

Reorganization of the Radiation Transfer Calculations in the ECMWF IFS

George Mozdzynski and
Jean-Jacques Morcrette

Research Department

April 2014

*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.*



Series: ECMWF Technical Memoranda

A full list of ECMWF Publications can be found on our web site under:

<http://www.ecmwf.int/publications/>

Contact: library@ecmwf.int

©Copyright 2014

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

Radiation computations in IFS are a relatively expensive part of a high resolution forecast model taking approximately 10 percent of the total time. This current level of cost is achieved by running the radiation scheme only once every forecast hour and by using a radiation grid resolution which is coarser than that of the model grid. This technical memorandum investigates a new approach for the radiation scheme where the radiation computations execute in separate tasks from the rest of the model, with the objective of improved scalability and performance. The approach serves as an example that could be applied to other parts of IFS as scalability becomes increasingly more challenging on future supercomputers.

1 Introduction

Radiation transfer (RT) is one of the most expensive processes to be represented in a large-scale atmospheric model. Over the years, at ECMWF, various approaches have been used to get a state-of-the-art RT representation into the Integrated Forecasting System (IFS) while maintaining the cost of radiation within acceptable limits. For example, Morcrette (2000) discussed various strategies related to either a temporal or a spatial (or both) sampling of the radiation inputs, which was in use at the time with the pre-CY36R2 radiation package. Following the introduction of a computer with a distributed memory architecture (Dent and Mozdzyński, 1996, Mozdzyński, 2007), and the implementation of the more expensive McRad radiation package (Morcrette et al., 2008a), a reduced grid for radiation was implemented in the IFS, a grid coarser than the grid on which all other physical processes are computed (Morcrette et al., 2008b).

As of Spring 2013, radiation transfer in the high-resolution IFS, run at T_L1279 L137, uses a radiation reduced grid at T_L511 with full radiation computations performed every hour, from which short-wave (SW) transmissivities are defined. The SW radiation fields at in-between time-steps are computed for the relevant solar zenith angle based on these transmissivities. Net long-wave (LW) fluxes are kept constant between two full radiation time-steps.

This memorandum discusses a re-organization of the radiation transfer calculations, being investigated as part of the EU 7th Framework Programme CRESTA project (Collaborative Research into Exascale Systemware, Tools and Applications) for a future computer system with hundreds to thousands times more cores than used today.

This reorganization has the potential for a sizeable increase in efficiency for whatever parametrization is presently called outside of the main stream (the **GP_MODEL-EC_PHYS-CALLPAR** call tree). Radiation calculations is one of the main candidates for such a treatment. The radiation-in-parallel re-organization also has some potential drawbacks, which are discussed and tentatively quantified in the following.

Section 2 presents the current and radiation-in-parallel configurations. Section 3 compares results for various frequencies for calling the so-called full radiation computations for a T_L511 L91 model configuration. This section also addresses the dependence of the results of horizontal resolution and time-step length. Section 4 presents the performance of the radiation-in-parallel scheme. A discussion of the results with some conclusions, and perspectives are given in Section 5.

2 Model description and experimental design

2.1 Current configuration

The path to radiation calculations is more complex than the one for the other physical parameterizations. The radiation driver **RADDRV** is called from **GP_MODEL**, the routine doing the computations in grid-point space (therefore calling **EC_PHYS** which calls **CALLPAR** from where all the other physical parametrizations are called). **RADDRV** gets the fields from the global repository (profiles of temperature, humidity, cloud fraction and condensates, surface skin temperature and surface albedos), provides them to **RADINTG**, where they are interpolated to the reduced radiation grid (and climatological information is added). From **RADINTG**, **RADLSWR** calls the LW and SW radiation schemes. The radiation fluxes produced within **RADLSWR**, once interpolated back to the physical grid, go back up the path to **GP_MODEL** where they feed **EC_PHYS** then **CALLPAR** from where all the other physical parametrizations are called.

2.2 Radiation-in-parallel configuration

In the radiation-in-parallel configuration the radiation transfer calculations execute in parallel with the rest of the model using separate MPI tasks as shown in Figure 1. From an MPI point of view, the radiation and model tasks have separate MPI communicators, while both can also use a global communicator (called **COMM_WORLD**) for exchanging data between model and radiation tasks. Data from the model required for the radiation transfer calculations is sent asynchronously from within **RADINTG** after input interpolation (if required). Similarly data from the radiation scheme is received within **RADINTG** before output interpolation (if required). An important and necessary requirement for this configuration is that the product of the radiation transfer calculations are returned to the model shifted by one radiation time-step. This shift is necessary to allow the radiation transfer calculations to execute completely independently during the time the model executes a full time-step. The current and radiation-in-parallel configurations are supported today by setting an environment variable (**RADPAR_RUN**). If the current configuration is required then **RADPAR_RUN** should either be set to 'no' or remain unset. If the radiation-in-parallel configuration is required then **RADPAR_RUN** should be set to 'yes' and in this case it is also necessary to set **NPROC_RADPAR** to the number of tasks to be used for the radiation scheme. All code to support the radiation-in-parallel configuration is contained in branch **mpm_CY40R1_radpar**.

2.3 Initial test configuration

In an earlier study, the radiation schemes were called within the chain discussed in 2.1, but the results were stored, so as to fill the model with radiation tendencies one radiation time-step earlier than the time-step seen by the rest of the physics. By doing so, it simulated the configuration described in 2.2 in which radiation tendencies are computed in parallel to the first part of dynamics. This test configuration was called **radparsim** and was supported by branch **mpm_CY38R2_radparsim**.

3 Results from the initial test configuration

A complete set of plots was produced (with plots for geopotential, temperature, winds and relative humidity at 1000, 850, 500, 200, 100, and 30 hPa). The following will only present the most representative plots.

Three questions are relevant for this development:

- What would be the benefits of having radiation computed at every time-step?
- What would be the impact of having radiation computed at every time-step but one time-step earlier than for the other physical parametrisations?
- How are the results from this last question affected by differing model horizontal resolutions (i.e., time-step lengths)?

To address these three questions, a series of experiments with the cycle 38R2 model have been run using different model configurations, in each case for 37 ten-day forecasts, every ten days from 20120101 (12UTC) to 20121226 (UTC):

- default 38R2 at all resolutions from T_L159 to T_L1279 calling full radiation computations every hour (hereafter called **REF1HR**);
- as above, but with full radiation computations at every model time-step (hereafter called **REF1TS**);
- 38R2+**radparsim** with full radiation computations every hour (hereafter called **RPS1HR**);
- 38R2+**radparsim** with full radiation computations at every model time-step (hereafter called **RPS1TS**).

In the following, all forecasts are scored against the operational analysis at the time.

3.1 Impact of radparsim

For the model at T_L511 L91, a comparison for **RPS1HR-REF1HR**, and **RPS1TS-REF1TS** for the geopotential height at 1000 and hPa is presented in Figure 2, and for temperature at 850 and 200 hPa in Figure 3. Both figures are for the Northern hemisphere. The first set **RPS1HR-REF1HR** (in red) corresponds to the current radiation configuration with the full radiation called every hour and interpolated at time-steps in-between. The second set **RPS1TS-REF1TS** (in blue) is for the radiation called at every time-step (1200 s at T_L511 L91). Not surprisingly, the difference between the **RPS-REF**, whether for geopotential height or temperature, is much smaller when the full radiation computations are called at every model time-step. This result holds for other variables (wind and relative humidity) and other areas (Southern hemisphere, Europe, Tropical area).

3.2 Impact of full radiation at every time-step

For the same model configuration (T_L511 L91), Figures 4 and 5 presents comparisons for **REF1TS-REF1HR** (in red) and **RPS1TS-RPS1HR** (in blue) for the geopotential height at 1000 and 500 hPa

(Fig.4) and temperature at 850 and 200 hPa (Fig.5) in the Northern hemisphere. In these figures, a relative improvement is seen in curves appearing above the zero line. Such an improvement is more prevalent for the **radparsim** configuration, which simply indicates that **RPS1HR** was likely farther from the verifying analysis and that **RPS1TS** moved closer to it. Here again, similar results are obtained for the other areas and variables.

3.3 Impact of increasing resolution/decreasing model time-step

Table 1 presents the model horizontal resolution, radiation grid resolution and time-step in the various configurations of the IFS from T_L159 to T_L1279 that were tested with **radparsim**.

Model resolution	159	255	319	399	511	639	799	1279
Radiation grid	63	95	159	159	159	255	319	511
Time-step	3600	2700	1800	1350	1200	900	720	600

Table 1: Radiation grid and time-step for the IFS at different horizontal resolutions.

Results of a comparison **RPS1TS-REF1HR** for the L91 model at resolutions T_L159 (red), T_L319 (blue), T_L639 (green) and T_L1279 are presented in Figures 6 and 7 for the geopotential height in the Northern hemisphere (Fig.6) and the tropical area (Fig.7), whereas comparison for temperature at 850 and 250 hPa for the Southern hemisphere is shown in Figure 8. Whatever the parameter, the main message is that at the lower resolutions **RPS** is slightly better in the tropics than **REF** and fairly neutral in the extratropics, but that the difference decreases (as expected) with increased horizontal resolution (i.e., smaller time-step).

3.4 Impact on radiation fields

Given that the **radparsim** approach directly affects the radiation calculations, it is worth looking at the potential impact on the radiation fields at the surface and top of the atmosphere. Figure 9 presents the impact on the outgoing long-wave radiation, and Figure 10 summarizes the impact on the surface net long-wave radiation, and the net short-wave radiation at the surface and top of the atmosphere. All plots are for the average over the first five days of the three forecasts started on 1st, 11th and 21st January 2012. These results indicate that such an approach with radiation computed from fields one model time-step earlier would not create unacceptable systematic decrease in forecast quality. Although different in essence, the perturbation so introduced is not incommensurate with what is already in the model (temporal sampling, reduced radiation grid, sequential call to physics subroutines, McICA approach, future McSI). With future higher and horizontal (and vertical) resolution in an exascale computer system, the model time-step is likely to further decrease, reducing further the impact of having such radiation fields one time-step earlier than the rest.

