation timestep. However, additional modifications are needed to better account for the variation of solar zenith angle between calls to the radiation scheme.

Section 3.1 describes how surface fluxes are modified in response to the changed surface albedo, including the effect of changed back-reflection from the atmosphere, and section 3.2 outlines how the four components of surface albedo are treated. Offline calculations on clear and cloudy profiles are used to test the predicted change to the net flux profile in section 3.3. In section 3.4, a correction to the model is described due to its current use of different values of solar zenith angle in different parts of the radiation calculation. Then in section 3.5, we describe the implementation of the [Manners et al. \(2009\)](#) scheme for better treating the variation of solar zenith angle between calls to the radiation scheme.

3.1 Accounting for back-reflection by the atmosphere

The simplest approach to correcting the surface net shortwave flux would be to assume that the surface downwelling flux $S_{\text{surf}}^{\downarrow}$ from the radiation scheme is correct and then compute a new upwelling flux as $S_{\text{surf}}^{\uparrow} = \alpha S_{\text{surf}}^{\downarrow}$, where α is the local value of the surface albedo. Hence the new net flux would be

$$S_{\text{surf}}^n = (1 - \alpha)S_{\text{surf}}^{\downarrow}. \quad (11)$$

However, this neglects the fact that the downwelling flux cannot be considered independent of the surface albedo; in reality a fraction of the enhanced reflection from the surface is scattered back down to the surface. To account for this, we treat the entire atmosphere as a single slab with a broadband transmittance τ and reflectance R such that the following relationships may be written between the surface and top-of-atmosphere (TOA) upwelling and downwelling fluxes:

$$S_{\text{surf}}^{\downarrow} = \tau S_{\text{TOA}}^{\downarrow} + R S_{\text{surf}}^{\uparrow}; \quad (12)$$

$$S_{\text{TOA}}^{\uparrow} = \tau S_{\text{surf}}^{\uparrow} + R S_{\text{TOA}}^{\downarrow}; \quad (13)$$

$$S_{\text{surf}}^{\uparrow} = \alpha S_{\text{surf}}^{\downarrow}. \quad (14)$$

Equation 12 states that the surface downwelling flux is the sum of transmission from the top-of-atmosphere downwelling flux and reflection of the surface upwelling flux by the atmosphere, while (13) states that the upwelling TOA flux is the sum of the transmission of the surface upwelling flux and reflection of the TOA downwelling flux by the atmosphere.

The idea is to use the boundary fluxes from the radiation scheme to compute τ and R , and then assume that they are independent of the surface albedo. Then the dependence of upwelling and downwelling surface fluxes on albedo can be computed. From (12)–(14) we find

$$\tau = \frac{S_{\text{surf}}^{\uparrow} S_{\text{TOA}}^{\downarrow} - S_{\text{surf}}^{\downarrow} S_{\text{TOA}}^{\uparrow}}{(S_{\text{TOA}}^{\downarrow})^2 - (S_{\text{surf}}^{\uparrow})^2}; \quad (15)$$

$$R = \frac{S_{\text{TOA}}^{\uparrow} S_{\text{TOA}}^{\downarrow} - S_{\text{surf}}^{\uparrow} S_{\text{surf}}^{\downarrow}}{(S_{\text{TOA}}^{\downarrow})^2 - (S_{\text{surf}}^{\uparrow})^2}. \quad (16)$$

It can also be shown from (12) and (14) that the surface net flux is given by

$$S_{\text{surf}}^n = S_{\text{TOA}}^{\downarrow} \frac{\tau(1 - \alpha)}{1 - \alpha R}. \quad (17)$$

This equation may then be used to compute a new surface net flux according to the local value of surface albedo.

Equation (17) is applied immediately after the radiation scheme is called to correct for spatial errors in albedo at the coarser radiation grid, but no attempt is made to account for the temporal changes of

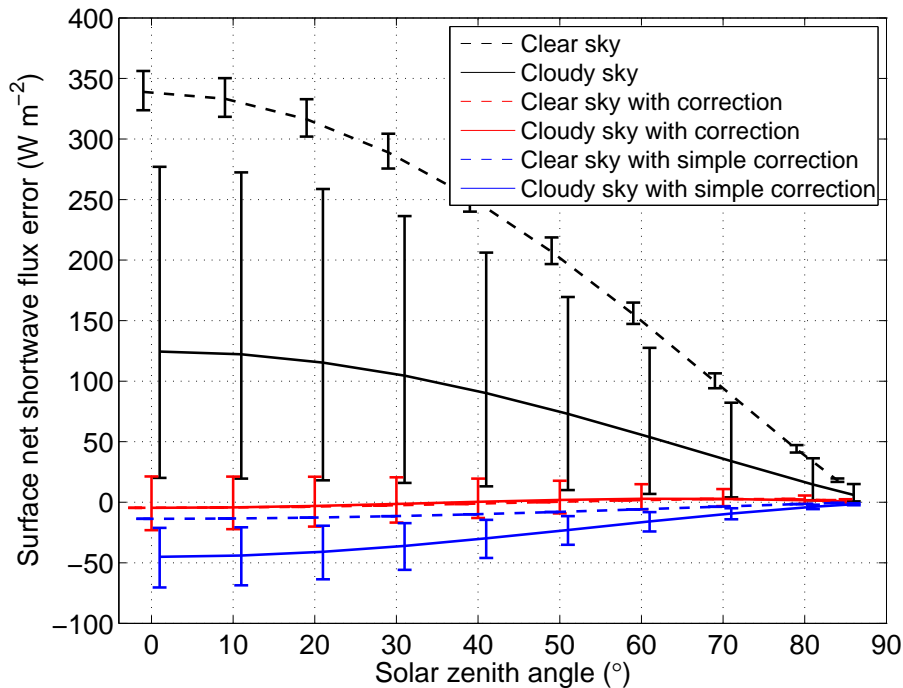


Figure 3: The error in surface net shortwave flux when a radiation calculation using a surface albedo of 0.08 is used over a surface with an albedo of 0.4, versus solar zenith angle, for all six standard atmospheres of McClatchey et al. (1972) and a range of cloud conditions. The solid lines depict the mean of the cloudy cases considered while the dashed lines depict the mean of the clear-sky cases. The error bars indicate the maximum and minimum errors for all conditions considered. The black lines depict the error with no attempt to make a correction for surface albedo, the red lines depict the correction proposed in this memorandum, and the blue lines depict a simpler correction where the downwelling shortwave flux is assumed to be constant with surface albedo.

surface albedo. Such variations are significant over the ocean via its dependence on solar zenith angle, and this relationship is represented in the model using the expression of Taylor et al. (1996). In principle, (17) could be applied every timestep with an ocean albedo updated according to the solar zenith angle, but in practice the large heat capacity of the ocean means that the resulting non-systematic errors in surface shortwave fluxes do not lead to significant differences in skin temperature, so this extra degree of complexity is not justified.

To test the new scheme, offline radiation calculations have been performed using RRTM-G with the 137 pressure levels of the IFS on the six standard atmospheres of McClatchey et al. (1972) for the full range of solar zenith angles and a variety of combinations of high/low and thick/thin cloud, as well as clear skies. In each case, calculations are performed with a surface albedo of 0.08, representing an ocean surface, and 0.4, representing a desert in the adjacent gridbox. The method described above is used to estimate what the net surface flux over the desert ought to be using the boundary fluxes from the calculation over the ocean.

The results are shown in Fig. 3. The black lines show the errors in net surface flux that would be made by the current version of the model, which makes no attempt to correct for the local value of surface flux. Unsurprisingly, the largest error is for an overhead sun in clear skies, where it is around 340 W m^{-2} . If we make the simple approximation that the downwelling shortwave flux is constant with surface albedo and use (11) to compute an updated net surface flux, then the resulting error is shown by the blue lines in Fig. 3. While a big improvement on the black lines, the neglect of back-reflection by the atmosphere,

particularly in cloudy situations, leads to systematic errors of up to almost 50 W m^{-2} . Finally, the red lines show that the reflectance-transmittance method described above provides an almost unbiased estimate of surface net flux for all solar zenith angles, regardless of whether there are clouds in the profile.

