AGU PUBLICATIONS

Journal of Advances in Modeling Earth Systems

RESEARCH ARTICLE

10.1002/2015MS000455

Mitigating errors in surface temperature forecasts using approximate radiation updates

Robin J. Hogan[1] and Alessio Bozzo[1]

[1]European Centre for Medium-Range Weather Forecasts, Reading, UK

Key Points:
- Weather models can have large temperature errors at coastlines
- A new scheme is proposed to correct for this via approximate radiation updates
- The computational cost is only of order 0.2% of the cost of the entire model

Correspondence to:
R. J. Hogan,
r.j.hogan@ecmwf.int

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

Received 17 MAR 2015
Accepted 20 MAY 2015
Accepted article online 24 MAY 2015
Published online 11 JUN 2015

Abstract Due to computational expense, the radiation schemes in many weather and climate models are called infrequently in time and/or on a reduced spatial grid. The former can lead to a lag in the diurnal cycle of surface temperature, while the latter 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. This paper describes a computationally efficient solution to these problems, in which the surface longwave and shortwave fluxes are updated every time step and grid point 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 grid point and time step, the upwelling flux and heating-rate profiles are updated using the new value of skin temperature. The computational cost of performing approximate radiation updates in the ECMWF model 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 8 month period reveals significant improvement in 2 m temperature forecasts at coastal stations compared to observations.

1. Introduction

Despite the use of key approximations such as the two-stream approximation [*Schuster*, 1905] and the correlated-*k* distribution method [*Lacis and Oinas*, 1991], the radiation scheme is one of the most computationally expensive parts of a weather or climate model and is generally too expensive to run every time step and grid point. *Morcrette et al.* [2008b] provided a history of the reduced radiation resolution in time and space at ECMWF, noting particularly the introduction by *Morcrette et al.* [2008a] of the Rapid Radiative Transfer Model for GCMs (RRTM-G) in 2007. 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. Current operational practice at ECMWF is to run the high-resolution model at a spectral resolution of T_L 1279 (around 16 km) with the radiation scheme run every 1 h (every 6 model time steps) at an effective resolution of T_L 511 (around 40 km), and to run the ensemble prediction system at a resolution of T_L 639 (around 32 km) with the radiation scheme run every 3 h (every 9 model time steps) at a resolution of T_L 255 (around 80 km). Thus, in both cases, the radiation scheme is run on 6.25 times fewer grid points than the rest of the model physics.

Morcrette [2000] examined the impact of temporal and spatial sampling of radiation on weather forecasts, and found: (a) negligible degradation of forecast skill in terms of 500 hPa geopotential in the first 7 days of the forecast, but biases emerging in seasonal forecasts, (b) locally significant changes to skin temperature and cloudiness, and (c) that the weakening of the coupling between rapidly varying cloud fields and the radiation field led to a change to the model's climate sensitivity. More recently at ECMWF, it has become apparent that forecasts of nighttime costal 2 m temperature can sometimes be too cold by in excess of 10 K. This is due to the averaging of the surface temperature field onto the coarser radiation grid before being input to the radiation scheme, leading to the upwelling longwave radiation output at coastal land points being overestimated due to contamination by warmer sea points in the vicinity. Excessive surface cooling then follows.

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

© 2015. The Authors.
This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

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 ECMWF model 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 time step and grid point 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 the longwave, *Morcrette* [2000] reported that an older version of the ECMWF model normalized the net flux at each half-level by T^4, where T is the half-level temperature, and then rescaled it by T^4 using the most recent value of T every model time step, but this scheme was discontinued at least 20 years ago.

In this paper, we propose two new methods for approximately updating the radiation fields between calls to the full radiation scheme, and examine their effects on weather 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 meters 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 8 months of daily 5 day forecasts are run with different model configurations to assess the improvement to the forecasts.

2. Longwave Method

The modifications to the numerical weather forecast model needed to provide an approximate update of the longwave net fluxes (surface and atmosphere) at every time step and grid point are in two parts. First, the radiation scheme is modified to output 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. Second, these extra variables are used to update the net fluxes every time step and grid point as described in section 2.3.

2.1. Extra Variables From the Radiation Scheme

The only output from the longwave radiation scheme that is used by the rest of the model to compute surface and atmospheric heating rates is the profile of net longwave flux at each model half-level including the surface, $L^n_{i-1/2}=L^{\downarrow}_{i-1/2}-L^{\uparrow}_{i-1/2}$, where i is the vertical layer index counting down from 1 at the top. The new method requires two additional outputs: (1) the surface downwelling flux $L^{\downarrow}_{\mathrm{surf}}$ already available from the scheme and (2) the partial derivative of upwelling longwave flux at all model half-levels with respect to the surface upwelling longwave flux, i.e., $\partial L^{\uparrow}_{i-1/2}/\partial L^{\uparrow}_{\mathrm{surf}}$. This is a partial derivative in the sense that we are treating the atmospheric temperature and composition constant.

We modify the ECMWF radiation scheme to compute the profile of $\partial L^{\uparrow}_{i-1/2}/\partial L^{\uparrow}_{\mathrm{surf}}$ 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

$$L^{\uparrow}_{i-1/2}=\varepsilon\sum_{g=1}^{n_g} B_g(T_{\mathrm{skin}})\tau_{i\ldots n,g}+f_g, \tag{1}$$

where ε is the surface emissivity (here assumed constant across the longwave spectrum), $B_g(T_{\mathrm{skin}})$ is the Planck function (as a flux in W m^{-2}) at skin temperature T_{skin} 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 $\tau_{i\ldots n,g}$ 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:

$$\tau_{i\ldots n,g}=\prod_{k=i}^{n}\tau_{k,g}. \tag{2}$$

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

$$\frac{\partial L^{\uparrow}_{i-1/2}}{\partial T_{\mathrm{skin}}}=\varepsilon\sum_{g=1}^{n_g}\frac{\partial B_g(T_{\mathrm{skin}})}{\partial T_{\mathrm{skin}}}\tau_{i\ldots n,g}. \tag{3}$$

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

$$L^{\uparrow}_{\mathrm{surf}}=\varepsilon\sigma T^4_{\mathrm{skin}}+(1-\varepsilon)L^{\downarrow}_{\mathrm{surf}}, \tag{4}$$

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

$$\frac{\partial L^{\uparrow}_{\mathrm{surf}}}{\partial T_{\mathrm{skin}}}=4\varepsilon\sigma T^3_{\mathrm{skin}}. \tag{5}$$

Combining with (3) yields

$$\frac{\partial L^{\uparrow}_{i-1/2}}{\partial L^{\uparrow}_{\mathrm{surf}}}=\frac{1}{4\sigma T^3_{\mathrm{skin}}}\sum_{g=1}^{n_g}\frac{\partial B_g(T_{\mathrm{skin}})}{\partial T_{\mathrm{skin}}}\tau_{i\ldots n,g}. \tag{6}$$

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 ECMWF longwave radiation scheme is particularly well suited to adding the exact computation of $\partial L^{\uparrow}_{i-1/2}/\partial L^{\uparrow}_{\mathrm{surf}}$ alongside the existing calculation of fluxes for two reasons. First, the individual g points in the Monte Carlo Independent Column Approximation [*Pincus et al.*, 2003] treat the atmosphere as plane-parallel, i.e., there is no partial cloudiness treated within each single monochromatic calculation. Second, 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. If these partial derivatives were to be computed in radiation scheme that represented longwave scattering, it would be commensurate with the overall accuracy of the new scheme to neglect scattering for the purposes of computing these derivatives, i.e., to still apply (6) directly.