4 Radiation-in-parallel performance

The main purpose of the radiation-in-parallel scheme is to improve the computational performance of an IFS model. To see how it performs in practice we looked at two cases, a small T_L159 model and a large $T_L3999L137$ model case.

4.1 T_L159 model performance study

To understand how the radiation-in-parallel configuration performs it is useful to consider a low resolution T_L159 model and see how this runs for increasing numbers of MPI tasks. In this study each MPI task uses 8 OpenMP threads. The default radiation resolution is T63 for this T_L159 model. Further, the number of grid columns for the T_L159 model is 35,718 and for the T63 radiation grid there are 6,114 columns. The frequency at which radiation transfer calculations are performed is by default every 3 hours for a T_L159 model, which has a time-step of 1 hour, i.e. NRADFR=3. Table 1 shows T_L159 model performance for increasing number of MPI tasks, for both the default radiation frequency (NRADFR=3) and where radiation is called every time-step (NRADFR=1). These runs were performed in a non-dedicated mode on TITAN, a CRAY XK7 at the Oak Ridge National Laboratory.

Tasks	NRADFR=3 (default)	NRADFR=1
8	2,316	1,539
32	8,142	5,533
128	24,598	18,438
512	45,117	39,775
1024	46,734	44,240

Table 1: T_L159 model performance in forecast days per day.

From this table we can see that a T_L159 model reaches its asymptotic performance at around 512 or 1024 tasks for the default (NRADFR=3) radiation configuration, and that performance is similar at 1024 tasks for both 3 hourly and hourly radiation configurations.

Now if we use the radiation-in-parallel scheme with NRADFR=1, we can achieve **58,858** forecast days per day (FD/D) when using 1024 tasks, an improvement of over 25 percent on the NRADFR=3 configuration. This was achieved with a split of 768 tasks for the model and 256 tasks for the radiation (the sum being 1024).

Model tasks	Radiation tasks	FD/D
512	512	53,004
640	384	55,837
704	320	58,048
736	288	57,212
768	256	58,858
800	224	57,920
832	192	55,946
896	128	48,941

Table 2: T_L159 model performance for 1024 tasks using the radiation-in-parallel scheme and NRADFR=1.

Table 2 shows the effect of using different combinations of model and radiation tasks that together sum to 1,024. We can see that using 224 to 320 radiation tasks achieve similar performance of over 57,000 FD/D. When we repeated the 768:256 split case this achieved 56,959 FD/D which suggests an error of about 2,000 FD/D for this T_L159 model configuration. This variability is most likely due to the effect of other jobs running on the TITAN system.

This T_L159 model case serves as a proof of concept for the radiation-in-parallel scheme. However, we can already see that performance is only improved when the original model is approaching or at its limit

of scalability, and what we are doing with the radiation-in-parallel scheme is getting the model to run on fewer tasks where it is more efficient, leaving the remaining tasks to execute the radiation scheme. For this to work requires the cost of the radiation scheme to be sufficiently large to allow it to utilise a reasonable number of tasks, but not so large that the model is squeezed into too few tasks. A possible balance could be 20 to 30 percent of the total tasks for the radiation scheme, but this is very much dependent on whether the original model is at its limit of scalability and the computational cost of the radiation resolution being used. Another issue that must be considered with the radiation-in-parallel scheme is the need to perform runs to find a good balance between the number of model and radiation tasks, which is required to be done for every model and radiation configuration.

4.2 $T_L3999L137$ model performance

Table 3 summarises $T_L3999L137$ model runs on TITAN with and without the radiation-in-parallel scheme. The total number of tasks was 22,624 for all runs, each task using 8 OpenMP threads, so in total 212,992 AMD Interlagos cores were used. As the model time-step used was 240 seconds, and by default radiation is called every hour this equates to the radiation frequency $NRADFR=15$. The radiation resolution was fixed at a practical T_L2047 (half the spectral resolution of the model). What is clear from these results is that running the $T_L3999L137$ model with radiation called every time-step is very costly, and this extra cost cannot be recovered with the radiation-in-parallel scheme. Nevertheless, the results show that there is still a 6 percent performance improvement when comparing radiation computed every time-step ($NRADFR=1$). Of course we could increase the radiation resolution to that of the model, or possibly reduce the radiation frequency. There are disadvantages in both of these. Increasing the radiation resolution reduces the overall FD/D, while reducing the radiation frequency to say $NRADFR=2$ could have a negative effect on the quality of results. It should be noted that the **radparsim** and **radpar** code branches described in Section 2 support any value of $NRADFR$.

RADPAR_RUN	Model tasks	Radiation tasks	Radiation Frequency	FD/D
no	22,624	-	$NRADFR=15$	206
no	22,624	-	$NRADFR=1$	124
yes	18,432	8,192	$NRADFR=1$	132
yes	18,944	7,680	$NRADFR=1$	128
yes	19,456	7,168	$NRADFR=1$	130
yes	19,968	6,656	$NRADFR=1$	120

Table 3: $T_L3999L137$ model performance on TITAN in forecast days per day with and without radiation-in-parallel.

4.3 Colocating model and radiation threads on the same core

Within the EU funded CRESTA project we were interested in the performance to be gained by colocating a model and radiation thread on the same core. This is only possible on processor architectures that support Hyper-Threading (IBM calls this Simultaneous multithreading or SMT). Why would we do this? The reason is that Hyper-Threading provides the best performance when threads executing on the same core are using different resources of the core, for example, one thread is doing memory intensive operations such as loads or stores, and another thread on the same core is doing floating point operations. As ECMWF's CRAY XC-30 Ivybridge cores support Hyper-Threading, we worked with our CRAY partner

in CRESTA who provided the necessary scripting to produce the required MPICH.RANK.REORDER file that was needed to specify the detailed thread placement.

Two model cases were tested on the XC-30, a T_L159 case and a T_L3999 case.

For the T_L159 case, this ran on a single node using 8 tasks with 6 OpenMP threads per task, so 48 threads in total. Table 4 summarises the runs performed. Colocating model and radiation threads gives about a 5 percent performance improvement for the radiation-in-parallel scheme (RADPAR_RUN=yes).

For the T_L3999 case this ran on 512 nodes using 4,096 tasks with 6 OpenMP threads per task, so 48 threads per node. Table 5 summarises the runs performed where colocating model and radiation threads gives an 8 percent performance improvement for the radiation-in-parallel scheme (RADPAR_RUN=yes).

It should be noted that colocating model and radiation threads on the same core has only been tested in the case where an equal number of threads are being used for model and radiation. If there would be a difference in this number then the scheme would need to revert to the non-colocated threads running on cores by themselves, which could be less efficient if the number of such threads becomes significant.

RADPAR_RUN	Colocated threads	Model tasks	Radiation tasks	Radiation Frequency	FD/D
no	no	8	-	NRADFR=3	2,066
no	no	8	-	NRADFR=1	1,333
yes	no	4	4	NRADFR=1	1,349
yes	yes	4	4	NRADFR=1	1,423

Table 4: T_L159 model performance on a CRAY XC-30 in forecast days per day.

RADPAR_RUN	Colocated threads	Model tasks	Radiation tasks	Radiation Frequency	FD/D
no	no	4,096	-	NRADFR=15	70.1
no	no	4,096	-	NRADFR=1	35.1
yes	no	2,048	2,048	NRADFR=1	36.5
yes	yes	2,048	2,048	NRADFR=1	39.4

Table 5: $T_L3999L137$ model performance on a CRAY XC-30 in forecast days per day.

5 Discussion, conclusions and perspectives

The radiation-in-parallel configuration, by making radiation calculations at the same time as the rest of the model, has some potential for hiding the cost of these radiative calculations. From runs performed on TITAN, it is clear the gains are very much dependent on the relative cost of the radiation scheme to the rest of the model and whether the model has reached or is beyond its asymptotic performance (as was the case with the T_L159 model study). If radiation is called every hour for a high resolution model with a default radiation resolution (half the model resolution) then there is no performance gain by using the radiation-in-parallel scheme. The reason for this is due the frequency at which the radiation computations are called. In the default T_L3999 model configuration, radiation computations are called once an hour which is NRADFR=15. However, in the radiation-in-parallel configuration we have chosen to run with NRADFR=1 to obtain the best quality of results. This factor of 15 is too high to hide by running radiation computations in parallel with the model. It would be possible to run the radiation-in-parallel scheme just once an hour such that the radiation computations run in parallel with the model in the last step of each 15 model time-steps. However, this would be wasteful as the dedicated radiation tasks would be idle for

the remaining 14 model time-steps. In the future it may be possible to use a graph based approach (G. Bosilca et al., 2012) where tasks are scheduled dynamically, thus utilizing the idle resources.

Developments to reduce the cost of the radiation computations are being considered. If these are successful then the savings made could be used to perform the radiation computations more frequently than the one hour used today or alternatively increase the spatial resolution of the radiation grid. If the former is selected then it may be possible to use the radiation-in-parallel scheme with no noticeable effect on the quality of results. But will the radiation-in-parallel scheme be faster than the current scheme for a high resolution model? We suspect that this will very much depend on whether the model is close to or beyond its asymptotic performance for a given core count and the relative cost of the radiation computations to the rest of the model.

The radiation-in-parallel scheme serves as an example of how different model processes could be executed in parallel. The climate community is used to coupling models (atmosphere, ocean, sea-ice, chemistry, etc.) using an MPMD style of programming involving a coupler such as OASIS for exchanging fields which typically involves interpolation. These coupled models face the same issues as the radiation-in-parallel scheme in terms of how processor resources are distributed to the component models and often face constraints such as how the component models are parallelized (e.g. some may be OpenMP only, others may just use MPI). The approach taken in the radiation-in-parallel scheme uses a different approach where the same executable is used by model and radiation, and the decision of how many tasks to use for both is simply passed via an environment variable. This approach avoids the need for a coupler which could be expensive when coupling is required at each time-step.