3.2 Computing the broadband albedo

If the albedo is unchanged from the value used in the radiation scheme then (17) will predict exactly the same surface net flux. However, care is required to ensure this is the case since the shortwave radiation scheme actually uses four separate albedo values: separate values for direct and diffuse radiation, and separate values in the ultraviolet/visible and near-infrared parts of the spectrum. Therefore, in order to use (17) at each gridpoint, we need to convert the four albedos at each gridpoint (denoted α_1 to α_4) into an equivalent broadband albedo α . This is done by modifying the shortwave radiation scheme to return not only the surface broadband downwelling flux S^\downarrow , but also the four components of the flux corresponding to direct and diffuse, and ultraviolet/visible and near infrared (denoted S_1^\downarrow to S_4^\downarrow). These components are interpolated from the radiation grid to the model grid and then used as follows to derive broadband albedo:

$$\alpha = \frac{\sum_{i=1}^4 \alpha_i S_i^\downarrow}{S^\downarrow}. \quad (18)$$

3.3 Updating the shortwave net flux profile

The method described so far appears to provide a good correction for the surface net flux, but does not provide guidance as to how to update the flux profile, and therefore the heating rate profile. The importance of scattering in the shortwave means that formally computing, say, the partial derivative of the net flux profile with respect to the surface net flux would be at least as computationally expensive as the original radiation code. The atmosphere is far more transparent to solar radiation than thermal infrared, especially when considering the solar radiation reflected from the surface, since by this point the radiation in the strongly absorbing parts of the spectrum has already been removed. Therefore we make the approximation that any excess upwelling solar radiation at the surface is lost to space, which means that the atmospheric heating rates are unchanged. (The same assumption was made by [Manners et al. \(2009\)](#) in a similar context; see section 3.5 below.) This is implemented by adding an offset to the net flux profile that is constant with height, such that the surface value matches that calculated by the reflectance-transmittance method described above.

To test the validity of this approximation, Figs. 4 and 5 show Tropical standard atmosphere scenarios, the first cloudy with a solar zenith angle of 40° and the second clear with an overhead sun, over surfaces with albedos of 0.08 (black lines) and 0.4 (blue lines). The red lines indicate the attempt to reproduce the $\alpha = 0.4$ profiles from the $\alpha = 0.08$ profile. In both scenarios, the surface net flux is closely reproduced, as are the surface upwelling and downwelling components. In the first scenario the modified net flux profile closely matches the true $\alpha = 0.4$ profile. This reflects the fact that in this case the heating rate profiles are very similar between the two surface albedos (Fig. 4b). In the second scenario the difference in heating rate profile (Fig. 5b) is more clearly evident at pressures higher than 800 hPa, which would not be captured by the approximate radiation update. However, in this example the boundary layer would be strongly convective, and so the sensible heat flux profile would likely adjust to produce a similar boundary-layer temperature profile.

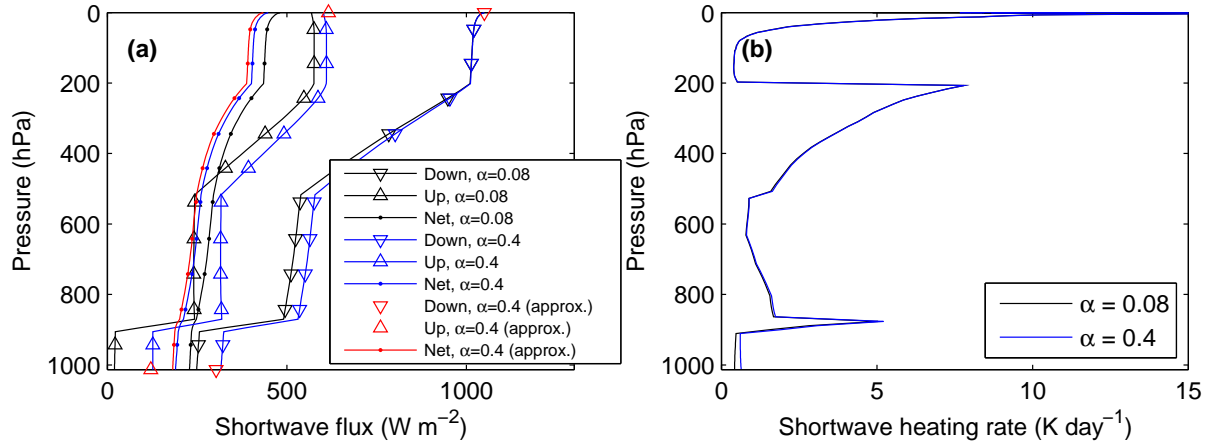


Figure 4: Evaluation of the approximate method to update the shortwave flux profile in response to a surface albedo change using the Tropical standard atmosphere with a solar zenith angle of $\theta_0 = 40^\circ$ and liquid and ice cloud layers as described in the caption of Fig. 1. The black lines depict the control scenario of an oceanic profile with a surface albedo of $\alpha = 0.08$. The blue lines depict the results for the same atmospheric profile but with a desert surface albedo of 0.4. The red symbols in panel a depict the approximate surface and top-of-atmosphere fluxes for a surface albedo of 0.4, estimated using the control scenario as input. The red line with dots depicts the approximate shortwave net flux profile. Since this is simply a shifted version of the control profile (black line with dots), the heating rate profile is unchanged and would be the same as the control profile, so is not shown in panel b.

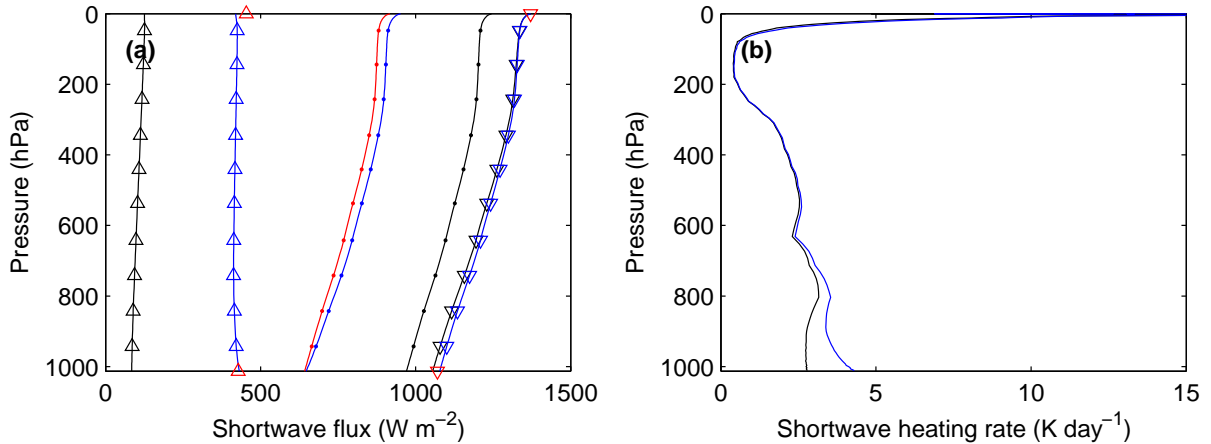


Figure 5: As Fig. 4 but for a cloud-free atmosphere with a solar zenith angle of 0° .

3.4 Computing solar zenith angle at the correct time

When the shortwave radiation scheme is called, the top-of-atmosphere incoming solar flux into a horizontal plane is set to unity, in order that the computed net shortwave flux profile is normalized. We refer to the time between calls to the radiation scheme as a *radiation* timestep. Then at every *model* timestep and gridbox, the incoming solar flux is computed from the local value of solar zenith angle θ_0 ; the normalized net flux profile is multiplied by the incoming solar flux to obtain the actual net flux profile, from which the local value of heating rate is computed (Morcrette, 2000). Thus, θ_0 is used three times:

1. Once every radiation timestep just before the radiation scheme is called, to compute the θ_0 -dependent albedo of the ocean surface using the empirical relationship of Taylor et al. (1996);

2. Once every radiation timestep in the radiation scheme itself, to compute the path-length of the direct solar beam through the atmosphere, and to compute the fraction of the solar beam that is scattered upwards and downwards;
3. At every model timestep and gridbox to compute the incoming solar flux into a horizontal plane at the top-of-atmosphere, in order to convert the normalized flux profile back into W m^{-2} .

