

RTTOV v12.1 Performance Test Log - Appendix to RTTOV v12 Test Plan (NWPSAF-MO-TV-038)

This document describes the tests performed to compare the performance of RTTOV v12.1 with that of RTTOV v11.3. This includes comparisons of execution speed and peak memory usage.

The direct, TL, AD and K models were run separately for the following test cases:

1. MW simulation (ATMS) without interpolation, 54L profile
2. MW simulation (ATMS) with interpolation, 101L profile
3. IR v7 predictor simulation (SEVIRI) with interpolation, 101L profile
4. Visible-only v9 predictor simulation (SEVIRI) with interpolation, 101L profile
5. All trace gas v9 predictor simulation (IASI) without interpolation, 101L profile
6. IR v7 predictor aerosol simulation (SEVIRI) with interpolation, 50L profile
7. IR v7 predictor cloud simulation (SEVIRI) with interpolation, 50L profile
8. PC-RTTOV simulation (IASI) without interpolation, PCscores only, 101L profile
9. PC-RTTOV simulation (IASI) without interpolation, with reconstructed radiances, 101L profile

The ATMS and SEVIRI coefficients are on 54L and the IASI coefficients are on 101L. In each case the profile surface type was set to sea with *calcemis* and *calcrefl* set to true for all channels. The following options/inputs were used:

- FASTEM6 for MW sea surface emissivity.
- ISEM for IR sea surface emissivity.
- All trace gas implies the six gases supported by RTTOV v11 (i.e. no SO₂).
- IR cloud scattering used Chou-scaling and the Baran 2014 ice scheme.
- Input units for cloud/aerosol profiles were those required by v11.3.
- The SEVIRI IR-only and ATMS simulations were run for all channels.
- The SEVIRI visible-only simulations were run for channels 1-3.
- The IASI clear-sky simulations were run for 183 channels.
- PC-RTTOV was run with ipcreg 1 (300 predictor channels) and 100 PC scores.
- PC-RTTOV reconstructed radiances were calculated for 300 channels.

All options were selected to make the simulations as similar as possible between v11.3 and v12.1.

1. Speed comparisons

Tests were performed on an Intel desktop (Nehalem) using ifort v16.0.1 and gfortran v4.4.7, a Cray XC-40 (Ivybridge) using the Cray Fortran compiler v8.3.4 and a BULL cluster (Broadwell) using ifort v16.1.150.

The tests were run for a large number of profiles, shown in Table 1.

Test case	Model(s)	Number of profiles
1	direct/TL/AD/K	50000
2	direct/TL/AD/K	50000
3	direct/TL/AD/K	50000
4	direct/TL/AD/K	50000
5	direct/TL/AD K	5000 1000
6	direct/TL/AD/K	50000
7	direct/TL/AD/K	50000
8	direct/TL/AD K	1000 100
9	direct/TL/AD K	1000 100

Table 1: Number of profiles used for each timing test.

All tests were run with one profile passed to RTTOV per call and the results are shown in Tables 2 (Intel desktop) and 3 (supercomputer platforms). In addition test case 1 for ATMS was run for 50 profiles per call: the results are shown in Tables 4 and 5. Timings were taken from the test suite and the timing results are shown as *ms per profile*.

The colour-coding is as follows:

- green => v12.1 more than 10% faster than v11.3
- yellow => v12.1 5-10% faster than v11.3
- orange => v12.1 5-10% slower than v11.3
- red => v12.1 more than 10% slower than v11.3

The optical depth calculation for v9 predictors has undergone optimisation in v12 and the impact can be seen in the IASI tests (5, 8 and 9) at least on the Intel desktop and Cray platforms. For the MW and clear-sky SEVIRI IR simulations there is a slight increase in run-time with v12. The timings are broadly the same for the visible clear-sky and IR cloudy simulations. The IR aerosol simulations show a slightly larger increase in run-time with v12 than the other simulation types.

The results are broadly in line with expectations. We will investigate the slight increase in run-time for clear-sky MW and IR simulations and IR aerosol simulations with a view to addressing these in the next minor update of v12.

PC-RTTOV vs HT-FRTC

Table 6 compares PC-RTTOV and HT-FRTC using ifort v12.0.4 on an Intel desktop. The HT-FRTC direct and K model runs are simulating 1000 and 10 profiles respectively in each call. Since HT-FRTC reads its coefficient data in each call it is much more efficient to process many profiles at once than to loop over profile individually. For PC-RTTOV the direct and K model simulations were repeated 1000 and 100 times for a single profile respectively (note the timings in the table are *per profile* so are directly comparable between the models). The difference in run-time for per-profile or multiple-profile simulations is not very large for PC-RTTOV, but it has a relatively large memory footprint which limits the maximum number of profiles which can be run at one time (especially for the K model). It is planned to investigate and improve the performance of the HT-FRTC K model.