Mogensen et al. (2012) describe a development at ECMWF where NEMO and IFS that were previously coupled with OASIS using an MPMD programming approach have been integrated into a single executable to share a common time-step loop. This development from a technical viewpoint has a lot in common with the radiation-in-parallel scheme, although the starting point was two separate applications rather than the IFS radiation scheme which is very much integrated within IFS in terms of the initialisation, time-stepping and interpolation of radiation input and output fields.

A question remains: what would the new approach do to the analysis results? Past experience when the frequency of the full radiation computations in the forward model used as trajectory for the analysis was changed from three hours to one hour did not indicate a particularly large impact. Given that the full radiation computations with the McRad model are only used in the trajectory calculations (and not in the adjoint calculations), it could be expected that a similar small impact on the analysis could result from such a future change.

The radiation-in-parallel method would have the built-in drawback of radiative heating rates being computed from the previous radiation time-step whereas the rest of the physical tendencies are sequentially computed from inputs derived from quantities at the current time-step.

The study presented here was meant to see whether the use of **radparsim** (which makes full radiation available one time-step earlier than the rest of the physics) creates unacceptable systematic decrease in forecast quality. Results presented in the previous sections with **radparsim** show that the effect is certainly not worse than the current temporal sampling of the radiation calculations. Furthermore it is shown that the effect decreases with increasing spatial resolution (decreasing time-step). Notably, the ECMWF model, by having a sequential approach to the computation of the effects of the physical processes, already has some "built-in temporal shift" between the different physical processes.

Finally this study shows that the impact of the lagged radiative computation is likely to be of the same order of magnitude or smaller than the error introduced by the current or near-future features of the

radiation scheme, namely:

- the McICA (the Monte-Carlo Independent Column Approximation, which was introduced in the IFS in June 2007 as part of the McRad package), which draws various cloud profiles to be distributed randomly in the various 112 g-points (the number of spectral elements) of the short-wave radiation scheme and the 140 g-points of the long-wave radiation scheme (Barker et al., 2003; Pincus et al., 2003; Morcrette et al., 2008);
- the McSI (Monte-Carlo Spectral Integration, presently tested by Alessio; Pincus and Stevens, 2009, 2013), which goes somewhat further, sampling g-points at each time-step and computing the fluxes from these restricted sets of spectral elements.

Furthermore, the success of the EPS system with its randomly perturbed physics is also proof that some large perturbations can be made to the diabatic heating rates without upsetting the overall model behaviour, so it can be thought that the proposed change in the way radiation computations could be handled in the future will be of no noticeable consequence to the quality of the forecasts.

With the perspective of even higher horizontal and vertical resolution in an exascale computer system, the model time-step is likely to further decrease, reducing further the impact of having such "radiation fields one time-step earlier than the rest".

6 Acknowledgements

This work has been supported by the CRESTA project that has received funding from the European Community's Seventh Framework Programme (ICT-2011.9.13) under Grant Agreement no. 287703. We would like to thank Alistair Hart (CRAY UK) for helping with the colocation experiments described in section 4.3. Finally, we thank Anton Beljaars, Alessio Bozzo, Mats Hamrud and Nils Wedi for their comments and suggestions.

References

- Barker, H.W., R. Pincus, and J.-J. Morcrette, 2003: The Monte-Carlo Independent Column Approximation: Application within large-scale models. *Proceedings GCSS/ARM Workshop on the Representation of Cloud Systems in Large-Scale Models*, May 2002, Kananaskis, AI, Canada, 10 pp.
- Dent, D., and G. Mozdzyński, 1996: ECMWF operational forecasting on a distributed memory platform: Forecast model. in *Making its Mark: Proceedings of the 7th Workshop on the Use of Parallel Processors in Meteorology*, G.-R. Hoffmann and N. Kreitz, Eds., 36-51.
- Mozdzyński, G., 2007: A new partitioning approach for ECMWF's Integrated Forecast System. *ECMWF Newsletter*, **114**, 17-23.
- Morcrette, J.-J., 2000: On the effects of the temporal and spatial sampling of radiation fields on the ECMWF forecasts and analyses. *Mon. Wea. 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. Wea. Rev.*, **136**, 4773-4798, doi: 10.1175/2008MWR2363.1

- Morcrette, J.-J., G. Mozdzyński and M. Leutbecher, 2008b: A reduced radiation grid for the ECMWF Integrated Forecasting System. *Mon. Wea. Rev.*, **136**, 4760-4772, doi: 10.1175/2008MWR2590.1
- Pincus, R., H.W. Barker, and J.-J. Morcrette, 2003: A fast, flexible, approximate technique for computing radiative transfer in inhomogeneous clouds. *J. Geophys. Res.*, **108D**, 4376, doi: 10.1029/2002JD003322.
- 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**, doi: 10.3894/JAMES.2009.1.1.
- Mogensen, K., Keeley, S., and Towers, P. 2012: Coupling of the NEMO and IFS models in a single executable. Technical Memorandum 673. April 2012.
- Pincus, R., and B. Stevens, 2013: Paths to accuracy for radiation parameterizations in atmospheric models. *J. Geophys. Res.*, , doi: 10.1029/2012JD ...
- G. Bosilca, A. Bouteiller, A. Danalis, T. Herault, P. Lemarinier, J. Dongarra, 2012: DAGuE: A generic distributed DAG engine for high performance computing, *Parallel Computing*, , Volume 38, Issue 1-2, pp. 37-51.

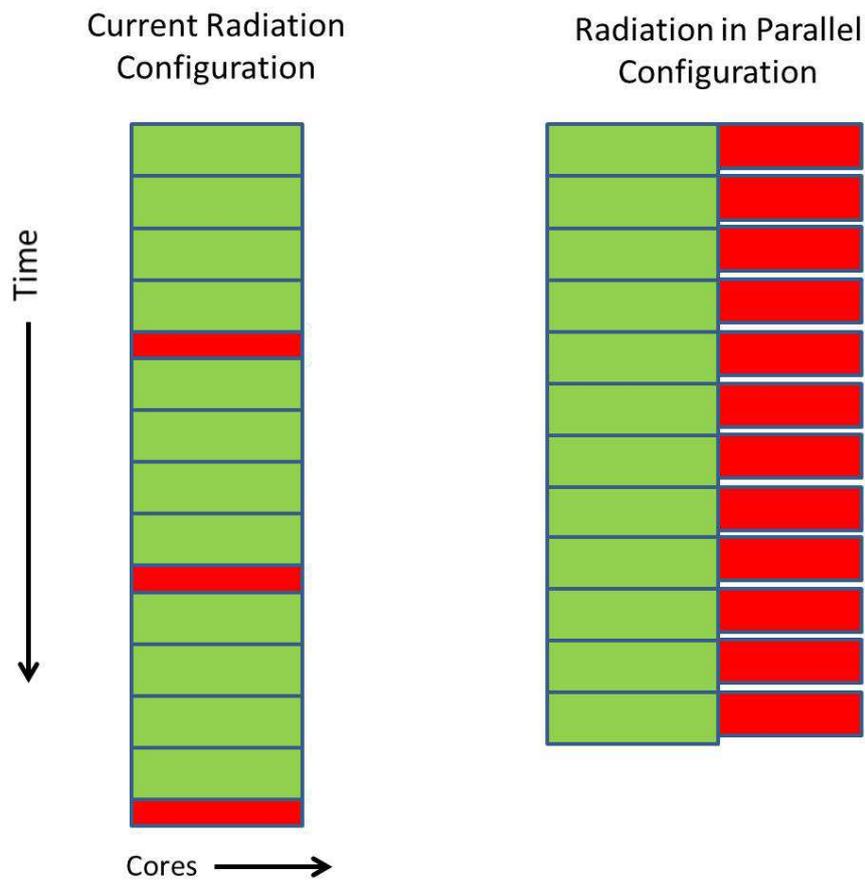


Figure 1: Graphical representation of current (left) and radiation-in-parallel (right) configurations. In the current configuration radiation transfer calculations (shown in red) are performed every NRADFR model (shown in green) time-steps, in this example case every 4 time-steps. In the radiation-in-parallel configuration radiation transfer calculations are performed at every timestep in parallel with the rest of the model.