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 meters 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 ECMWF model. 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 ECMWF model, 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 neighboring 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 d^{-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

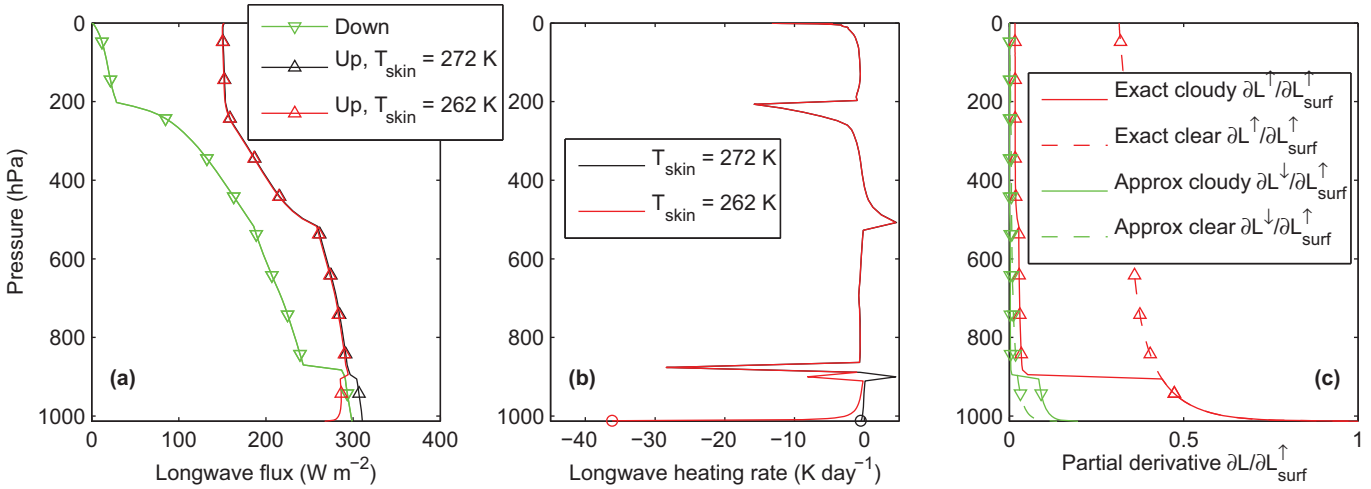

Figure 1. Illustration of the impact of skin temperature (T_{skin}) on (a) longwave flux and (b) heating-rate profiles by applying RRTM-G to the Mid-Latitude Winter standard atmosphere with a liquid cloud (mixing ratio 0.2 g kg^{-1}, effective radius 10 μm, cloud fraction 0.75) between 860 and 900 hPa and an ice cloud (mixing ratio 0.05 g kg^{-1}, effective radius 50 μm, 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 very closely by taking the control profile and updating the upwelling fluxes using (6) and (7). The heating rates of the lowest model layer are shown by the circles in Figure 1b. (c) Partial derivative of upwelling and downwelling fluxes with respect to the surface upwelling flux for clear and cloudy versions of the profile shown in Figure 1a: (red) the exact derivative of upwelling flux and (green) the assumed derivative of downwelling flux accounting approximately for the warming of the atmosphere.

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 d^{-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 Figure 1c depict the partial-derivative profile $\partial L^{\uparrow}/\partial L^{\uparrow}_{\mathrm{surf}}$ 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, again suggesting that it is important to adjust the atmospheric heating-rate profile in response to a change to surface upwelling flux. In the cloudy case, most of the remainder is then absorbed at cloud base. The green lines are discussed in the next section.

2.3. Updating the Net Longwave Flux Profile

Section 2.1 described how a longwave radiation scheme can be modified to provide the profile of $\partial L^{\uparrow}_{i-1/2}/\partial L^{\uparrow}_{\mathrm{surf}}$, in addition to the profile of net longwave flux $L^{n}_{\mathrm{ref},i-1/2}$ and the upwelling and downwelling flux components at the surface $L^{\uparrow}_{\mathrm{ref,surf}}$ and $L^{\downarrow}_{\mathrm{ref,surf}}$ (where "ref" indicates reference values output by the full radiation scheme called every 1 or 3 h, which are subsequently modified to respond to local surface conditions). Since the radiation scheme is run on a lower resolution horizontal grid than the rest of the model, these variables need to be interpolated back onto the native model grid where they are available for several time steps until the radiation scheme is called again.

The current version of the ECMWF model uses the net longwave flux profile to compute the profile of atmospheric heating rate in each of these intervening time steps, 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 skin temperature changes. This is the principal cause of 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 skin temperature, we may compute a new upwelling longwave flux from (4) and

$$L^{\uparrow}_{i-1/2} = L^{\uparrow}_{\mathrm{ref},i-1/2} + \left(L^{\uparrow}_{\mathrm{surf}} - L^{\uparrow}_{\mathrm{ref,surf}}\right) \frac{\partial L^{\uparrow}_{i-1/2}}{\partial L^{\uparrow}_{\mathrm{surf}}}. \tag{7}$$

If skin temperature were the only thing to change then (7) would very closely match what would be output from the radiation scheme if it were run at high temporal and spatial resolutions. 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 meters of the atmosphere is strongly coupled to skin temperature, due both to turbulent heat fluxes and to longwave radiative exchange (see Figure 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:

$$L^{\downarrow}_{\mathrm{surf}} - L^{\downarrow}_{\mathrm{ref,surf}} = \gamma\left(L^{\uparrow}_{\mathrm{surf}} - L^{\uparrow}_{\mathrm{ref,surf}}\right). \tag{8}$$

A value of γ=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 time step. A more refined approach would be to make γ dependent on near-surface temperature and humidity, similar to empirical methods to estimate surface downwelling longwave flux by considering the effective emissivity of the atmosphere [e.g., *Crawford and Duchon*, 1999], but from the experiments performed in this paper we find a fixed value of γ to be adequate.

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 $\partial L^{\downarrow}_{\mathrm{surf}}/\partial L^{\uparrow}_{\mathrm{surf}} = \gamma$ and the top-of-atmosphere value is $\partial L^{\downarrow}_{\mathrm{TOA}}/\partial L^{\uparrow}_{\mathrm{surf}} = 0$. This is achieved by

$$\frac{\partial L^{\downarrow}_{i-1/2}}{\partial L^{\uparrow}_{\mathrm{surf}}} = \gamma \frac{\partial L^{\uparrow}_{i-1/2}/\partial L^{\uparrow}_{\mathrm{surf}} - \partial L^{\uparrow}_{\mathrm{TOA}}/\partial L^{\uparrow}_{\mathrm{surf}}}{1 - \partial L^{\uparrow}_{\mathrm{TOA}}/\partial L^{\uparrow}_{\mathrm{surf}}}. \tag{9}$$

The green lines in Figure 1c 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

$$L^{n}_{i-1/2} = L^{n}_{\mathrm{ref},i-1/2} + \left(L^{\uparrow}_{\mathrm{surf}} - L^{\uparrow}_{\mathrm{ref,surf}}\right)\left(\frac{\partial L^{\downarrow}_{i-1/2}}{\partial L^{\uparrow}_{\mathrm{surf}}} - \frac{\partial L^{\uparrow}_{i-1/2}}{\partial L^{\uparrow}_{\mathrm{surf}}}\right). \tag{10}$$

A much simpler approach than described in this section would be not to update atmospheric heating rates at all; (4) would still be applied at the surface, but then the correction to surface net flux would be applied to the entire net flux profile. This is equivalent to assuming that any excess surface flux is transmitted directly to TOA, and is the approach currently implemented at the Met Office. In section 4.2, we compare these two approaches in their ability to predict the diurnal cycle of surface temperature.

3. Shortwave Method

Unlike skin 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 (and thereby correct for errors in the spatial interpolation of shortwave albedo and fluxes) need to be applied only once per radiation time step. However, we have also implemented the *Manners et al.* [2009] method at each model time step to better account for the variation of solar zenith angle between calls to the radiation scheme; their scheme updates the direct (unscattered) component of the solar radiation reaching the surface according to the known variation of path length through the atmosphere with solar zenith angle.