In the third case, since the resulting fluxes are used to provide heating rates to evolve temperatures from model timesteps n to $n + 1$, θ_0 should be computed at the time corresponding to $n + 1/2$, and this is indeed done. In the first and second cases, θ_0 is supposed to be an average value for duration of a radiation timestep, i.e. the f model timesteps in which the output from the radiation scheme is used. Therefore it is calculated at a time corresponding to $f/2$ model timesteps in the future, i.e. model timestep $n + (1 + f)/2$. Unfortunately, if the radiation scheme is run every model timestep, i.e. $f = 1$, then this means θ_0 is computed for the *end* of the model timestep, $n + 1$, so will not be the same as the value computed in the third case. This has now been corrected so that in the first and second cases θ_0 is computed at a time corresponding to model timestep $n + f/2$.

3.5 Implementation of Manners et al. (2009) correction for solar zenith angle

As just described, the radiation scheme uses a solar zenith angle appropriate for a time half-way between calls to the radiation scheme. In the current ECMWF ensemble prediction system, the radiation scheme is called every 3 hours, which can lead to significant instantaneous errors due to the path-length of the direct beam through the atmosphere being computed at a time up to 1.5-h different from the times at which the fluxes are used. Manners et al. (2009) proposed a correction for this problem, which has been implemented in the IFS. In addition to computing normalized net flux profiles, the radiation scheme outputs the normalized downwelling direct surface solar flux, $\hat{S}_{\text{surf,dir}}^{\downarrow}$, which can be written as

$$\hat{S}_{\text{surf,dir}}^{\downarrow}(\mu_0) = \frac{S_{\text{surf,dir}}^{\downarrow}}{S_{\text{TOA}}^{\downarrow}} = \exp\left(-\frac{\delta}{\mu_0}\right), \quad (19)$$

where δ is the zenith broadband optical depth of the atmosphere and μ_0 is the cosine of the solar zenith angle used by the radiation scheme. Assuming that the atmospheric optical depth remains constant between calls to the radiation scheme, this quantity can be corrected to obtain normalized downwelling direct surface flux at another solar zenith angle (characterized by μ'_0) via

$$\hat{S}_{\text{surf,dir}}^{\downarrow}(\mu'_0) = \hat{S}_{\text{surf,dir}}^{\downarrow}(\mu_0)^{\mu_0/\mu'_0}. \quad (20)$$

We require the *total* (direct plus diffuse) normalized downwelling surface flux, at a changed solar zenith angle, i.e. $\hat{S}_{\text{surf}}^{\downarrow}(\mu'_0)$. Manners et al. (2009) found that this could be estimated quite well by assuming that if the direct flux at the surface is reduced by a certain amount, then half of this amount is scattered down to the surface in the form of diffuse radiation and half is scattered back to space. This leads to

$$\hat{S}_{\text{surf}}^{\downarrow}(\mu'_0) = \hat{S}_{\text{surf}}^{\downarrow}(\mu_0) + \frac{1}{2} \left[S_{\text{surf,dir}}^{\downarrow}(\mu_0)^{\mu_0/\mu'_0} - S_{\text{surf,dir}}^{\downarrow}(\mu_0) \right]. \quad (21)$$

Manners et al. (2009) assumed no change to atmospheric absorption, so it is necessary to simply add a constant offset to the entire net shortwave flux profile in order that the surface value is consistent with (21). This scheme has been added to the IFS and found to perform well (see section 4.2).

4 Case study

4.1 Correction of coastal errors

The approximate longwave and shortwave updates have been implemented in the IFS and in this section are tested in a case study for a period where the operational forecast had produced nighttime minimum 2-m temperature errors too cold by more than 10 K in the coastal gridpoints in the vicinity of Long Island and Connecticut, and indeed at the time, forecasters at La Guardia airport alerted ECMWF of the poor forecast. This error was associated with longwave errors, but the same forecast exhibited strong temperature overestimates around the coast of Arabia associated with shortwave errors. We run 3-day forecast experiments initialized at 12 Z on 3 January 2014 using IFS Cycle 40R2 at T_L1279 resolution in three configurations:

Control The default configuration in which the unmodified radiation scheme is called every hour and intermittently in space;

High-resolution radiation As the control except that the radiation scheme is run at every timestep and gridpoint, and the fix described in section 3.4 is applied to ensure that solar zenith angle is computed consistently;

Approximate radiation update As the control except that the procedures described in sections 2 and 3 are applied to provide approximate updates to the longwave and shortwave surface fluxes and, in the case of the longwave, the heating rate profile.

Figures 6 and 7 depict snapshots of the skin temperature from the three model configurations in two target regions that highlight the nighttime and daytime errors, respectively, in the default model configuration. The differences between Figs. 6a and 6b demonstrate how coastal nighttime land temperature can be substantially underestimated due purely to a longwave error: the net longwave flux over the sea is applied over the adjacent land. Figure 7 reveals a large daytime skin temperature overestimate at the coast of Oman. This error is due to a combination of a longwave effect where the lower sea temperature leads to a low upwelling longwave radiation that is incorrectly applied to the warmer land, and a shortwave effect where the lower sea albedo predicts too high a shortwave absorption when applied to the coastal desert immediately adjacent to the sea.

Figures 6c and 7c show that the use of an approximate radiation update produces a virtually identical skin temperature map to the simulation running the radiation scheme at all timesteps and gridpoints, but of course with a much smaller computational cost. To understand in more detail how the approximate update modifies the surface fluxes, Figs. 8 and 9 depict the timeseries of surface net fluxes, skin temperature and 2-m temperature for the full 72 hours of the forecasts, for the points indicated by the white circles in Figs. 6a and 7a where the greatest temperature errors were found.

Considering first Long Island, the control experiment for the night of 4 January shows a 2-m temperature underestimate of up to 10 K and a skin-temperature underestimate of up to 25 K, compared to the high-resolution radiation benchmark experiment. Figure 8a reveals that this is caused by an underestimate in net longwave radiation by at least 50 W m^{-2} for the first 36 hours of the forecast. This is associated with the interpolation of the output of the radiation scheme from a nearby sea point at a time when the skies were largely cloud-free. The approximate longwave update (red line) reduces this error to less than 10 W m^{-2} , which then reduces the temperature errors to less than 2 K. While the errors in this winter case are predominantly a longwave phenomenon, large shortwave errors are also present associated with

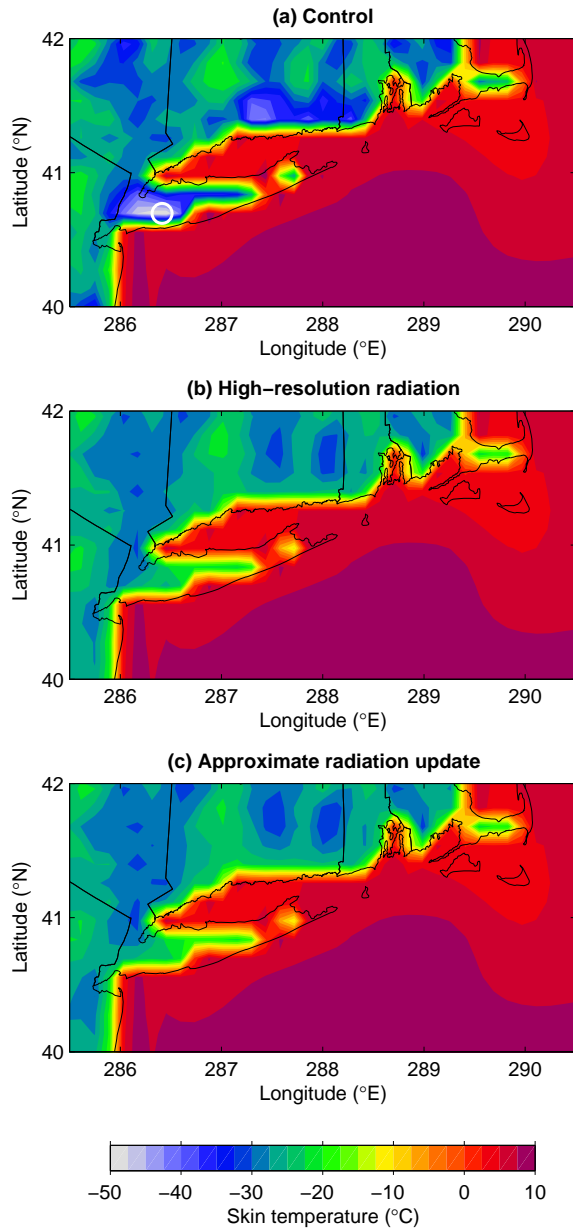


Figure 6: Skin temperature at 12 UTC (0700 local time) on 4 January 2014 for a region around the coast of Long Island and Connecticut from forecasts initialized 24 hours previously: (a) control T_{L1279} model, (b) the same model but with the radiation scheme run at every timestep and gridpoint, (c) as panel a but with the new scheme to perform an approximate update to the surface fluxes and the heating-rate profile. The time series of surface variables at the point indicated by the white circle is shown in Fig. 8.