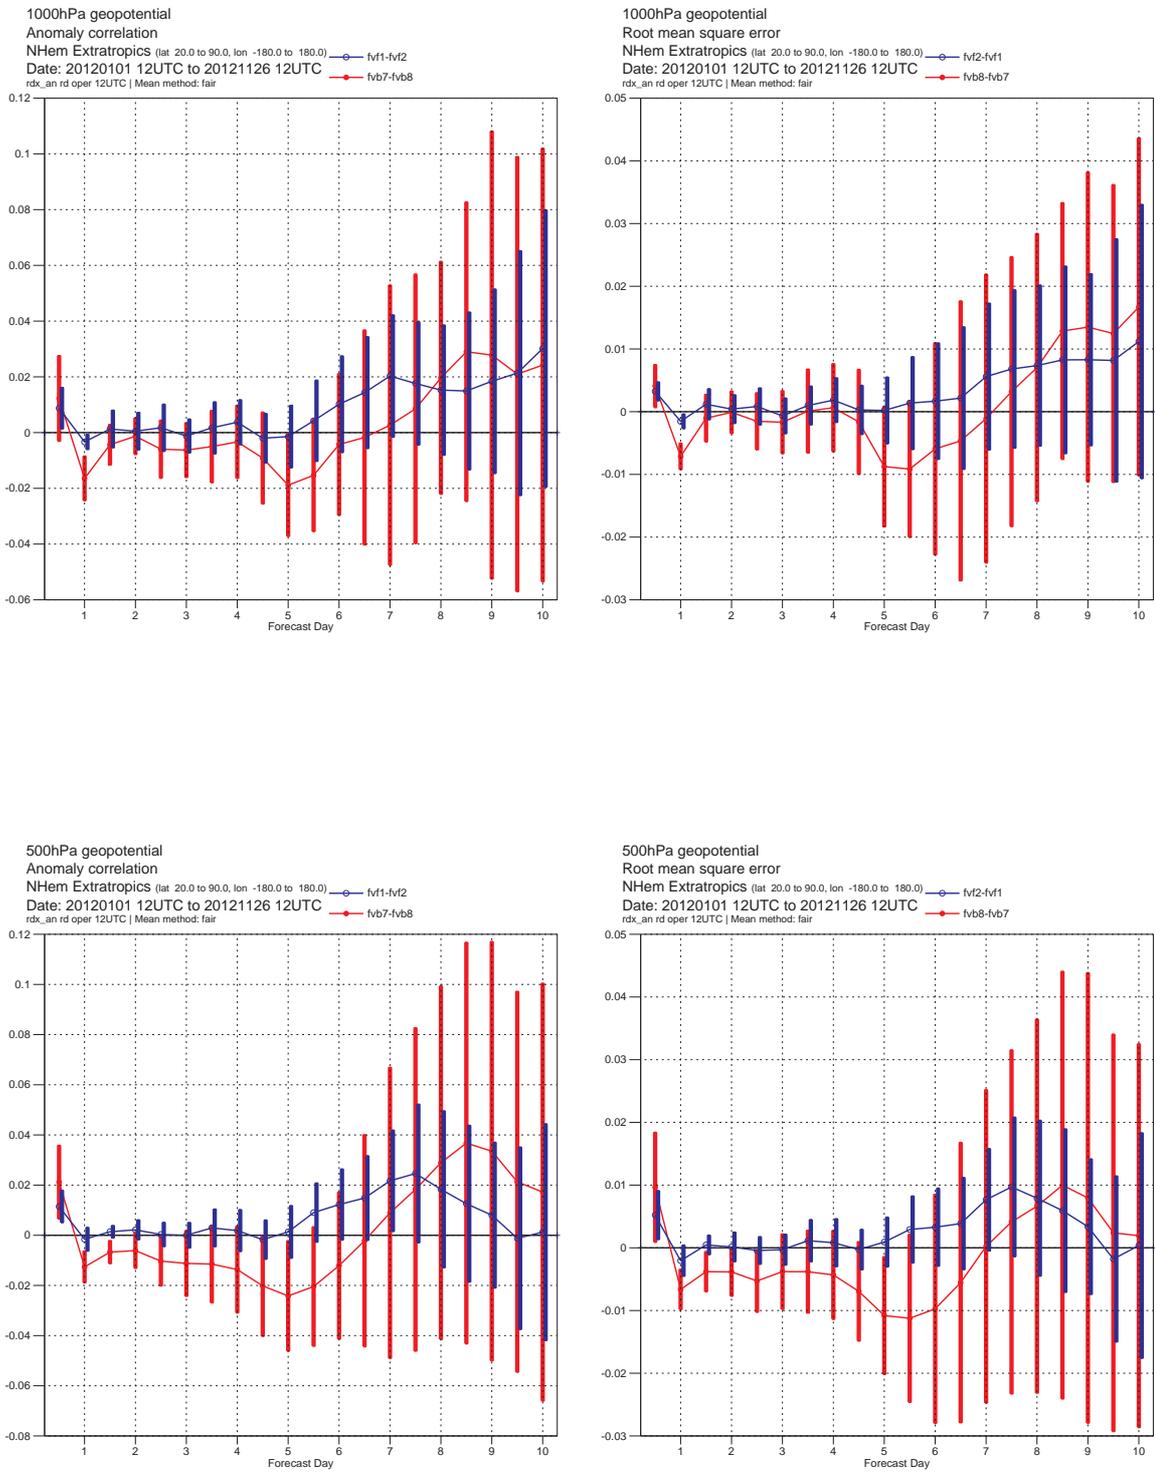


Figure 2: The difference in errors in geopotential at 1000 and 500 hPa for the T_L511 L91 **RPS1HR-REF1HR** (red), and the **RPSITS-REFITS** (blue) sets of forecasts, for the Northern hemisphere. Left panels are for the anomaly correlation, right panels for the r.m.s. error

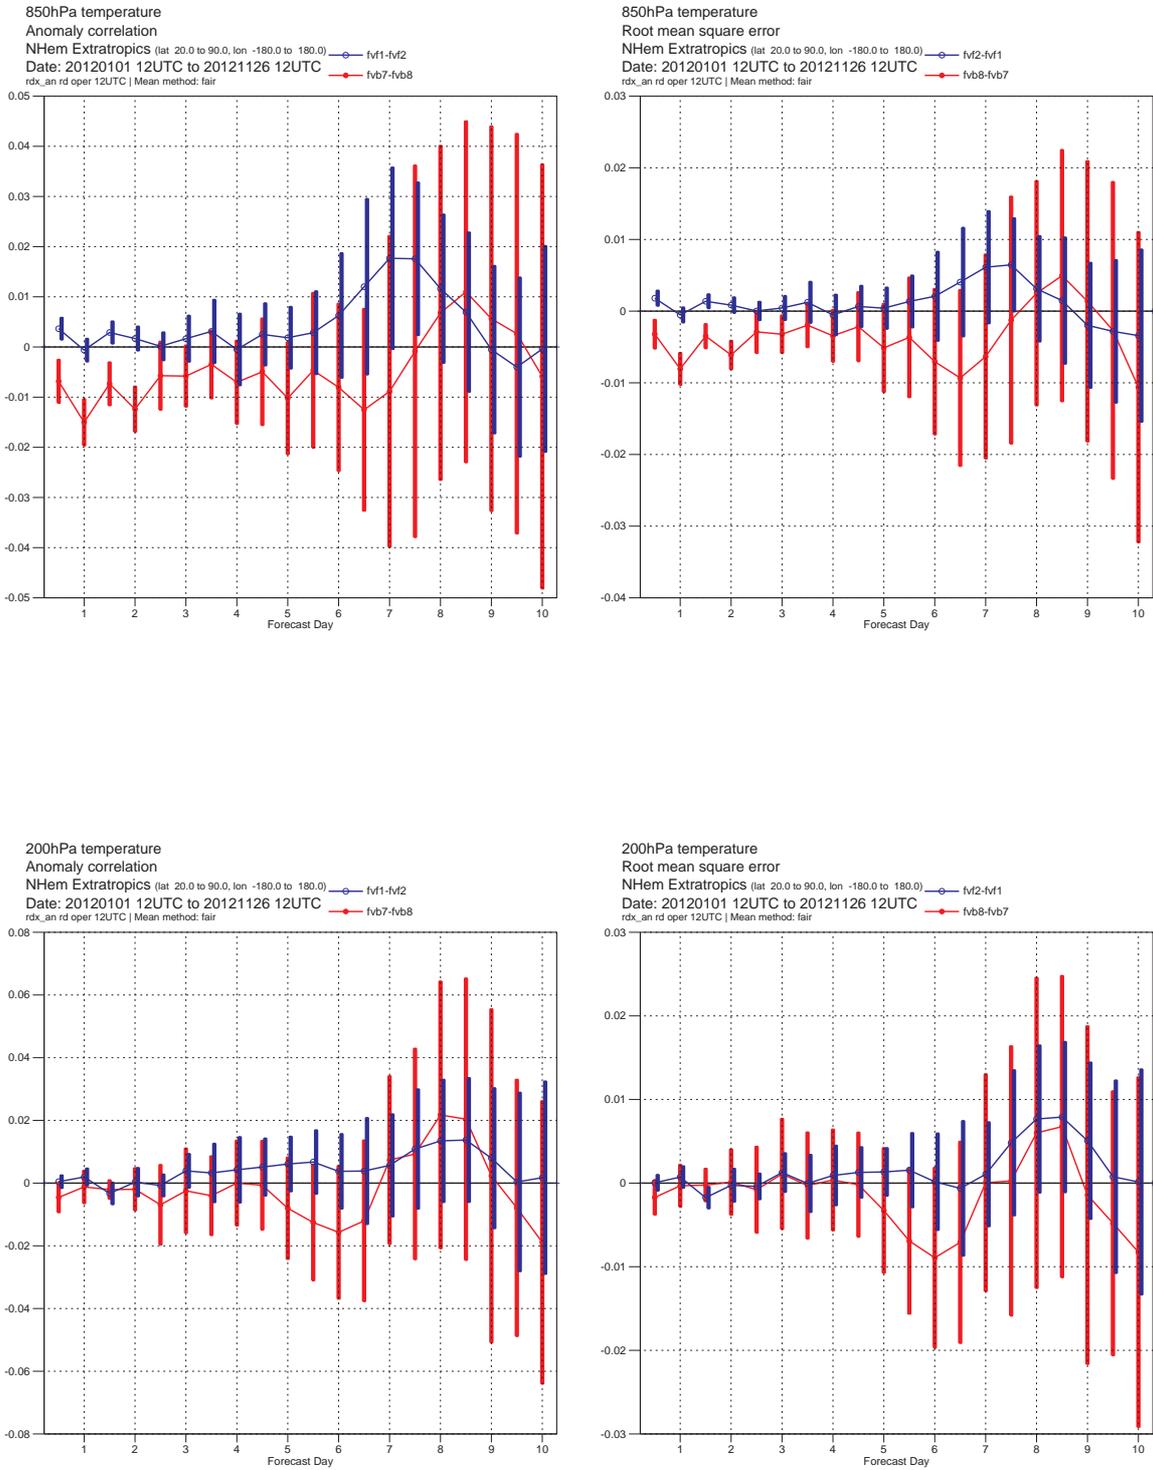


Figure 3: The difference in errors in temperature at 850 and 200 hPa for the T_{L511} L91 RPS1HR-REF1HR (red), and the RPSITS-REFITS (blue) sets of forecasts, for the Northern hemisphere. Left panels are for the r.m.s. error, right panels for the mean error