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

$$S^n_{\rm surf} = (1-\alpha) S^{\downarrow}_{\rm surf}. \tag{11}$$

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:

$$S^{\downarrow}_{\rm surf} = \tau S^{\downarrow}_{\rm TOA} + R S^{\uparrow}_{\rm surf}, \tag{12}$$

$$S^{\uparrow}_{\rm TOA} = \tau S^{\uparrow}_{\rm surf} + R S^{\downarrow}_{\rm TOA}, \tag{13}$$

$$S^{\uparrow}_{\rm surf} = \alpha S^{\downarrow}_{\rm surf}. \tag{14}$$

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) and (13) we find

$$\tau = \frac{S^{\uparrow}_{\rm surf} S^{\downarrow}_{\rm TOA} - S^{\uparrow}_{\rm surf} S^{\uparrow}_{\rm TOA}}{(S^{\downarrow}_{\rm TOA})^2 - (S^{\uparrow}_{\rm surf})^2}, \tag{15}$$

$$R = \frac{S^{\uparrow}_{\rm TOA} S^{\downarrow}_{\rm TOA} - S^{\uparrow}_{\rm surf} S^{\downarrow}_{\rm surf}}{(S^{\downarrow}_{\rm TOA})^2 - (S^{\uparrow}_{\rm surf})^2}. \tag{16}$$

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

$$S^n_{\rm surf} = S^{\downarrow}_{\rm TOA} \frac{\tau(1-\alpha)}{1-\alpha R}. \tag{17}$$

This equation may then be applied at the full model horizontal resolution immediately after the radiation scheme is called at a coarser resolution, in order to compute a new surface net flux according to the high-resolution surface albedo field.

To test surface fluxes predicted by the new scheme, offline radiation calculations have been performed using RRTM-G with the 137 pressure levels of the ECMWF model 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 would be using the boundary fluxes from the calculation over the ocean.

The results are shown in Figure 2. 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 Figure 2. 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.

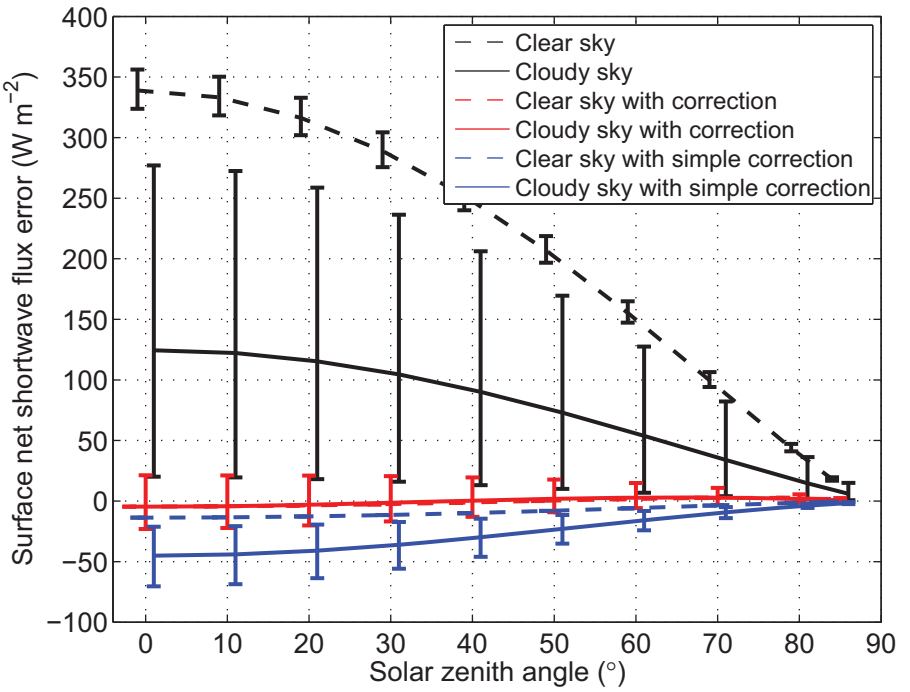

Figure 2. The error in surface net shortwave flux when a radiation calculation using a surface albedo of 0.08 is used over a surface with an albedo of 0.4, versus solar zenith angle, for all six standard atmospheres of *McClatchey et al.* [1972] and a range of cloud conditions. The solid lines depict 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 bias with no attempt to make a correction for surface albedo, the red lines depict the correction proposed in this paper (equation (17)), and the blue lines depict a simpler correction where the downwelling shortwave flux is assumed to be constant with surface albedo (equation (11)).

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 for shortwave radiation schemes that use separate albedos for different parts of the spectrum and/or for direct and diffuse radiation. In the case of ECMWF, $n = 4$ albedos are used. Therefore, in order to use (17) at each grid point, we need to convert the n individual albedos at each grid point (denoted α_1 to α_n) 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 n components of this flux corresponding to each of the n albedos (denoted $S_1^{\downarrow}$ to $S_n^{\downarrow}$). These components are interpolated from the radiation grid to the model grid and then used as follows to derive broadband albedo:

$$\alpha = \frac{\sum_{i=1}^{n} \alpha_i S_i^{\downarrow}}{S^{\downarrow}}. \tag{18}$$

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 the partial derivative of the net shortwave flux profile with respect to the surface net flux (as we have done in the longwave) 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 same assumption as *Manners et al.* [2009] in a similar context, that any excess upwelling solar radiation at the surface is lost to space, which means that the atmospheric heating rates are unchanged. 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.

4. Case Study

4.1. Correction of Coastal Errors

The approximate longwave and shortwave updates have been implemented in the ECMWF model 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 grid points in the vicinity of Long Island and Connecticut, and indeed at the time, forecasters at La Guardia airport alerted ECMWF of the poor forecast. It will be shown that this error was caused by errors in the surface longwave fluxes. The same forecast exhibited strong temperature overestimates around the coast of Arabia, and it will be shown that these were predominantly associated with shortwave errors. We run 3 day forecast experiments initialized at 12 UTC on 3 January 2014 using ECMWF model cycle 40R2 at T_L 1279 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 time step and grid point. For the purposes of this section, this simulation is treated as the "truth" that the approximate radiation update simulation should attempt to match.

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 3 and 4 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 Figures 3a and 3b are in the range -26 to $+5$ K, demonstrating how coastal nighttime land temperature can be substantially underestimated purely due to a longwave error: the net longwave flux over the sea is applied over the adjacent land. The differences between Figures 4a and 4b are in the range -2 and $+9$ K, illustrating a large daytime skin temperature overestimate at a desert coastline. 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 3c and 4c show that approximate radiation updates produce much more similar skin temperatures to the simulation running the radiation scheme at all time steps and grid points, but of course with a much smaller computational cost. The difference between the approximate updates and high-resolution radiation is up to ± 2.1 K in Figure 3, and up to ± 0.8 K in Figure 4. To understand in more detail how the approximate updates modify the surface fluxes, Figures 5 and 6 depict the time series of surface net fluxes, and skin and 2 m temperatures for the full 72 h of the forecasts, for the points indicated by the white circles in Figures 3a and 4a 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 26 K, compared to the high-resolution radiation benchmark experiment. It can be seen that this bias builds up entirely during the night and therefore must be a longwave effect. Figure 5a reveals an underestimate in net longwave radiation by at least 50 W m^{-2} for the first 36 h of the forecast, 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 use of approximate longwave updates (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 the high albedo of the snow cover in Long Island. In the first 2 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 error is a strong daytime overestimate in skin temperature. Figure 6a shows that the control forecast overestimates the surface net shortwave flux by up to 200 W m^{-2},