RTTOV interpolator

Tables 7 and 8 show the impact of the RTTOV interpolator on execution time using ifort v12.0.4 on an Intel desktop. The ATMS test with a 54L profile to provide a fair comparison with/without interpolation. Results for each interpolation mode are shown and also results when pressure is an active TL/AD/K variable. (NB Interpolation mode 2 should not be used for TL/AD simulations so these are omitted from the tables - see the user guide). Note that this represents more or less the worst case since for more expensive simulation types the interpolation takes a relatively smaller proportion of the total runtime.

2. Memory comparisons

Peak memory usage was measured using valgrind's massif tool. Tests were performed on an Intel desktop using the gfortran compiler. Table 9 shows the results of the tests. The colour-coding is as follows:

- green => peak memory for v12.1 90% or less than that for v11.3
- yellow => peak memory for v12.1 90-95% than that for v11.3
- orange => peak memory for v12.1 105-110% than that for v11.3
- red => peak memory for v12.1 110% or more than that for v11.3

For the IASI tests the required channels were read in from the full HDF5 coefficient files. The memory usage is dominated by the coefficient reading for these tests as these coefficient files are very large. The IASI tests were repeated after extracting the coefficients for the required channels for each test to separate coefficient files. The results are shown in Table 10.

Again, the results are broadly in line with expectations. The modest increases in memory requirements for the clear-sky simulations are not considered cause for concern: these are at least partly due to an increase in the size of the coefficients structure. In the case of instruments like IASI we would usually expect users with performance-critical applications to extract the channel subset they wish to use to a separate coefficient file and the results show v12 uses less memory in this scenario. Similar considerations apply to the IR scattering case: the v12 cloud/aerosol coefficient files contain additional parameters for the DOM solver and some additional memory is required internally for the optional cloud/aerosol unit conversion. If the number of profiles passed in to each RTTOV call increases then the relative difference in memory usage between v11 and v12 decreases so the relative differences shown in the tables below represent more or less the worst case.

Test case	Model	Intel gfortran v11.3	Intel gfortran v12.1	Intel gfortran v12.1:v11.3	Intel ifort v11.3	Intel ifort v12.1	Intel ifort v12.1:v11.3
	Direct	0.41	0.43	1.05	0.278	0.282	1.01
1	TL	0.692	0.732	1.06	0.513	0.53	1.03
MW, no interp	AD	0.762	0.884	1.16	0.589	0.663	1.13
	K	1.021	1.158	1.13	0.853	0.869	1.02
	Direct	0.684	0.707	1.03	0.438	0.447	1.02
2	TL	1.131	1.185	1.05	0.8	0.818	1.02
MW, interp	AD	1.316	1.46	1.11	0.919	0.987	1.07
	K	1.825	1.98	1.08	1.271	1.327	1.04
	Direct	0.315	0.344	1.09	0.217	0.232	1.07
3	TL	0.548	0.596	1.09	0.41	0.432	1.05
IR v7 clear	AD	0.623	0.708	1.14	0.468	0.514	1.10
	K	0.851	0.948	1.11	0.62	0.655	1.06
	Direct	0.493	0.513	1.04	0.52	0.54	1.04
4	TL	0.85	0.84	0.99	0.818	0.842	1.03
Visible clear	AD	0.937	0.99	1.06	0.854	0.904	1.06
	K	1.177	1.218	1.04	1.009	1.031	1.02
	Direct	7.404	6.202	0.84	4.636	4.144	0.89
5	TL	13.044	10.452	0.80	8.91	7.338	0.82
All gas clear	AD	13.562	11.898	0.88	10.982	8.538	0.78
	K	37.04	37.76	1.02	29.39	30.57	1.04
	Direct	0.282	0.318	1.13	0.216	0.236	1.09
6	TL	0.509	0.574	1.13	0.406	0.435	1.07
IR aerosol	AD	0.577	0.674	1.17	0.454	0.502	1.11
	K	0.842	0.957	1.14	0.638	0.69	1.08
	Direct	0.637	0.672	1.05	0.402	0.399	0.99
7	TL	1.133	1.196	1.06	0.763	0.795	1.04
IR cloud	AD	1.274	1.389	1.09	0.949	0.982	1.04
	K	1.689	1.817	1.08	1.122	1.156	1.03
	Direct	12.78	10.88	0.85	8.37	7.38	0.88
8	TL	22.39	18.06	0.81	15.6	13.17	0.84
PC	AD	23.83	20.16	0.85	19.25	15.04	0.78
	K	74.1	74.3	1.00	62.7	63.1	1.01
	Direct	12.97	10.83	0.84	8.24	6.72	0.82
9	TL	22.53	18.33	0.81	15.86	12.38	0.78
PC, rec. rads	AD	23.96	20.42	0.85	19.91	14.77	0.74
	K	113.4	114	1.01	98.7	93.4	0.95

Table 2: Speed test results, 1 profile per call, Intel desktop. Timings are ms per profile.