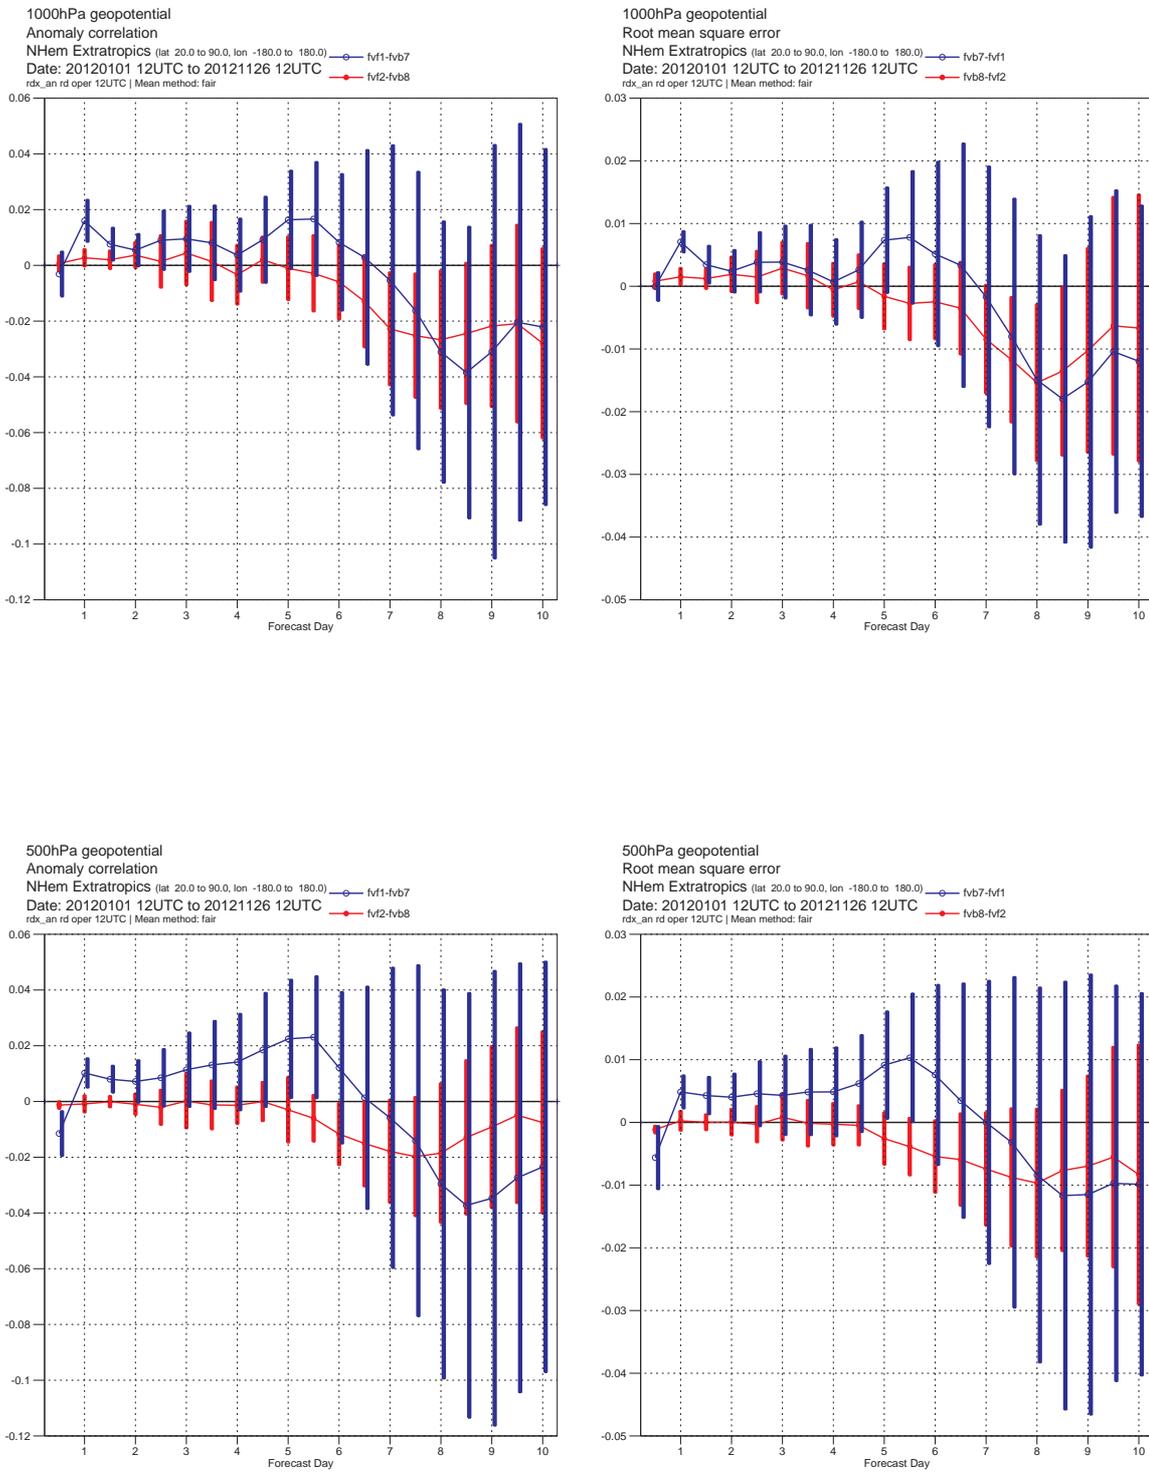


Figure 4: The difference in errors in geopotential at 1000 and 500 hPa for the T_{L511} L91 **REFITS-REFIHR** (red), and the **RPSITS-RPSIHR** (blue) sets of forecasts, for the Northern hemisphere. Left panels are for the anomaly correlation, right panels for the r.m.s. error