Figure 7: As Fig. 6 from the same forecasts but at 10 UTC (1400 local time) on 4 January 2014 for a region around the southern coast of Oman. The time series of surface variables at the point indicated by the white circle is shown in Fig. 9.

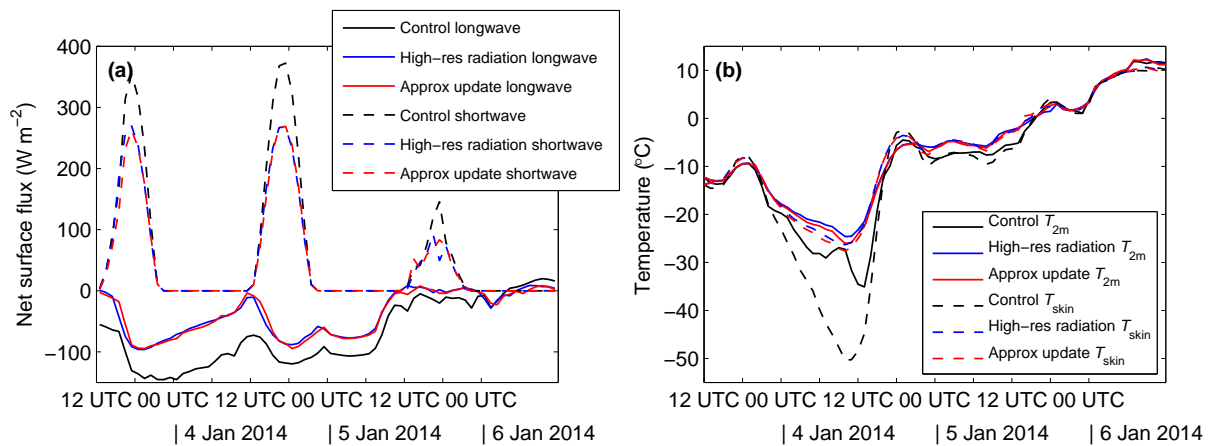


Figure 8: Time series of (a) surface net shortwave and longwave fluxes and (b) 2-m temperature and skin temperature at 40.7°N 286.4°E (the point on Long Island indicated by the white circle in Fig. 6a) for the full 72 hour forecasts using same three model configurations.

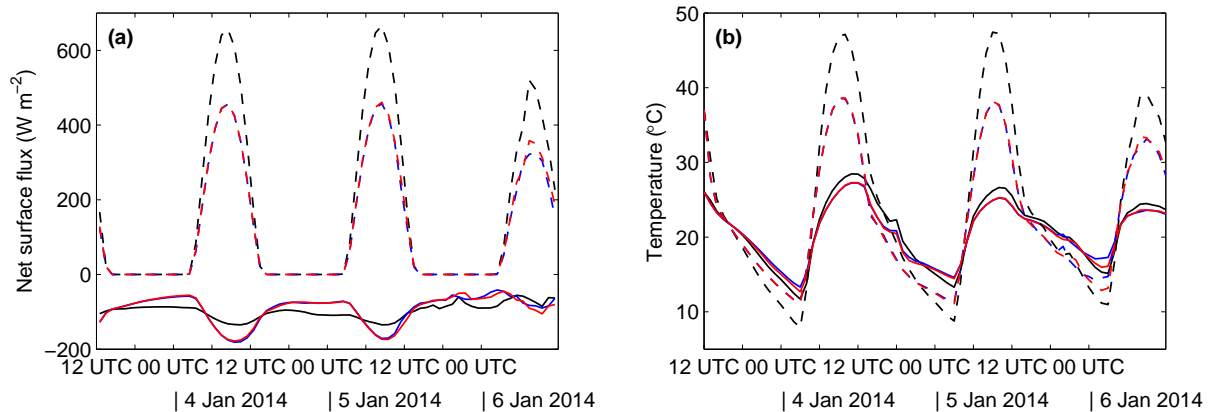


Figure 9: As Fig. 8 but for 18.9°N 57.1°E (the point on the coast of Oman indicated by the white circle in Fig. 7a).

the high albedo of the snow cover in Long Island. In the first two days of the forecast the approximate shortwave update provides an almost perfect correction for this, with the modest differences compared to the high-resolution benchmark on the third day being associated with differences in the cloud field in the two forecasts.

Considering second the coast of Oman, the main errors are strong daytime overestimates in skin temperature. These errors are primarily associated with an overestimate in net shortwave surface flux of up to 200 W m^{-2} , which is again almost perfectly corrected by the approximate shortwave update. However, the longwave fluxes also show significant errors, with the timeseries from the control experiment having much less diurnal variation, typical for an ocean surface with much less variation in surface temperature. In the daytime the longwave error is of the same sign as the shortwave error, although of lower magnitude, so contributes to the temperature overestimate. The approximate longwave update is again able to closely match the fluxes from the high resolution radiation simulation, and as a consequence the temperature timeseries in Fig. 9b is closely reproduced.

Finally, we address the question raised in section 2.2 of whether the large atmospheric heating rates that can arise from the approximate longwave update are countered by other parts of the physics package, especially the turbulent transport. Figure 10 shows the near-surface profiles of temperature and temper-

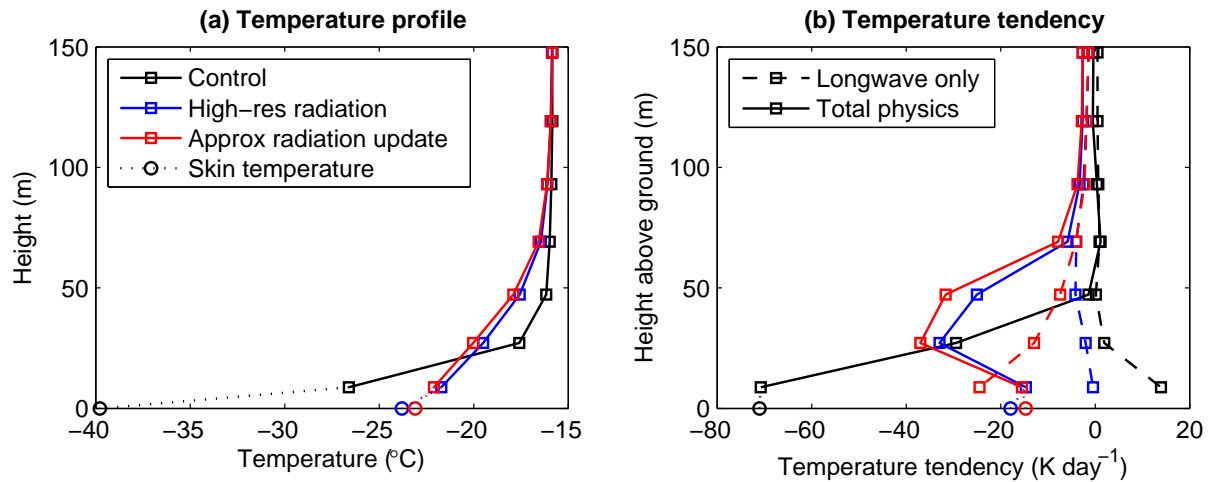


Figure 10: Profiles of (a) temperature and (b) temperature tendency at 07 UTC (0200 local time) on 14 January 2014 for the three model configurations over the point on Long Island shown in Figs. 6a and 8. The skin temperature and its tendency is shown by circles. In panel b the temperature tendency is shown for the entire physics package as well as for longwave radiation only.

ature tendency at the time when the Long Island surface temperature was cooling most rapidly. Panel a shows that the approximate longwave update is able to closely reproduce the temperature profile from the high-resolution radiation experiment. The dashed lines in panel b reveal substantial differences in the atmospheric heating rates, with the approximate radiation update predicting cooling rates as low as 20 K day^{-1} in the lowest model level compared to near-zero in the high-resolution radiation experiment. However, the temperature tendency associated with all physical processes reveals much better agreement between these two experiments. The control experiment, on the other hand, shows a much stronger near-surface cooling by the physics package in excess of 70 K day^{-1} , even though the longwave tendency is for a warming. The explanation for this behaviour is that the near-surface temperatures are strongly coupled to the surface by turbulent transport. Since the approximate longwave update produces a good estimate of the surface temperature tendency and hence a good surface temperature forecast (the red circles in Fig. 10), the turbulent transport simply adjusts to give approximately the same total temperature tendency, even though the longwave heating rates are different. Thus we can conclude that even though the approximate longwave update accounts only for differences in surface temperature and not air temperature, errors in longwave heating rates are mitigated by other parts of the physics package.