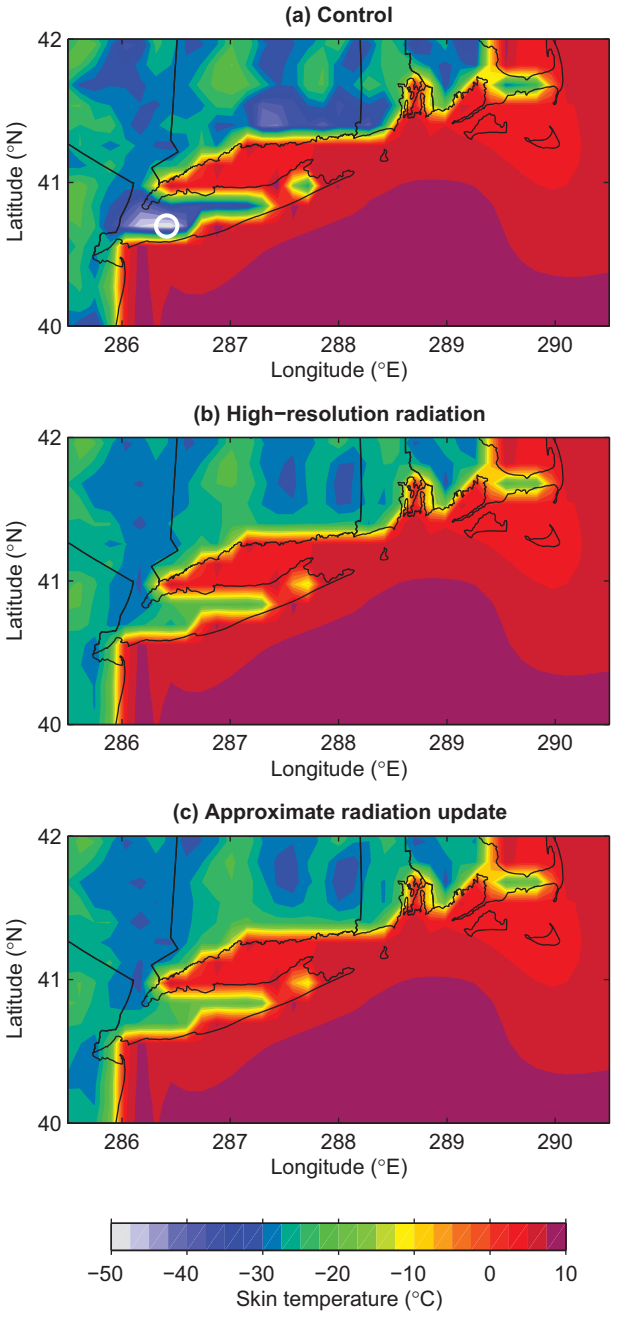

Figure 3. Skin temperature at 11 UTC (0600 local time) on 4 January 2014 for a region around the coast of Long Island and Connecticut from forecasts initialized 24 h previously: (a) control T_L 1279 model, (b) the same model but with the radiation scheme run at every time step and grid point, (c) as Figure 3a 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 Figure 5.

an error that is almost perfectly corrected by the approximate shortwave update. The longwave fluxes also show significant errors, with the time series 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 (peaking at around 50 W m^{-2}), 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 time series in Figure 6b is closely reproduced.

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 (specifically the longwave updates) can also correct for temporal errors associated with the radiation scheme being called infrequently in time. Such errors are particularly apparent for the model configuration used in ensemble forecasting, in which the radiation scheme is called only every 3 h. We have replicated the ensemble model configuration (T_L 639 resolution, but with radiation at a spatial resolution of T_L 255) 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 h is shown in Figure 7.

The black line in Figure 7a demonstrates the limitation of the current longwave scheme: surface net longwave flux, L^n_{surf}, 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 time step and gridbox, we see in Figure 7b that this leads to instantaneous errors of up to almost 100 W m^{-2}. *Bozzo et al.* [2015] showed a similar comparison of a diurnal cycle of L^n_{surf} over the Nevada desert from model versions with

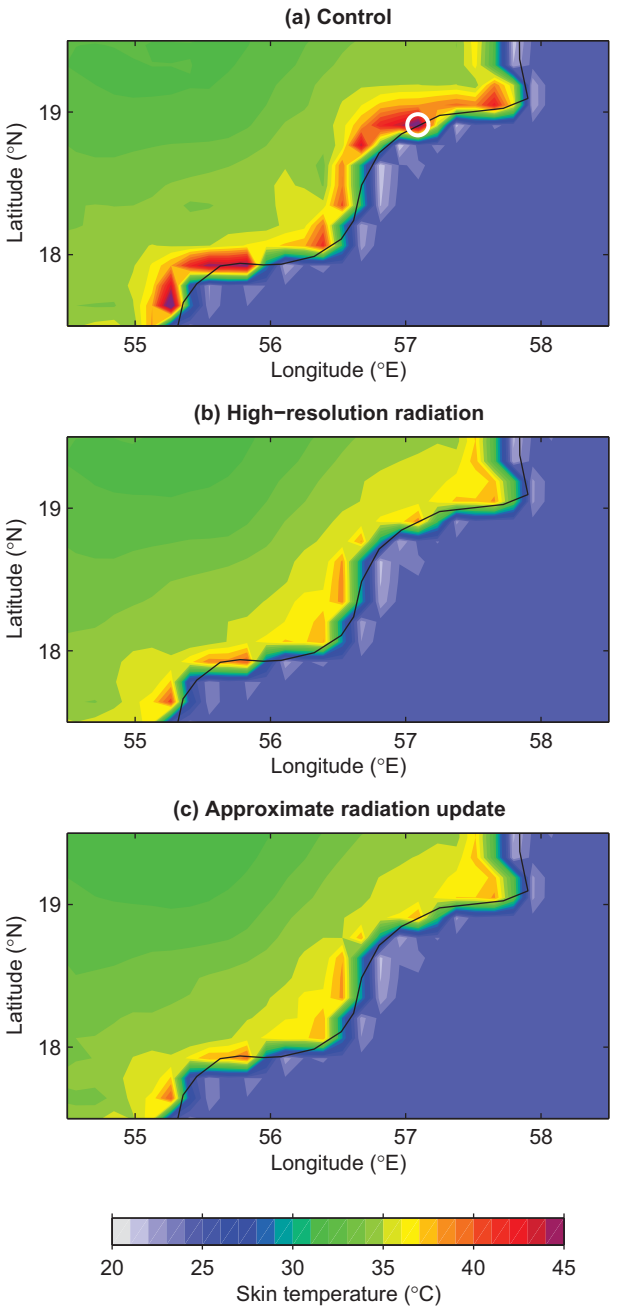

Figure 4. As Figure 3 from the same forecasts but at 09 UTC (1300 local time) on 4 January 2014 for a region around the southern coast of Oman. The time series of surface variables at the point indicated by the white circle is shown in Figure 6.

different radiation time steps, and found that calling the radiation scheme every model time step led to the best agreement with observed surface fluxes. The red solid line shows that a 3 h radiation time step but with approximate updates to the radiative fluxes every model time step reduces this error to typically no more than 10 W m^{-2}. If we apply a simpler longwave update in which surface upwelling flux is updated according to the local skin temperature but any excess is lost to space and the downwelling flux is unchanged then the result is shown by the red dashed line. It can be seen that this provides a poorer approximation to the high-resolution radiation simulation.

The black line in Figure 7d shows that the current shortwave scheme also has substantial 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 3 h. Consider the situation in the morning when solar zenith angle reduces with time: the solar zenith angle used by the radiation scheme is computed for a time halfway between radiation time steps, so for model time steps near the start of the 3 h period it will be too small. This means that the path length through the atmosphere will be too short and too much direct solar radiation will reach the surface. The opposite effect occurs at the end of the 3 h period. The red line shows that these errors are largely removed via the use of approximate radiation updates, but in this case the specific reason is our implementation of the scheme of *Manners et al.* [2009] that approximately corrects the direct solar flux to account for the (known) error in path length through the atmosphere. Note that unlike the longwave errors considered above, these shortwave errors are essentially random so will not lead to a significant error in the diurnal cycle of surface temperature.