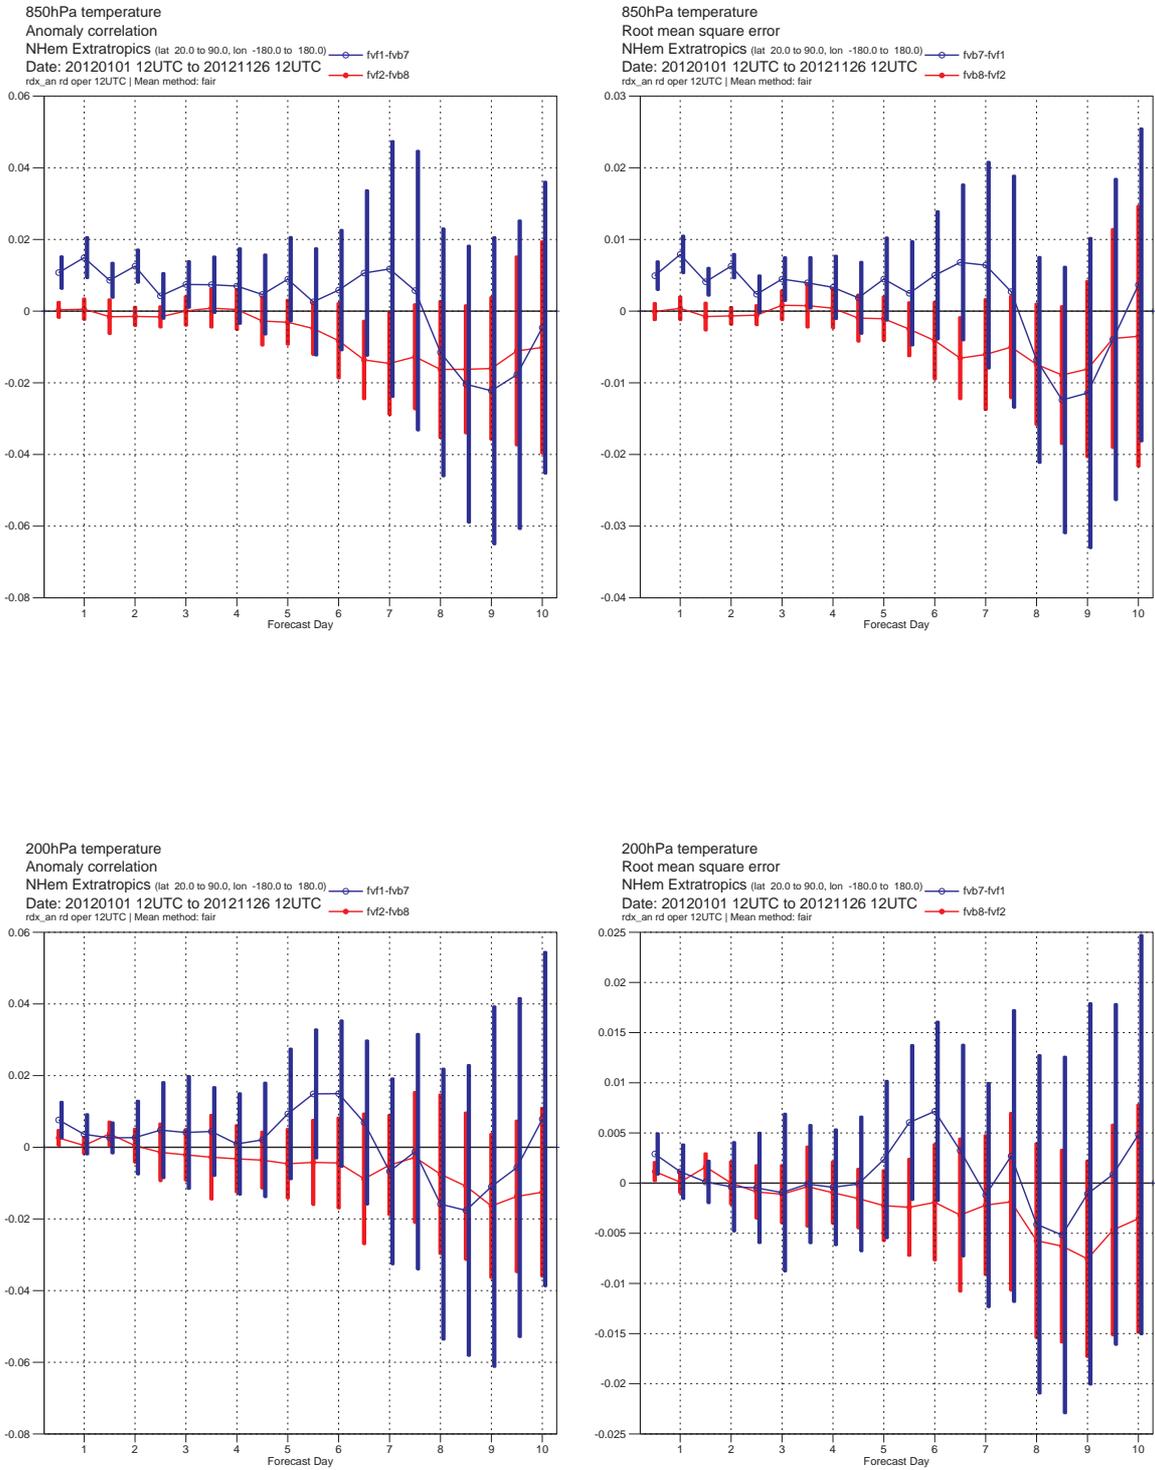


Figure 5: The difference in errors in temperature at 850 and 200 hPa for the $T_{L511} L91$ **REFITS-REF1HR** (red), and the **RPSITS-RPS1HR** (blue) sets of forecasts, for the Northern hemisphere. Left panels are for the r.m.s. error, right panels for the mean error

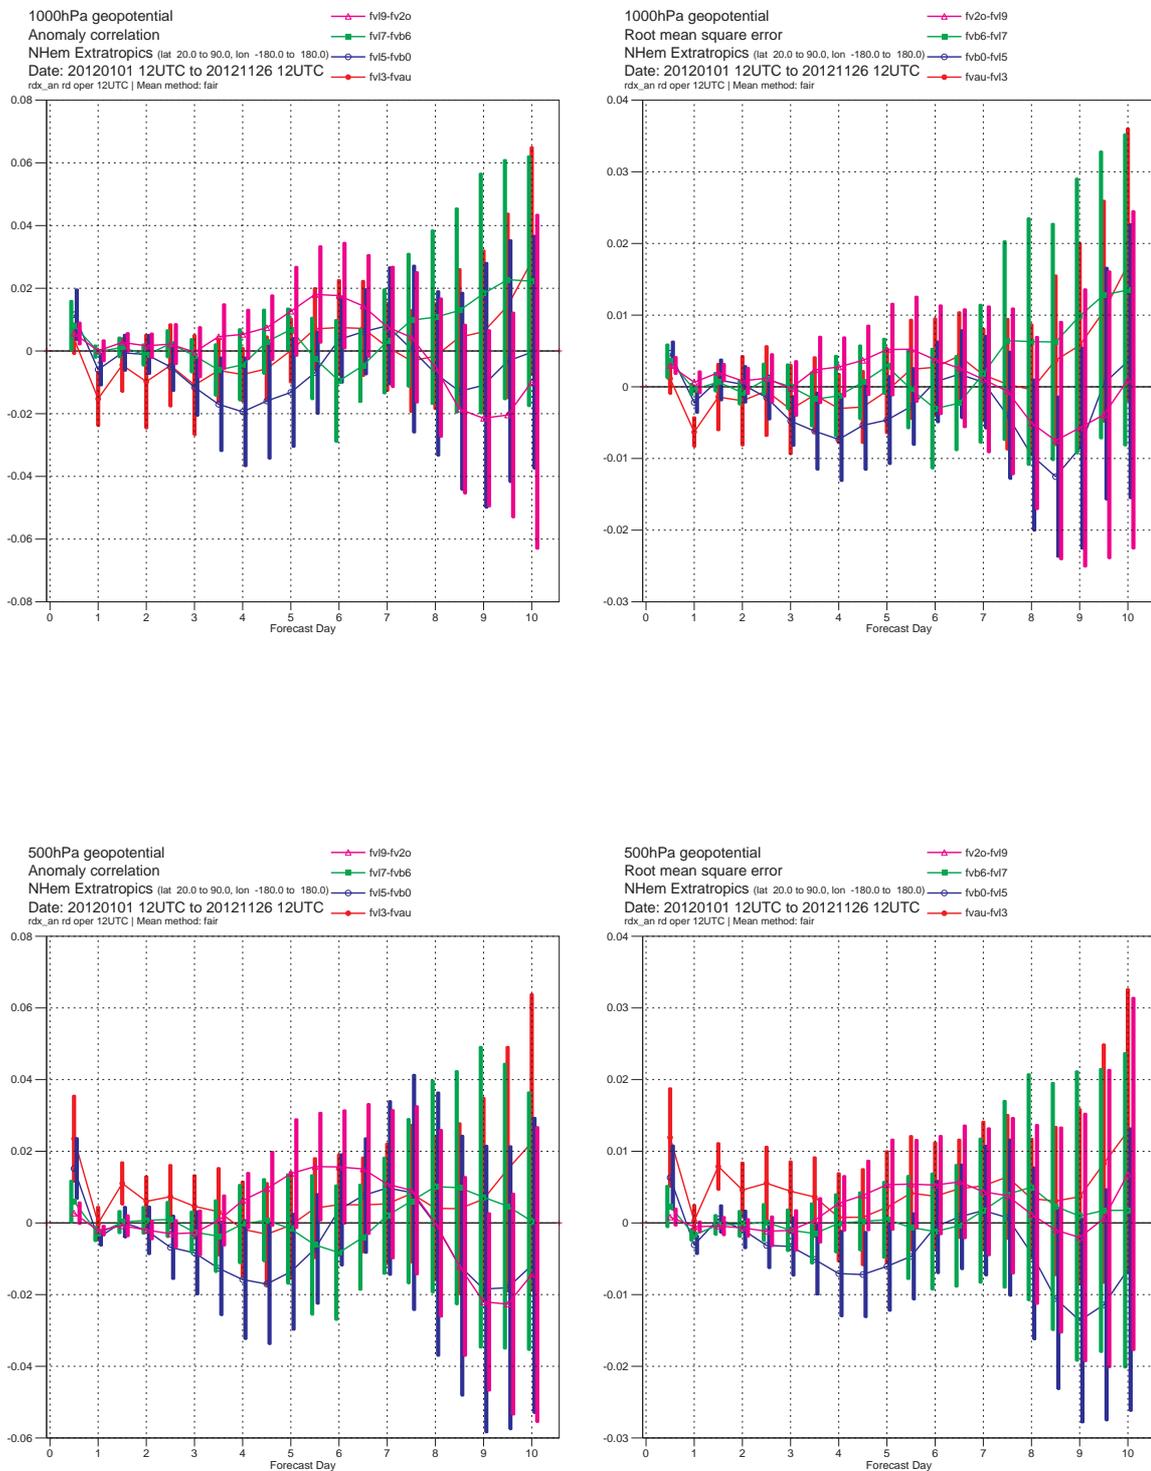


Figure 6: The difference in errors in geopotential height at 1000 and 500 hPa for the **RPSITS-REF1HR**, at T_L159 (red), T_L319 (blue), T_L639 (green), and T_L1279 (magenta) models at L91 vertical resolution. Results are for the Northern hemisphere. Left panels are for the anomaly correlation, right panels for the r.m.s error