4.2 Correction of the diurnal cycle of surface temperature

Section 4.1 demonstrated the performance of approximate updates in correcting spatial errors due to the radiation scheme being called intermittently in space. Here we demonstrate that these updates can also correct for temporal errors associated with the radiation scheme being called infrequently in time. These errors are particularly apparent for the model configuration used in ensemble forecasting, in which the radiation scheme is called only every 3 hours. We have replicated the ensemble model configuration (T_L639 resolution, but with radiation at a spatial resolution of T_L255) for the same case as studied in section 4.1. Errors in surface fluxes and temperature due to the infrequent radiation calls are most obvious over deserts, so we have chosen a point over a desert region of Western Australia; note that in early January the sun passes only a few degrees from zenith at this point, and it is also the time of year when the sun-earth distance is shortest. The evolution of net surface fluxes over 48 hours is shown in

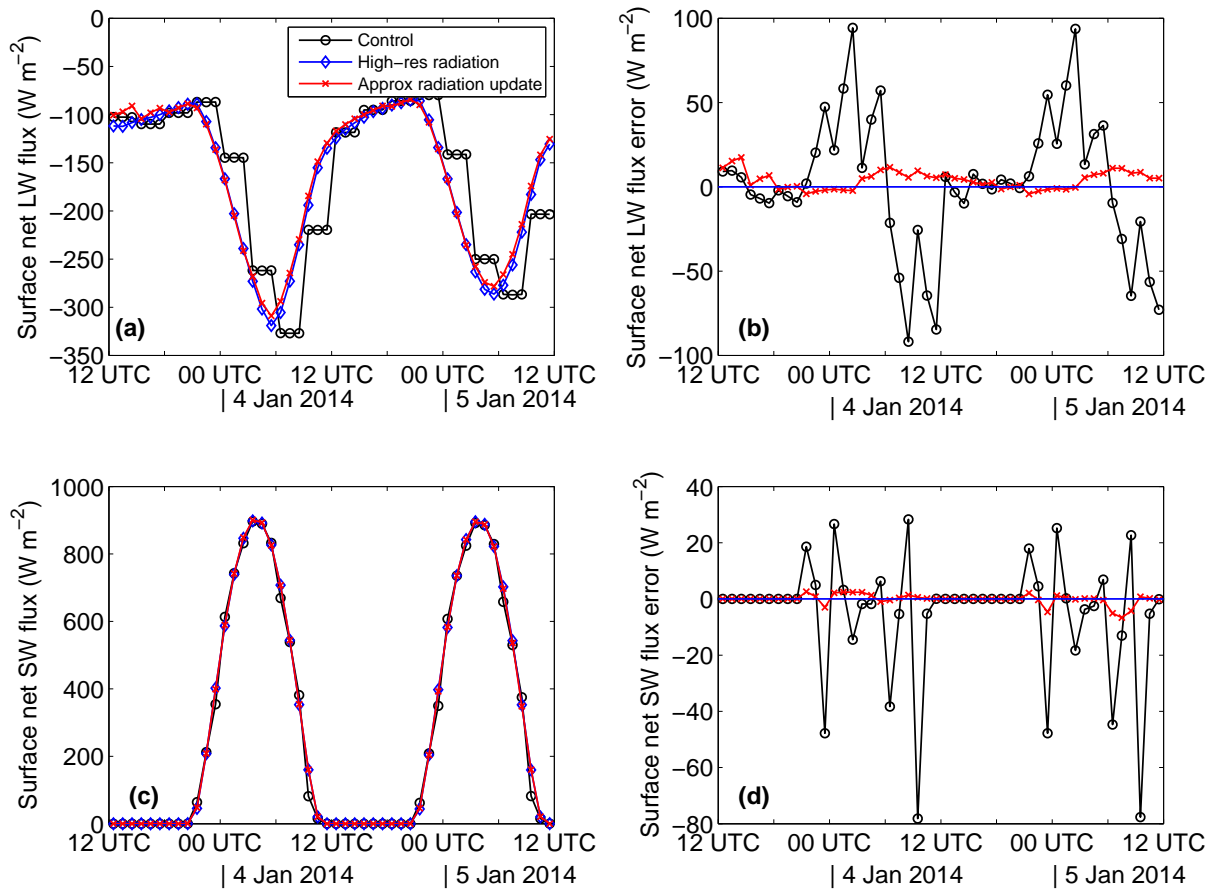


Figure 11: Time series of surface net fluxes at 26.25°S , 123.25°E (a point in the Australian Desert) for the same forecast period as in Figs. 6–10, but for a model resolution of $T_L 639$ and a default radiation timestep of 3 hrs. (a) Longwave fluxes from the control model configuration, the “high-resolution radiation” simulation with the radiation scheme run at every gridbox and timestep, and the new approximate radiation update. (b) Longwave flux minus the values from the high-resolution radiation simulation, treated as truth. (c)–(d) As (a)–(b) but for shortwave radiation.

Fig. 11.

The black line in Fig. 11a demonstrates the limitation of the current longwave scheme: surface net longwave flux, L_{surf}^n , is assumed constant between calls to the radiation scheme, leading to 3-h “steps”. Comparing to the same model configuration but with the radiation scheme run every model timestep and gridbox, we see in Fig. 11b that this leads to instantaneous errors of up to almost 100 W m^{-2} . The red line shows that a 3-h radiation timestep but with approximate updates to the radiative fluxes every model timestep reduces this error to typically no more than 10 W m^{-2} . Bozzo et al. (2015) showed a similar comparison of a diurnal cycle of L_{surf}^n over the Nevada desert from model versions with different radiation timesteps, and found that calling the radiation scheme every model timestep led to the best agreement with observed surface fluxes.

The black line in Fig. 11d shows that the current shortwave scheme has similar errors, with hourly fluxes in error by up to almost 80 W m^{-2} . These fluctuating errors are due to the fact that the solar zenith angle used in the radiation scheme to compute the path-length of the direct solar beam through the atmosphere is assumed constant for three hours. The red line shows that these errors are largely removed via the scheme of Manners et al. (2009) as discussed in section 3.5.