Figure 8 confirms that the much improved fluxes depicted in Figure 7 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 by 1–2 K. This arises due to the longwave net flux being updated only every 3 h, which means that radiation emitted from the surface does not

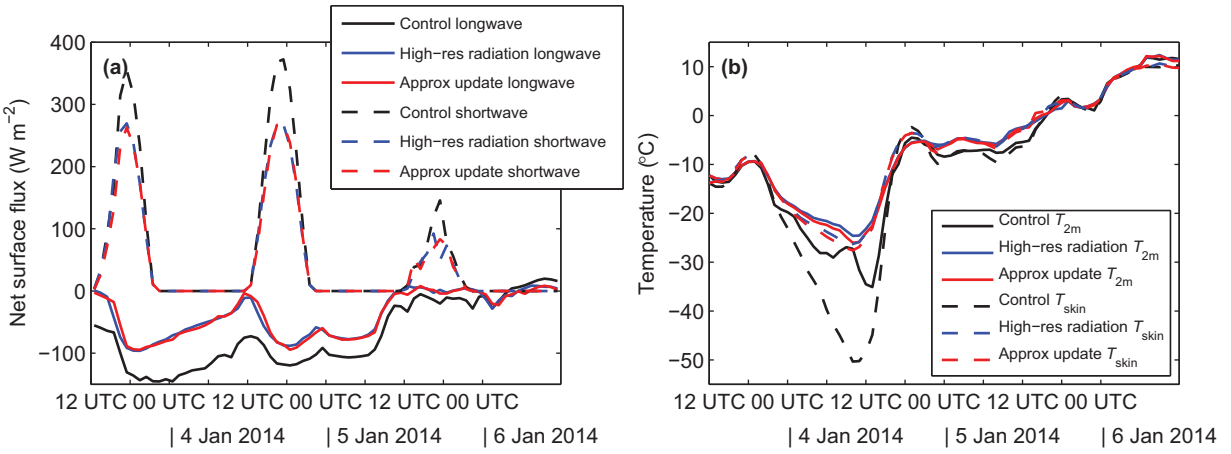

Figure 5. Time series of (a) surface net shortwave and longwave fluxes and (b) 2 m temperature and skin temperature at 40.7°N, 286.4°E (the point on Long Island indicated by the white circle in Figure 3a) for the full 72 h forecasts using same three model configurations.

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. The red dashed lines show the result of the simpler longwave update, and the results are clearly much worse than implementing the full scheme described in section 2.

One puzzling aspect is the immediate evolution of a positive temperature error in both the control and approximate update experiments in the first 3 h 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_L 639 forecast is initialized from a T_L 1279 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.

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 00 UTC each day for four Northern Hemisphere summer

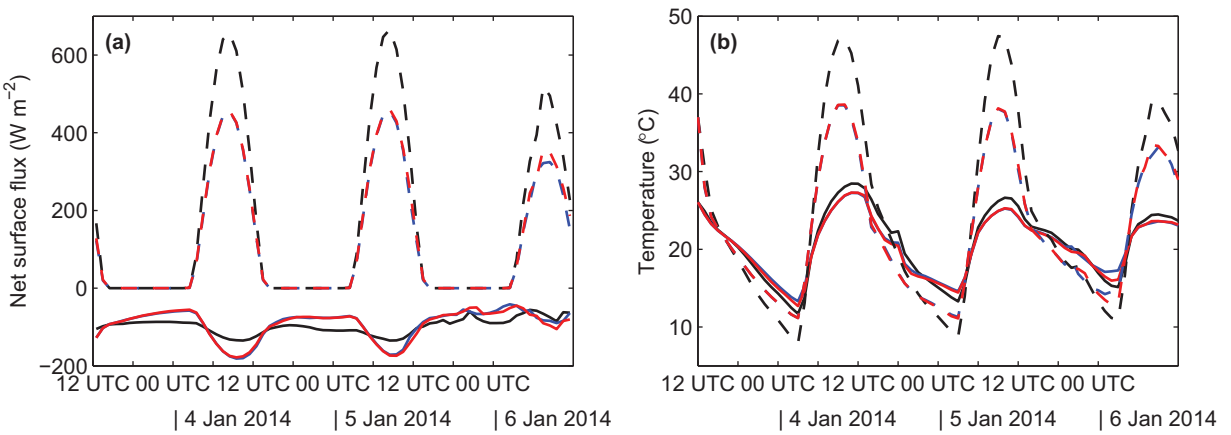

Figure 6. As Figure 5 but for 18.9°N, 57.1°E (the point on the coast of Oman indicated by the white circle in Figure 4a).

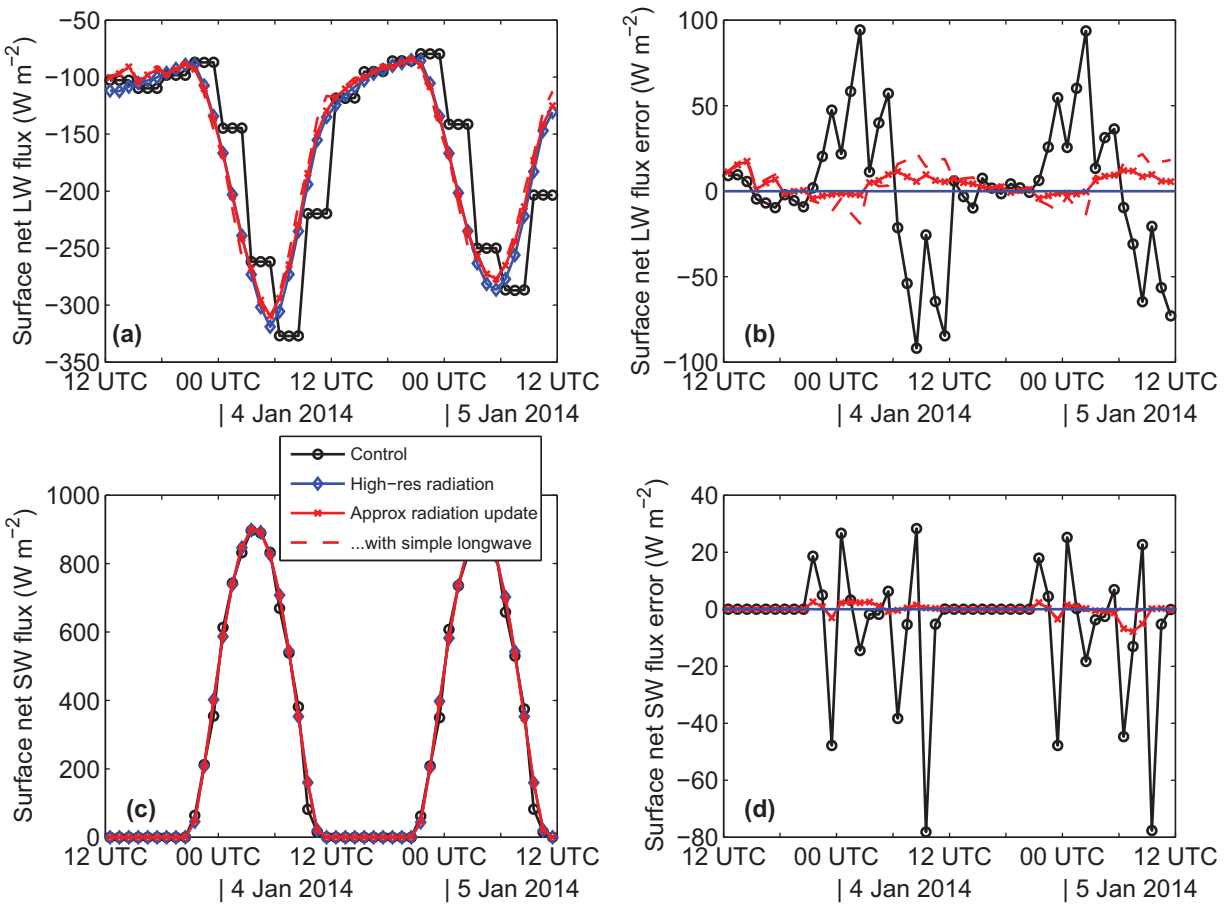

Figure 7. Time series of surface net fluxes at 26.25°S, 123.25°E (a point in the Australian Desert) for the same forecast period as in Figures 3–6, but for a model resolution of T_L 639 and a default radiation time step of 3 h. (a) Longwave fluxes from the control model configuration, the "high-resolution radiation" simulation with the radiation scheme run at every gridbox and time step, and the new approximate radiation update. The dashed red line depicts the results of a simpler approximate longwave scheme where (4) is applied at the surface but excess flux is assumed lost to space with no change to atmospheric heating rates. (b) Longwave flux minus the values from the high-resolution radiation simulation, treated as truth. (c and d) As Figures 7a and 7b but for shortwave radiation.