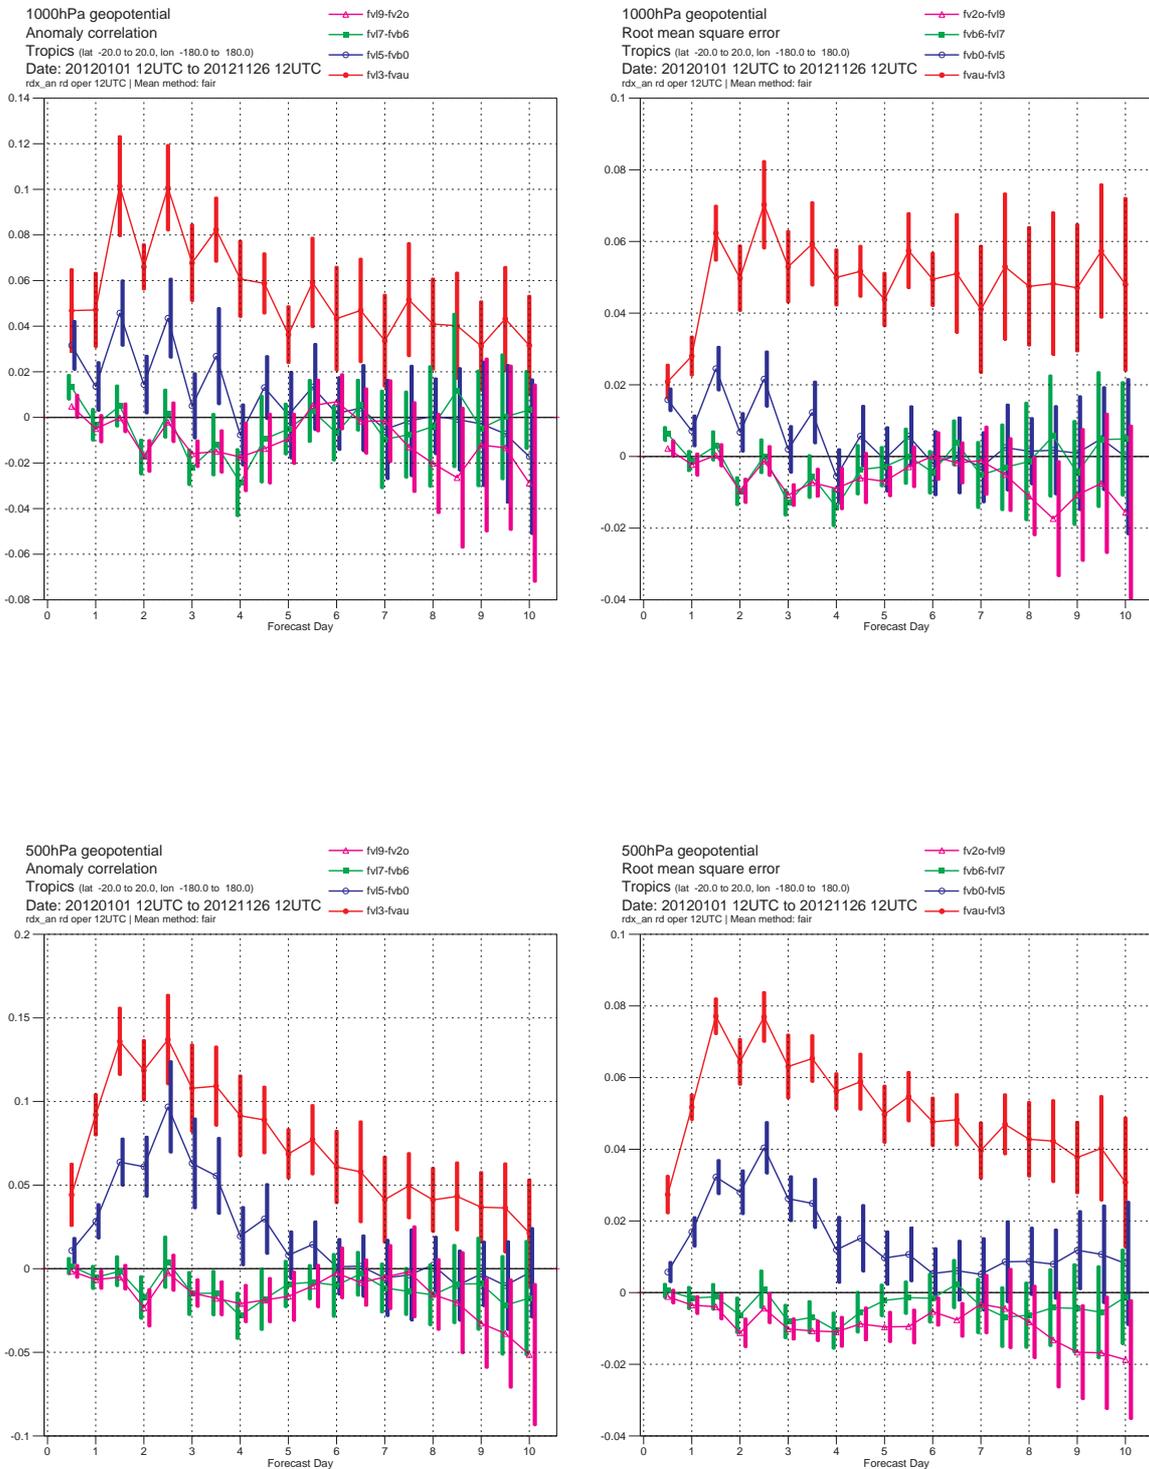


Figure 7: The difference in errors in geopotential height at 1000 and 500 hPa for the **RPSITS-REF1HR**, at T_L159 (red), T_L319 (blue), T_L639 (green), and T_L1279 (magenta) models at L91 vertical resolution. Results are for the Tropics. Left panels are for the anomaly correlation, right panels for the r.m.s. error

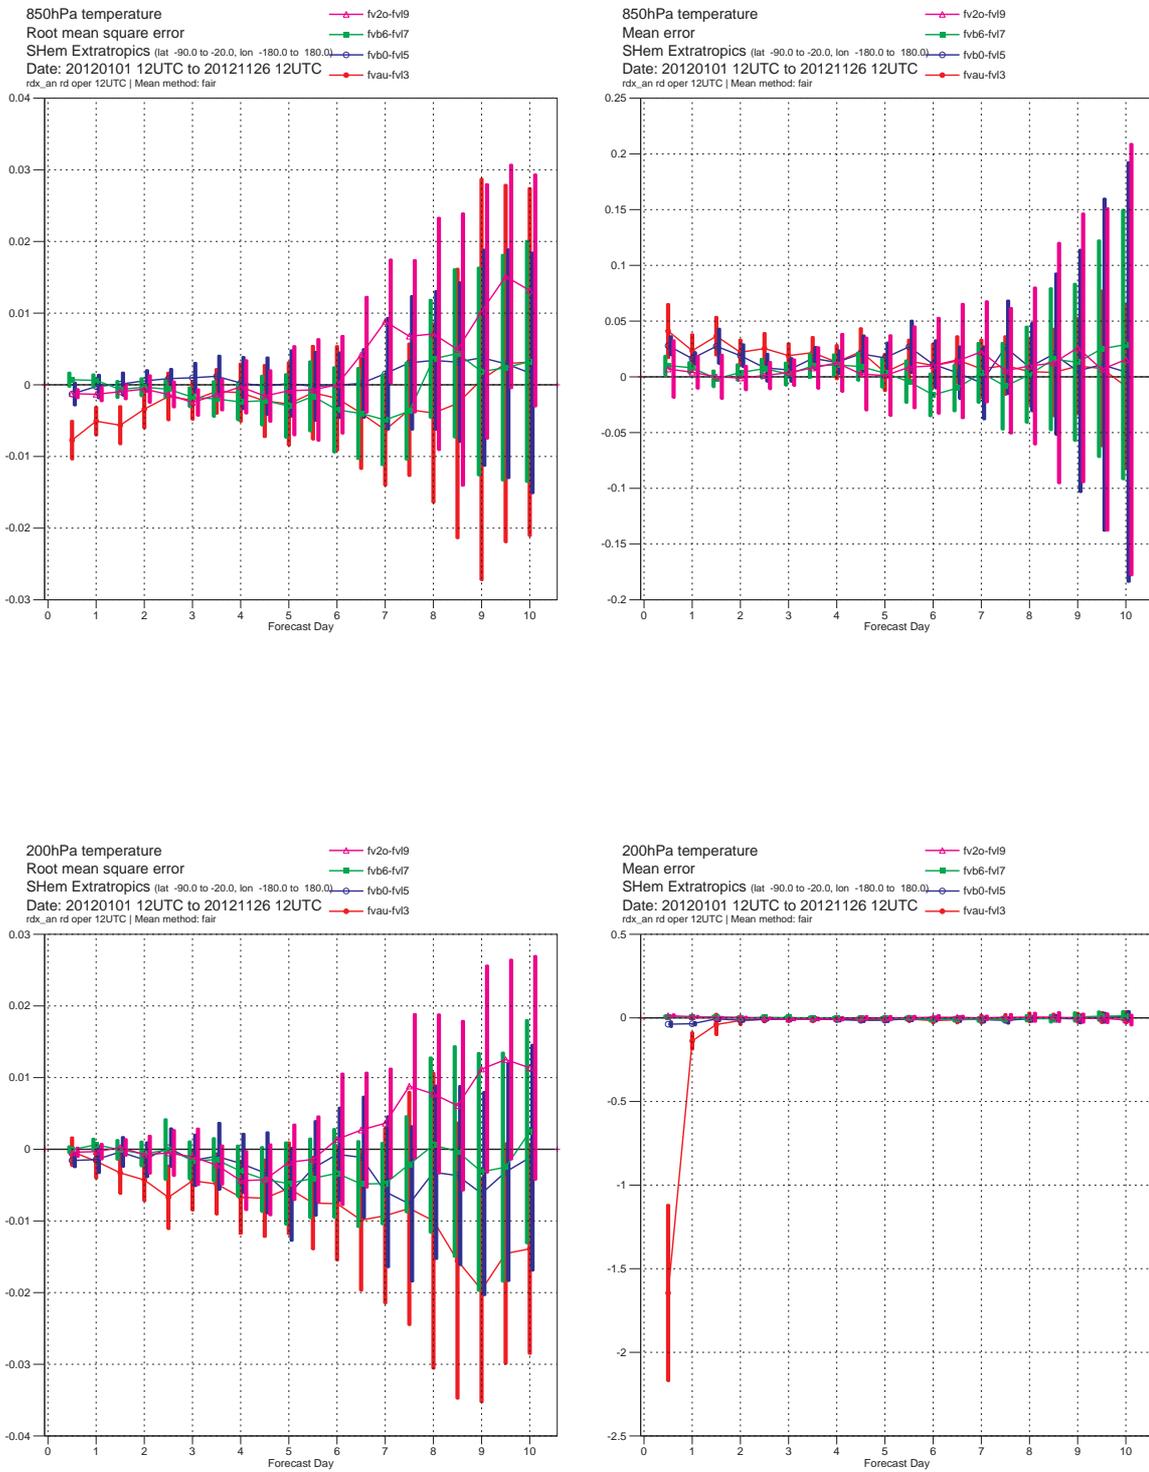


Figure 8: The difference in errors in temperature at 850 and 200 hPa for the **RPSITS-REF1HR**, at T_L159 (red), T_L319 (blue), T_L639 (green), and T_L1279 (magenta) models at L91 vertical resolution. Results are for the Southern hemisphere. Left panels are for the r.m.s. error, right panels for the mean error