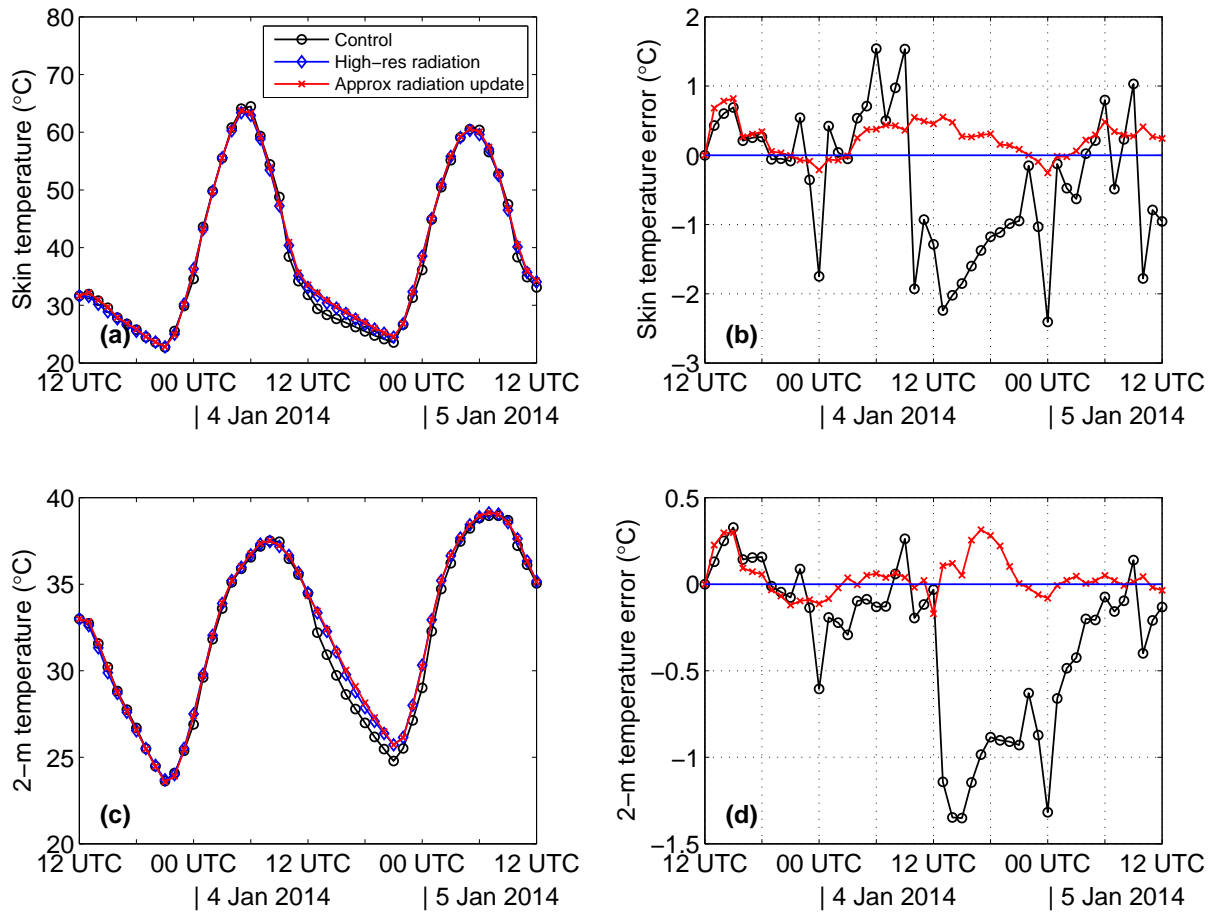


Figure 12: As Fig. 11 but for (a)–(b) skin temperature and (c)–(d) 2-m temperature.

Figure 12 confirms that the much improved fluxes depicted in Fig. 11 lead to a much improved forecast of skin temperature and 2-m temperature. The largest and most persistent error in the control model (black lines) appears to be at night when skin temperature is underestimated 1–2 K. This arises due to the longwave net flux being updated only every 3 hours, which means that radiation emitted from the surface does not reduce correctly in response to the falling surface temperature, so the radiative cooling is too rapid. In the daytime, temperatures are too low in the morning but typically too high by sunset. This is associated with the longwave emission not increasing correctly in response to the increasing surface temperature. It can be seen from the red lines that approximate updates provide much improved temperature forecasts.

One puzzling aspect is the immediate evolution of a positive temperature error in both the control and approximate-update experiments in the first three hours of the forecast. Indeed, in the first hour the skin temperature increases, even though the sun is below the horizon. This is believed to be a spin-up effect, specifically that skin temperature in this T_L639 forecast is initialized from a T_L1279 analysis with different orography, and is not in fact a value that exactly balances the terms in the surface energy balance equation. It takes at least the time between calls to the radiation scheme to recover from this effect.

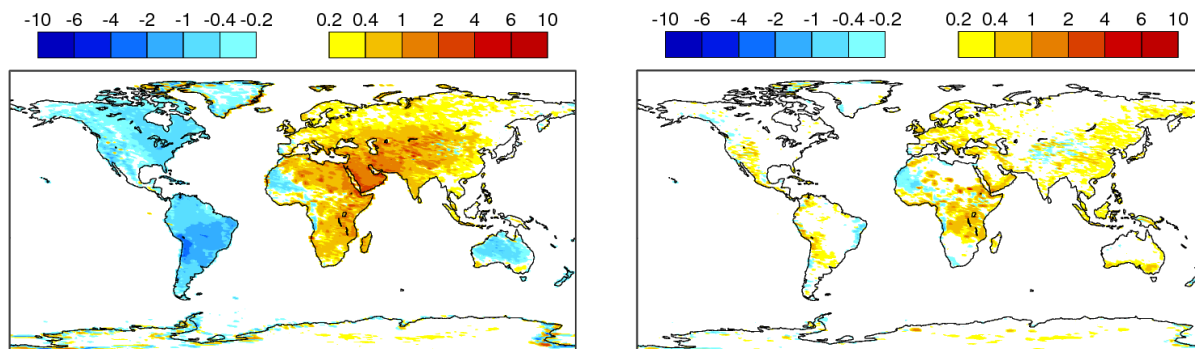


Figure 13: (Left) Difference in mean skin temperature (K) at 12 UTC between T_L639 forecasts with the radiation scheme called every 3 hours and intermittently in space, and forecasts with the radiation scheme called every timestep and gridbox. (Right) The same but with approximate updates to the radiation fields every timestep and gridpoint. One forecast is initialized from the analysis at 00 UTC every day for four months between 1 June and 30 September 2012, and skin temperature values 36 h into each forecast have been extracted. The approximate update assumes a longwave downwelling factor of $\gamma = 0.2$.

5 Long-term evaluation of new scheme

For a more rigorous evaluation of the performance of the scheme, we perform a set of 5-day forecasts initialized from the operational ECMWF analysis at midnight each day for four Northern Hemisphere summer months (June to September 2012) and four winter months (December 2012 to March 2013). This configuration matches exactly that used by [Bozzo et al. \(2015\)](#), except that here we use model cycle 40R2. Two model resolutions are used matching those used in section 4: T_L1279 with the radiation scheme called by default every 1 h, and T_L639 with the radiation scheme called by default every 3 h. The results are analysed in terms of changes to the diurnal cycle of surface temperature (section 5.1), evaluation against 2-m temperature measurements at European coastal stations (section 5.2) and forecast skill as a function of time into the forecast (section 5.3).

5.1 Diurnal cycle of surface temperature

Errors in the diurnal cycle of surface temperature are only really concerning in the current ensemble configuration of the model at T_L639 , where the radiation scheme is called only every 3 h, so we restrict our analysis in this section to this resolution. Following the approach of [Bozzo et al. \(2015\)](#) and [Sandu et al. \(2014\)](#), the left panel of Fig. 13 depicts the difference in mean skin temperature at 12 UTC (36 h into the forecast) between the control version of the model and a version with the radiation scheme called every timestep and gridpoint (treated as “truth” for the purposes of this study), for the June–September period. It can be seen that temperatures are underestimated in land regions before local noon (west of the Greenwich meridian) and overestimated after local noon. This can be explained by the longwave mechanisms discussed in section 4.2: assuming the surface net longwave fluxes are constant for 3 h leads to a lag in the surface cooling rate. In particular, during the evening the cooling rate is too strong, leading to an underestimate in the night-time minimum and morning temperatures. When the surface starts to heat up during the day, the longwave lag means that the longwave cooling rate is not as strong as it ought to be and the surface warming is too rapid. By afternoon the temperatures are then overestimated.

The right panel of Fig. 13 shows the same but for forecasts using approximate radiation updates, and it can be seen that the errors are significantly reduced. The remaining errors are largely in regions of

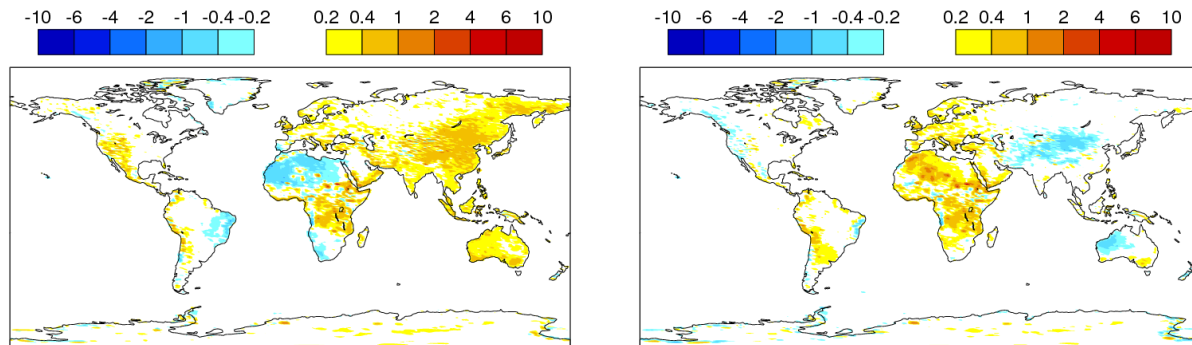


Figure 14: (Left) As the right panel of Fig. 13 but for approximate updates using a longwave downwelling factor of 0. (Right) The same but using a longwave downwelling factor of 0.4.

orography, particularly in Africa. It is also worth noting that coastal temperature errors are apparent in the control simulation (left panel), particularly on the east coast of Greenland and around the coast of Arabia, but these are largely removed in the right panel.