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 ECMWF model cycle 40R2. Two model resolutions are used matching those used in section 4: T_L 1279 with the radiation scheme called by default every 1 h, and T_L 639 with the radiation scheme called by default every 3 h. The results are analyzed 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_L 639, 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], the left plot of Figure 9 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 time step and grid point (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 solely by the longwave mechanisms discussed in section 4.2: assuming the surface net longwave fluxes are constant for 3 h leads to a

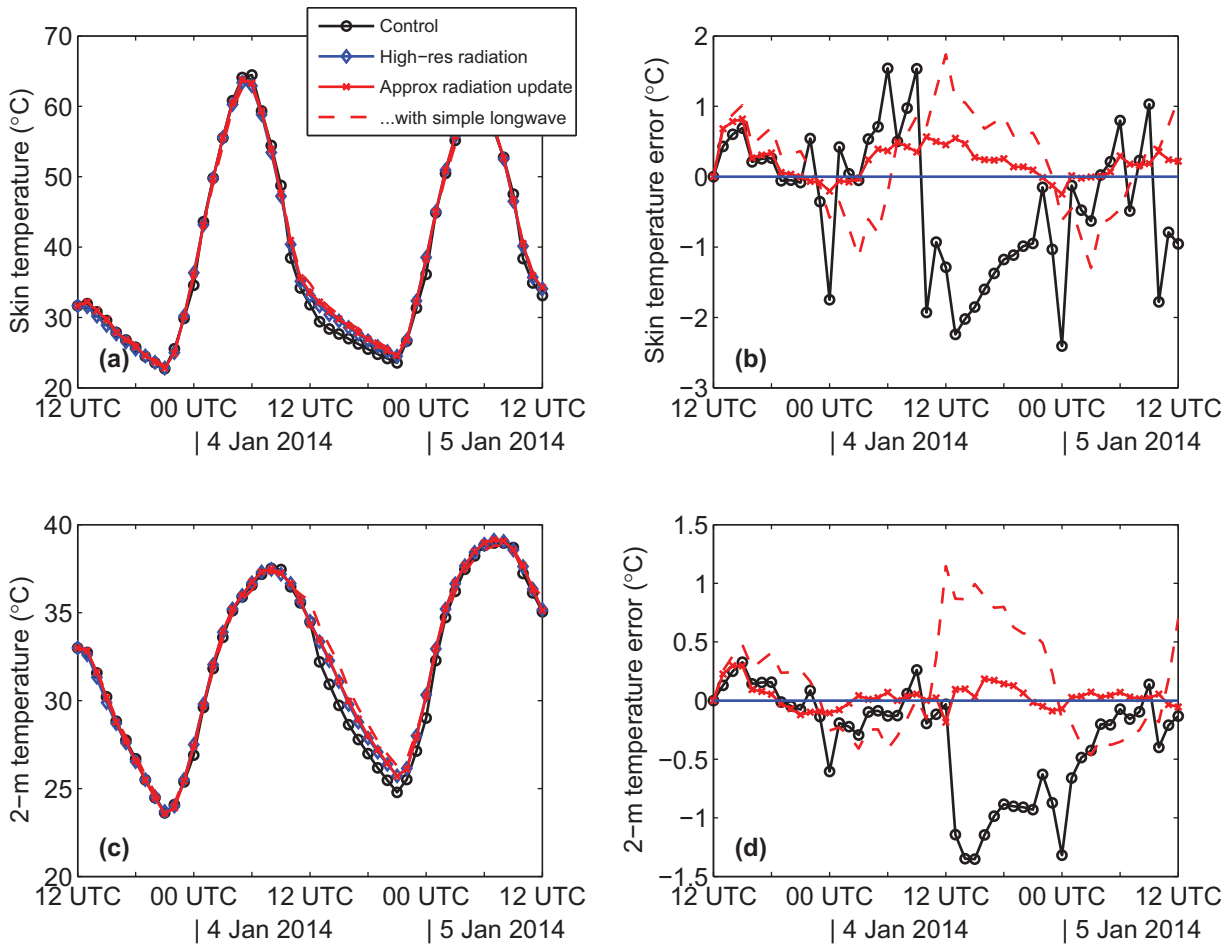

Figure 8. As Figure 7 but for (a and b) skin temperature and (c and d) 2 m temperature.

lag in the surface cooling rate. In particular, during the evening, the cooling rate is too strong, leading to an underestimate in the nighttime 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.

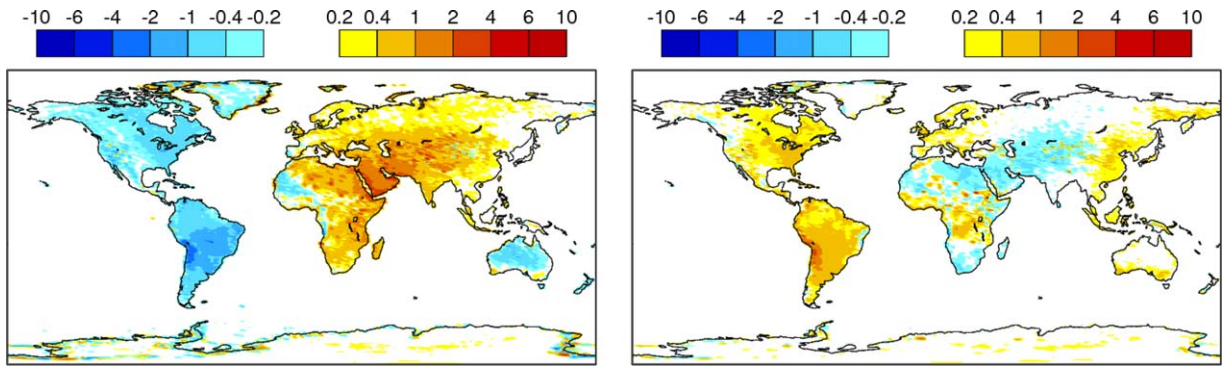

Figure 9. (left) Difference in mean skin temperature (K) at 12 UTC between T_L 639 forecasts with the radiation scheme called every 3 h and intermittently in space, and forecasts with the radiation scheme called every time step and gridbox. (right) The same but with approximate updates to the radiation fields every time step and grid point. One forecast is initialized each day from the analysis at 00 UTC, for 4 months between 1 June and 30 September 2012, and skin temperature values 36 h into each forecast have been extracted. The approximate update scheme assumes a longwave downwelling factor of $\gamma=0.2$ in (8).

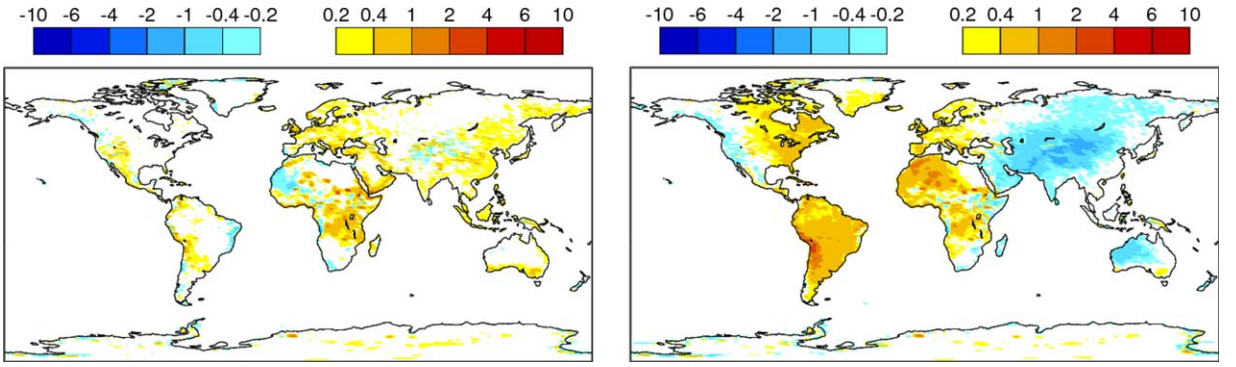

Figure 10. (left) As the right plot of Figure 9 but for approximate updates using a longwave downwelling factor of $\gamma = 0$. (right) The same but using a longwave downwelling factor of $\gamma = 0.4$.