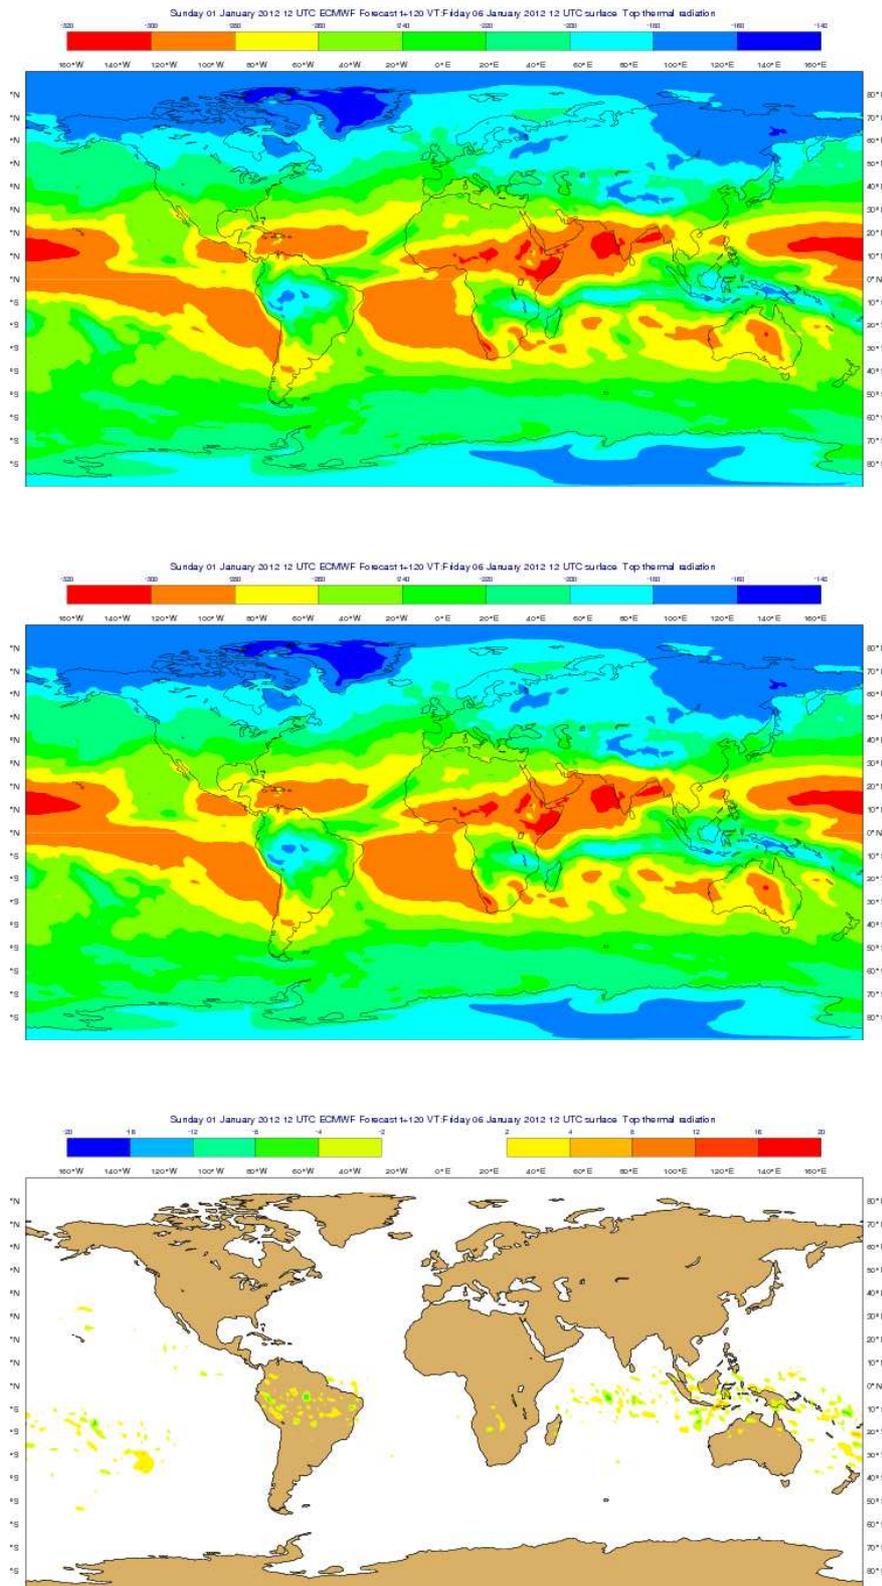


Figure 9: The outgoing long-wave radiation at the top of the atmosphere (in Wm^{-2}) averaged over the first five days of the three forecasts for 1, 11 and 21 January 2012. Top panel is for the control configuration **REF1HR**, middle panel is for the **radparsim** configuration called every model time-step **RPSITS**, and bottom panel is the difference **RPSITS-REF1HR**. Intervals are positive and negative, 2, 4, 8, 12, 16, 20 Wm^{-2} .

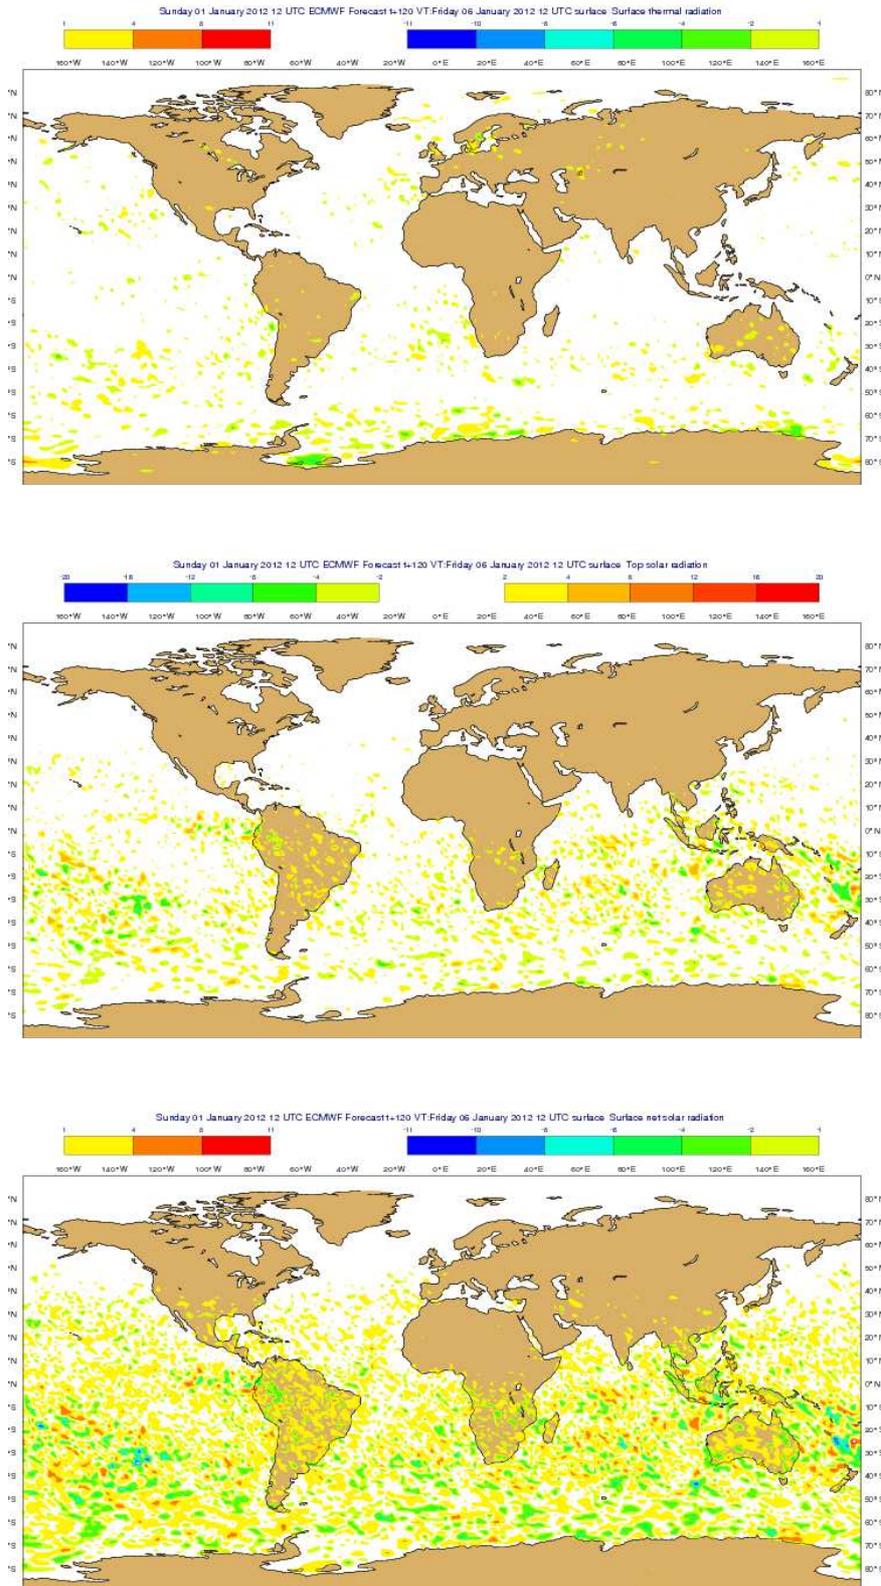


Figure 10: As in bottom panel of Figure 9, but for the net long-wave radiation at the surface (top panel), the net short-wave radiation at the top of the atmosphere (middle panel) and net short-wave radiation at the surface (bottom panel). All figures correspond to **RPSITS-REF1HR**. Intervals are positive and negative at 2, 4, 8, 12, 16, 20 Wm^{-2} for the top figure, at 1, 2, 4, 6, 8, 10, 11 Wm^{-2} for the bottom two figures.