In this case, approximate radiation updates have been applied with a longwave downwelling factor of $\gamma = 0.2$ (described in section 2.3). Figure 14 shows the result of applying approximate radiation updates with values of 0 and 0.4, and it can be seen that while the diurnal cycle errors are still better than the control simulation, they are worse than 0.2. Therefore we recommend $\gamma = 0.2$ for future use.

5.2 Comparison of coastal temperatures to observations

Using exactly the same approach as Bozzo et al. (2015), we have selected European coastal observing stations whose closest model gridpoint in the control version of the model uses a nearby sea point in the radiation scheme, and therefore the points where we would expect the largest forecast errors due to incorrect surface temperature and albedo. Note that the stations selected are different depending on the resolution of the model, but for each resolution and season analysed, between 55 and 95 stations are available.

Figure 15 shows the probability distribution of 2-m temperature error for the two model resolutions at midday in summer months and midnight in winter months. The distribution of errors in the control model (black lines) is significantly skewed in each case, with a tail of large positive temperature errors in summer at midday, and a tail of large negative temperature errors in winter at midnight. As found by Bozzo et al. (2015), this is much improved when the radiation scheme is called every timestep and gridpoint (blue lines), as indicated also by the improvement in both bias and root-mean-squared error indicated in the legends of Fig. 15. The results from using approximate radiation updates (red lines) are generally quite similar to calling radiation every timestep and gridpoint, confirming observationally that it leads to significantly better forecasts. Evaluation of midnight forecasts in summer and midday forecasts in winter leads to much less difference between any of the lines (not shown).

Figure 16 compares these simulations to 2-m temperature observations at a Norwegian coastal site during December 2012. There is a significant cold bias in the control model configuration, with maximum differences in excess of 10 K towards the end of the month. This bias is largely removed both in forecasts with the radiation scheme run every timestep and gridpoint, and at much lower computational expense in forecasts using approximate radiation updates.

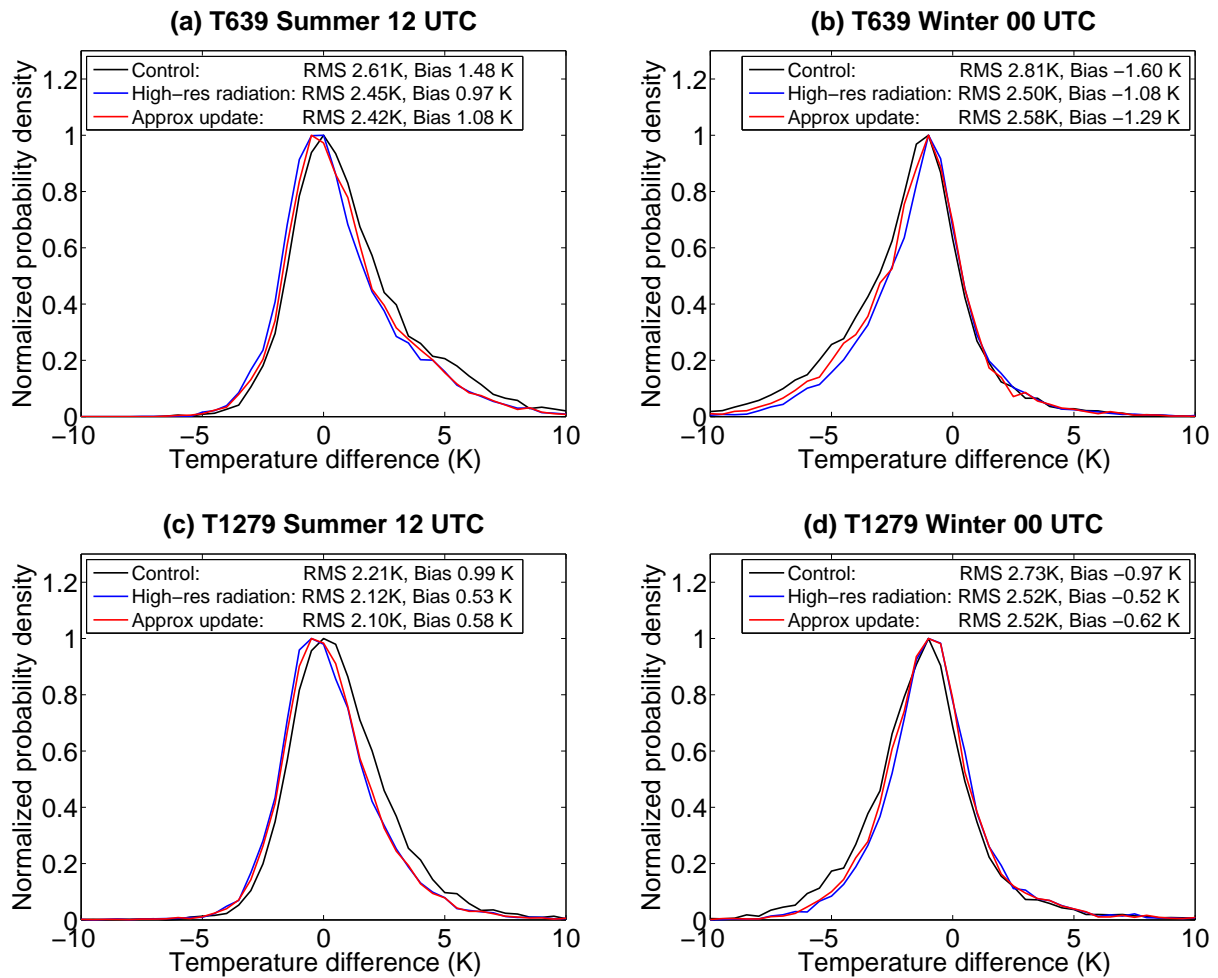


Figure 15: Probability density functions (normalized to the maximum value) of the difference between forecast and observed 2-m temperature at European coastal stations. The left column shows results for June to September 2012 forecasts at midday, 36 h into each daily forecast, while the right column shows results for December 2012 to March 2013 forecasts at midnight, 24 h into each daily forecast. The top row shows results for the T_L639 resolution model and the bottom row the T_L1279 model.

5.3 Forecast skill

In this section we assess whether the improvement in the interaction between radiation and the surface has led to better forecasts. Since the improved physics is almost exclusively over land, we focus on the Northern Hemisphere extra-tropics. Figure 17 depicts the root-mean-square error in 1000-hPa temperature forecasts, separately for the two 4-month periods, assessed by comparing against the operational analyses. At T_L639 , calling radiation every timestep and gridpoint leads to a significantly better forecast in both summer and winter, although the difference is larger in summer. In both cases, the use of approximate radiation updates leads to errors that lie around half-way between the control and high-resolution radiation simulations. This suggests that the fast interaction of clouds and radiation, which the new scheme does not improve, is also important for improving forecasts.

The difference between the various T_L1279 simulations is very small in terms of temperature. This is believed to be because the control model configuration calls the radiation scheme every hour rather than every 3 h, which is the main factor degrading the performance of the T_L639 control forecasts.




Figure 16: Comparison of midnight and midday 2-m temperature forecasts (at lead times of 24 h and 36 h, respectively) against observations at Sortland, Norway (68.7°N, 15.42°E) for December 2012. The T_{L639} model simulations were the same as used to produce Fig. 15b.

Figure 18 shows the root-mean-square error in 500-hPa geopotential-height forecasts for the same period as used in Fig. 17. The only case where there is a significant improvement against the control forecasts is in summertime T_{L639} simulations with high resolution radiation (i.e. calling the radiation scheme every timestep and gridpoint). In wintertime and T_{L1279} simulations, none of the changes to the radiation scheme make a significant change to forecast skill.

