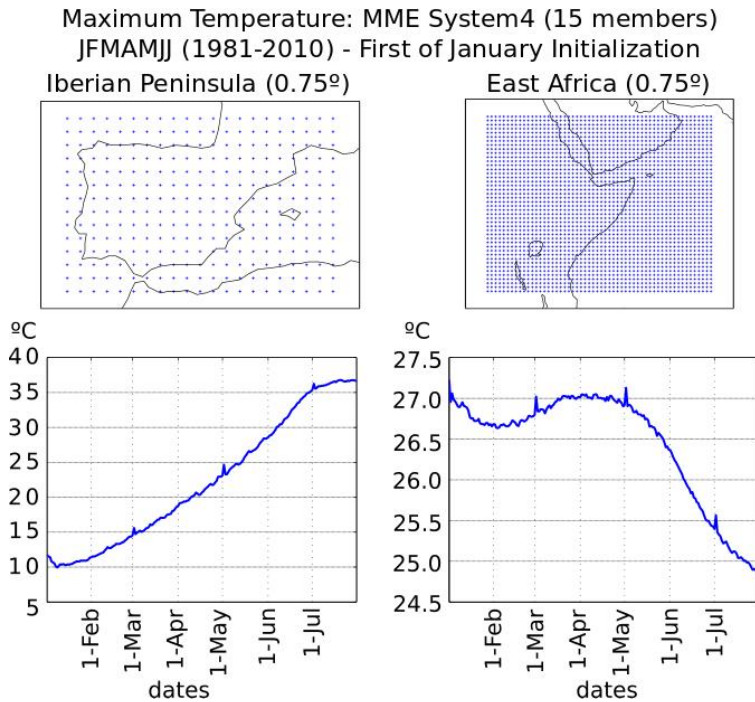


- [28-May-2013]. **Inhomogeneities in maximum temperature data from System4 hindcast:** The figure below shows 7-months long forecasts (January-to-July) of maximum temperature issued the first of January (thus leadtime = 0). Blue lines correspond to the multi-member (15 members) ensemble mean averaged over two different illustrative regions (the Iberian Peninsula on the left and East Africa on the right) for the thirty-years hindcast period 1981-2010. As can be seen, there is a systematic jump (of up to more than 1°C in the Iberian Peninsula) every two months since the initialization date (thus the 2nd of March, May and July). This also occurs for other initialization dates and locations (e.g. hemispheric averages).



Tim Stockdale, from ECMWF, confirmed this is a bug in the IFS cycle used in System4:

The cause of the problem is a bug in the way in which the IFS re-sets the arrays used to keep track of maximum/minimum temperatures and wind gusts, during restarts. It is a somewhat subtle issue, and although our previous version (System 3) used the same code for re-setting arrays, it does not seem to suffer the same problem. We use restarts when running the coupled model to ensure operational reliability and efficiency, and these can be made to give binary identical results over the restart point (i.e. the fact that a restart took place should be completely invisible). The bug means that the three fields referred to (Tmax, Tmin, max wind gust) are affected by the restart. The bug is still present in the latest version of the IFS, and we will ensure that a fix is made so that by the time we produce System 5 this problem does not exist.

The result of the bug is that the max/min arrays are not reset at the restart point. Thus, the value given for 0Z on the 2nd March is actually the maximum value over the previous 48h instead of over the previous 24h. Roughly half the time it will be correct (ie the max occurred in the previous 24h anyway) and half the time it will be too warm (because the daily max in the 24h was below the preceding daily max, but the model kept the higher value from the day before). Overall there is thus a warm bias. In all cases fields are physically reasonable, but obviously the statistics are wrong.