The right plot of Figure 9 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 orography, particularly in Africa. It is also worth noting that coastal temperature errors are apparent in the control

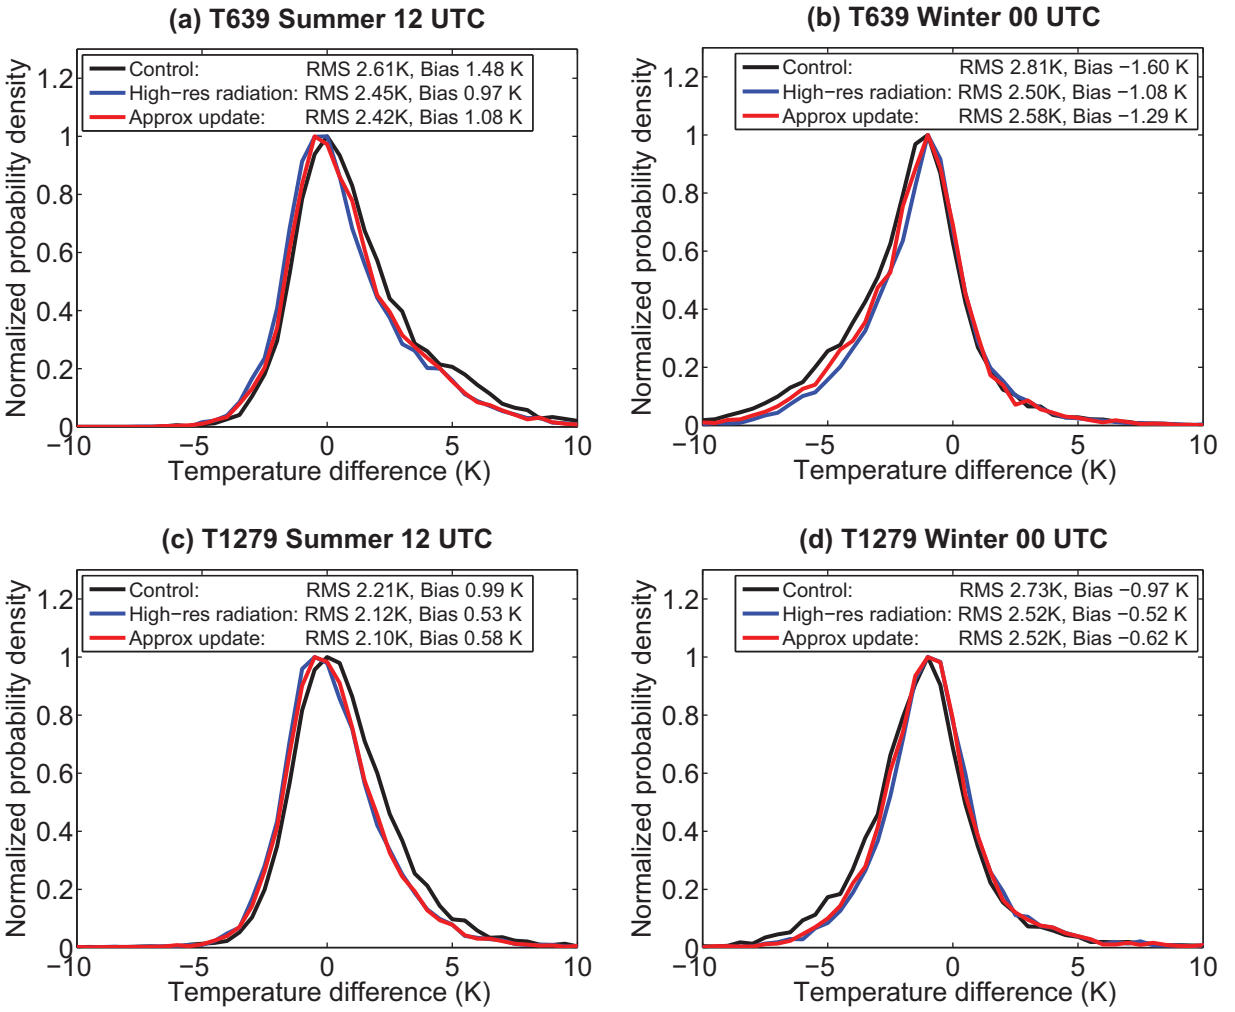

Figure 11. Probability density functions (normalized to the maximum value) of the difference between forecast and observed 2 m temperature at European coastal stations. (left column) Results for June–September 2012 forecasts at 12 UTC, 36 h into each daily forecast and (right column) results for December 2012 to March 2013 forecasts at 00 UTC, 24 h into each daily forecast. Results for (top row) the T_L 639 resolution model and (bottom row) the T_L 1279 model.

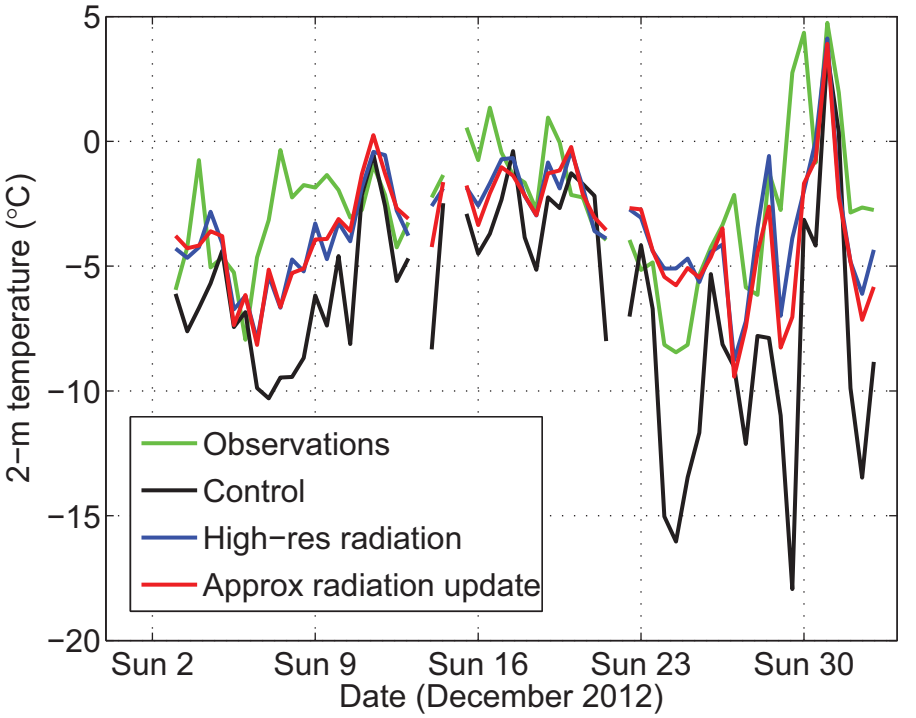

Figure 12. Comparison of 00 UTC and 12 UTC 2 m temperature forecasts (at lead times of 24 and 36 h, respectively) against observations at Sortland, Norway (68.7°N, 15.42°E) for December 2012. The T_L 639 model simulations were the same as used to produce Figure 11b.

simulation (left plot), particularly on the east coast of Greenland and around the coast of Arabia, but these are largely removed in the right plot.