6 Conclusions

In this memorandum, modifications to the IFS have been described that mitigate some of the problems associated with its calling of the radiation scheme intermittently in time and space: approximate updates are performed to radiative fluxes every timestep and gridpoint. This allows fluxes to respond correctly both to rapid spatial variation in surface albedo and temperature (particularly at coastlines), and to temporal evolution of surface temperature between calls to the radiation scheme. The two principal new ideas are (1) to store the partial derivative of upwelling longwave flux with respect to surface upwelling flux in order that an increase or decrease of the emission from the surface correctly leads to an increase or decrease in atmospheric absorption, and (2) to recompute surface shortwave fluxes in response to a changed surface albedo in a way that accounts for back-reflection from the atmosphere. We also implement the Manners et al. (2009) method to account for the change in direct solar path length through the atmosphere as solar zenith angle changes between calls to the radiation scheme.

We find that these changes lead to significant improvements in 2-m temperature forecasts at coastal land points, as well as improving a lag in the diurnal cycle of surface temperature. The Manners et al. (2009) scheme significantly reduces the random errors in surface shortwave downwelling flux. The computational cost of performing these updates is only around 2% of the cost of the radiation scheme itself.




Figure 17: Comparison of the root-mean-squared error of daily 1000-hPa temperature forecasts in the Northern Hemisphere extra-tropics, versus lead time, for various model configurations for the periods (top panel) June to September 2012 and (bottom panel) December 2012 to March 2013.




Figure 18: As Fig. 17 but for 500-hPa geopotential height in the Northern Hemisphere extra-tropics.

Bozzo et al. (2015) reported that similar improvements could be obtained by calling the radiation scheme more frequently but with only a subset of the spectral g points each call, an idea originally proposed in the context of large-eddy modelling by Pincus and Stevens (2009). However, spectral sampling leads to a large noise in the instantaneous fluxes; although this averages out completely after a few hours, it makes instantaneous radiative fluxes output from the model much more difficult to interpret by forecast users (such as the solar energy industry in this case). The approximate radiation updates described in this memorandum do not suffer from this problem.

While approximate updates lead to better surface temperature forecasts, there is no evidence that they improve the skill of 500-hPa geopotential height forecasts, even though running the radiation scheme every timestep rather than every 3-h does produce better height forecasts. This is believed to be because the new scheme responds to changes in surface properties but not clouds. A future research topic will be to investigate whether the full radiation scheme can be made more efficient by reducing the number of spectral g points, thereby allowing it to be run more frequently in time and space. We note that the current RRTM-G scheme uses 224 points across the shortwave and longwave spectrum, whereas fewer than half this number have been found sufficiently accurate for other successful correlated- k schemes (Fu and Liou, 1992; Cusack et al., 1999), and there is the potential for even fewer using the full-spectrum correlated k method (Pawlak et al., 2004; Hogan, 2010). Speed-up of the radiation scheme may also be possible by running radiation in parallel to the rest of the model, either within the current CPU framework (Mozdzynski and Morcrette, 2014) or using GPUs (Price et al., 2014).

Acknowledgements

We are grateful to Irina Sandu, Anton Beljaars, Linus Magnusson, Tim Hewson and James Manners (Met Office) for useful discussions. Robert Pincus (University of Colorado) is thanked for providing his offline radiation code that was used in the production of Figs. 1–5.

References

- Bozzo, A., R. Pincus, I. Sandu and J.-J. Morcrette, 2015: Impact of a spectral sampling technique for radiation on ECMWF weather forecasts. *J. Adv. Modeling Earth Sys.*, **6**, doi:10.1002/2014MS000386.
- Clough, S. A., M. W. Shephard, E. J. Mlawer, J. S. Delamere, M. J. Iacono, K. Cady-Pereira, S. Bouk-abara, P. D. Brown, 2005: Atmospheric radiative transfer modeling: a summary of the AER codes. *J. Quant. Spectrosc. Radiat. Transfer*, **91**, 233–244.
- Cusack, S., J. M. Edwards and J. M. Crowther, 1999: Investigating k distribution methods for parameterizing gaseous absorption in the Hadley Centre Climate Model. *J. Geophys. Res.*, **104**, 2051–2057.
- Fu, Q., and K.-N. Liou, 1992: On the correlated k -distribution method for radiative transfer in nonhomogeneous atmospheres. *J. Atmos. Sci.*, **49**, 2139–2156.
- Hogan, R. J., 2010: The full-spectrum correlated- k method for longwave atmospheric radiation using an effective Planck function. *J. Atmos. Sci.*, **67**, 2086–2100.
- Lacis, A., and V. Oinas, 1991: A description of the correlated k -distribution method for modeling nongray gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres. *J. Geophys. Res.*, **96**, 9027–9063.

- Manners, J., J.-C. Thelen, J. Petch, P. Hill and J. M. Edwards, 2009: Two fast radiative transfer methods to improve the temporal sampling of clouds in numerical weather prediction and climate models. *Q. J. R. Meteorol. Soc.*, **135**, 457–468.
- McClatchey, R. A., R. W. Fenn, J. E. A. Selby, F. E. Volz and J. S. Garing, 1972: *Optical properties of the atmosphere* (3rd ed.), Air Force Cambridge Research Laboratories, Rep. No. AFCRL72-0497, L. G. Hanscom Field.
- Mlawer, E. J., S. J. Taubman, P. D. Brown, M. J. Iacono, and S. A. Clough, 1997: Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave. *J. Geophys. Res.*, **102**, 16 663–16 682.
- Morcrette, J.-J., 2000: On the effects of the temporal and spatial sampling of radiation fields on the ECMWF forecasts and analyses. *Mon. Weath. Rev.*, **128**, 876–887.
- Morcrette, J.-J., E. J. Mlawer, M. J. Iacono and S. A. Clough, 2001: Impact of the radiation transfer scheme RRTM in the ECMWF forecasting system. *ECMWF Newsletter*, No. 91, Reading, United Kingdom, 2–9.
- Morcrette, J.-J., H. W. Barker, J. N. S. Cole, M. J. Iacono, R. Pincus, 2008: Impact of a New Radiation Package, McRad, in the ECMWF Integrated Forecasting System. *Mon. Weath. Rev.*, **136**, 4773–4798.
- Morcrette, J.-J., G. Mozdzyński and M. Leutbecher, 2008: A reduced radiation grid for the ECMWF Integrated Forecasting System. *Mon. Weath. Rev.*, **136**, 4760–4772.
- Mozdzyński, G., and J.-J. Morcrette, 2014: Reorganization of the radiation transfer calculations in the ECMWF IFS. *Tech. Memo. 721*, ECMWF, Reading, UK, 20 pp.
- Pawlak, D. T., E. E. Clothiaux, M. F. Modest and J. N. S. Cole, 2004: Full-spectrum correlated-*k* distribution for shortwave atmospheric radiative transfer. *J. Atmos. Sci.*, **61**, 2588–2601.
- Pincus, R., and B. Stevens, 2009: Monte Carlo spectral integration: a consistent approximation for radiative transfer in large eddy simulations. *J. Adv. Modeling Earth Sys.*, **1**, doi:10.3894/JAMES.2009.1.1.
- Price, E., J. Mielikainen, M. Huang, B. Huang, H.-L. A. Huang and T. Lee, 2014: GPU-accelerated longwave radiation scheme of the Rapid Radiative Transfer Model for General Circulation Models (RRTMG). *IEEE J. Sel. Topics Appl. Earth Obs. Remote Sens.*, **7**, 3660–3667.
- Sandu, I., N. Wedi, A. Bozzo, P. Bechtold, A. Beljaars and M. Leutbecher, 2014: On the near surface temperature differences between HRES and CTL ENS forecasts. *ECMWF Technical Report*, available internally at ECMWF from <http://intra.ecmwf.int/publications/library/do/references/show?id=1493>.
- Schuster, A., 1905: Radiation through a foggy atmosphere. *Astrophys. J.*, **21**, 1–22.
- Taylor, J. P., J. M. Edwards, M. D. Glew, P. Hignett and A. Slingo, 1996: Studies with a flexible new radiation code – 2. Comparisons with aircraft shortwave observations. *Q. J. R. Meteorol. Soc.*, **122**, 839–861.