In this case, approximate radiation updates have been applied with a longwave downwelling factor of $\gamma=$ 0.2 (described in section 2.3). Figure 10 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, although there is scope to refine the parameterization in future by making this parameter dependent on atmospheric properties such as near-surface temperature and humidity.

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 grid point 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 analyzed, between 55 and 95 stations are available. The distribution of stations is shown in *Bozzo et al*. [2015, Figure 6].

Figure 11 shows the probability distribution of 2 m temperature error for the two model resolutions at 12 UTC in summer months and 00 UTC 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 12 UTC, and a tail of large negative temperature errors in winter at 00 UTC. As found by *Bozzo et al*. [2015], this is much improved when the radiation scheme is called every time step and grid point (blue lines), as indicated also by the improvement in both bias and root-mean-squared error indicated in the legends of Figure 11. The results from using approximate radiation updates (red lines) are generally quite similar to calling radiation every time step and grid point, confirming observationally that it leads to significantly better forecasts. Evaluation of 00 UTC forecasts in summer and 12 UTC forecasts in winter leads to much less difference between any of the lines (not shown).

Figure 12 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

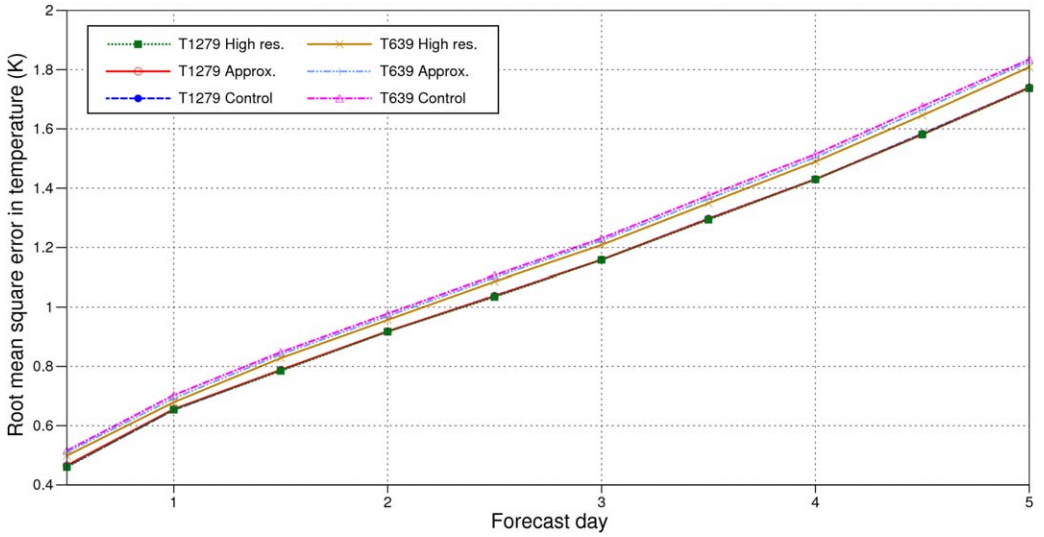

Figure 13. Comparison of the root-mean-squared error of daily 1000 hPa temperature forecasts in the Northern Hemisphere extratropics, versus lead time for June–September 2012 inclusive. The model configurations at the two horizontal resolutions are the control configuration (radiation scheme called at a reduced spatial resolution every 1 h at T_L 1279 and every 3 h at T_L 639), approximate radiation updates using the scheme described in this paper, and high-resolution radiation in which the radiation scheme is called every time step and grid point.

differences in excess of 10 K toward the end of the month. This bias is largely removed both in forecasts with the radiation scheme run every time step and grid point, and at much lower computational expense in forecasts using approximate radiation updates.

5.3. Forecast Skill

In this section, we briefly examine 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 extratropics. Figure 13 depicts the root-mean-square error in 1000 hPa temperature forecasts at all land and sea points for June to September 2012, inclusive, assessed by comparing to the operational analyses. At T_L 639, calling radiation every time step and grid point leads to significantly better forecasts, while the use of approximate radiation updates also leads to a significant improvement, but with errors that lie around halfway between the control and the 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. A similar picture is found for the December 2012 to March 2013 period (not shown).

When considering other model variables, in particular geopotential height, running the radiation scheme every gridbox and time step leads to a small reduction in root-mean-square error at T_L 639 compared to the control, but forecasts using approximate radiation updates are essentially indistinguishable from the control (not shown). This implies that the improvements to surface and near-surface temperature forecasts, particularly at coastlines, are at too small a scale in a global context to measurably improve the evolution of weather systems.

The difference between the various T_L 1279 simulations is very small in terms of temperature (Figure 13) and other variables, and only marginally statistically significant. 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_L 639 control forecasts.

6. Conclusions

In this paper, modifications to the ECMWF model 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 time step and grid point. 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; specifically, the approximate update of longwave fluxes remedies the underestimate of nighttime minimum temperatures, while the combination of approximate updates to shortwave and longwave fluxes remedies the overestimate of daytime maximum temperatures. The longwave part of the new scheme also improves a lag in the diurnal cycle of surface temperature. *Manners et al*. [2009] scheme is found to significantly reduce 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. *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 modeling by *Pincus and Stevens* [2009]. However, this spectral sampling leads to a large noise in the instantaneous fluxes; although this averages out after a few hours, it makes instantaneous radiative fluxes output from the model unsuitable for direct evaluation against observations or for applications such as solar energy forecasting. The approximate radiation updates described in this paper do not suffer from this problem, and are scheduled to be incorporated into a forthcoming version of the operational ECMWF model instead.

While approximate updates lead to better surface temperature forecasts, they appear not to significantly improve the skill of geopotential-height forecasts, even though there is some improvement from running the radiation scheme every time step rather than every 3 h. 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 spectra, 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].

Acknowledgments
The model experiments are based on the ECMWF IFS code, which is a comprehensive NWP model documented at http://old.ecmwf.int/ifsdocs/. The code changes are described in the current manuscript. We are grateful to Irina Sandu, Anton Beljaars, Linus Magnusson, Tim Hewson (ECMWF), 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 Figures 1 and 2.

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. Model. Earth Syst.*, *6*, 1288–1300, doi:10.1002/2014MS000386.

Crawford, T. M., and C. E. Duchon (1999), An improved parameterization for estimating effective atmospheric emissivity for use in calculating daytime downwelling longwave radiation, *J. Appl. Meteorol.*, *38*, 474–480.

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., *Rep. No. AFCRL72–0497*, Air Force Cambridge Res. Lab., L. G. Hanscom Field, Mass.

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

Morcrette, J.-J., H. W. Barker, J. N. S. Cole, M. J. Iacono, and R. Pincus (2008a), Impact of a new radiation package, McRad, in the ECMWF integrated forecasting system, *Mon. Weather Rev.*, *136*, 4773–4798.

Morcrette, J.-J., G. Mozdzynski, and M. Leutbecher (2008b), A reduced radiation grid for the ECMWF integrated forecasting system, *Mon. Weather Rev.*, *136*, 4760–4772.

Mozdzynski, G., and J.-J. Morcrette (2014), Reorganization of the radiation transfer calculations in the ECMWF IFS, *Tech. Memo. 721*, 20 pp., Eur. Cent. for Medium-Range Weather Forecasts, Reading, U. K.

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. Model. Earth Syst.*, *1*, 1, doi:10.3894/JAMES.2009.1.1.

Pincus, R., H. W. Barker, and J.-J. Morcrette (2003), A fast, flexible, approximate technique for computing radiative transfer in inhomogeneous cloud fields, *J. Geophys. Res.*, *108*(D13), 4376, doi:10.1029/2002JD003322.

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. Top. Appl. Earth Obs. Remote Sens.*, *7*, 3660–3667.

Schuster, A. (1905), Radiation through a foggy atmosphere, *Astrophys. J.*, *21*, 1–22.