Test case	Model	Cray v11.3	Cray v12.1	Cray v12.1:v11.3	BULL v11.3	BULL v12.1	BULL v12.1:v11.3
	Direct	0.159	0.169	1.06	0.156	0.171	1.10
1	TL	0.291	0.317	1.09	0.301	0.335	1.11
MW, no interp	AD	0.34	0.348	1.02	0.361	0.387	1.07
	K	0.488	0.496	1.02	0.569	0.615	1.08
	Direct	0.241	0.26	1.08	0.235	0.252	1.07
2	TL	0.456	0.481	1.05	0.461	0.498	1.08
MW, interp	AD	0.519	0.533	1.03	0.538	0.577	1.07
	K	0.722	0.752	1.04	0.848	0.895	1.06
	Direct	0.129	0.147	1.14	0.126	0.142	1.12
3	TL	0.24	0.26	1.09	0.247	0.267	1.08
IR v7 clear	AD	0.274	0.286	1.04	0.285	0.307	1.08
	K	0.359	0.383	1.07	0.433	0.467	1.08
	Direct	0.206	0.218	1.06	0.23	0.248	1.08
4	TL	0.353	0.368	1.04	0.406	0.42	1.03
Visible clear	AD	0.371	0.388	1.05	0.44	0.45	1.02
	K	0.459	0.474	1.03	0.55	0.567	1.03
	Direct	2.762	2.05	0.74	1.846	1.856	1.01
5	TL	5.286	3.994	0.76	3.68	3.736	1.02
All gas clear	AD	6.972	4.26	0.61	5.37	4.678	0.87
	K	22.16	17.36	0.78	14.71	15.64	1.06
	Direct	0.13	0.147	1.13	0.143	0.16	1.12
6	TL	0.244	0.274	1.12	0.277	0.313	1.13
IR aerosol	AD	0.278	0.294	1.06	0.326	0.372	1.14
	K	0.39	0.401	1.03	0.475	0.568	1.20
	Direct	0.261	0.286	1.09	0.26	0.277	1.06
7	TL	0.489	0.501	1.02	0.531	0.565	1.06
IR cloud	AD	0.598	0.579	0.97	0.638	0.653	1.02
	K	0.726	0.758	1.04	0.85	0.89	1.05
	Direct	5.47	4	0.73	3.47	3.47	1.00
8	TL	10.55	7.71	0.73	6.64	6.92	1.04
PC	AD	12.66	8.8	0.70	9.25	8.29	0.90
	K	35.9	30.9	0.86	39.7	39.3	0.99
	Direct	4.95	4.03	0.81	3.24	3.35	1.03
9	TL	9.83	7.77	0.79	6.49	6.6	1.02
PC, rec. rads	AD	13.73	8.19	0.60	9.1	8.13	0.89
	K	51.7	49.3	0.95	77.1	76.1	0.99

Table 3: Speed test results, 1 profile per call, supercomputer platforms. Timings are ms per profile.

Test case	Model	Intel gfortran v11.3	Intel gfortran v12.1	Intel gfortran v12.1:v11.3	Intel ifort v11.3	Intel ifort v12.1	Intel ifort v12.1:v11.3
	Direct	0.361	0.376	1.04	0.223	0.234	1.05
1	TL	0.607	0.639	1.05	0.432	0.431	1.00
MW, no interp	AD	0.678	0.789	1.16	0.507	0.548	1.08
	K	0.973	1.105	1.13	0.781	0.851	1.09

Table 4: Speed test results, 50 profiles per call, Intel desktop. Timings are ms per profile.

Test case	Model	Cray v11.3	Cray v12.1	Cray v12.1:v11.3	BULL v11.3	BULL v12.1	BULL v12.1:v11.3
	Direct	0.124	0.136	1.10	0.106	0.118	1.11
1	TL	0.249	0.256	1.03	0.211	0.236	1.12
MW, no interp	AD	0.298	0.289	0.97	0.262	0.284	1.09
	K	0.545	0.527	0.97	0.462	0.508	1.10

Table 5: Speed test results, 50 profiles per call, supercomputer platforms. Timings are ms per profile.

Model	Simulated output	PC-RTTOV	HT-FRTC	HT-FRTC:PC-RTTOV
Direct	PC scores	3.66	7.97	2.18
	Rec. rads	3.57	8.05	2.25
K	PC scores	37.90	2657.00	70.11
	Rec. rads	60.50	2676.00	44.23

Table 6: Speed test results comparing PC-RTTOV and HT-FRTC. Timings are ms per profile.

Interp mode	Model	No interp	Interp	Interp:No interp
1	Direct	0.275	0.315	1.14
	TL	0.513	0.591	1.15
	AD	0.644	0.740	1.15
	K	0.881	1.043	1.18
2	Direct	0.275	0.306	1.11
	TL	N/A		
	AD			
	K			
3	Direct	0.275	0.314	1.14
	TL	0.513	0.586	1.14
	AD	0.644	0.722	1.12
	K	0.881	1.024	1.16
4	Direct	0.275	0.334	1.21
	TL	0.513	0.627	1.22
	AD	0.644	0.773	1.20
	K	0.881	1.077	1.22
5	Direct	0.275	0.327	1.19
	TL	0.513	0.612	1.19
	AD	0.644	0.770	1.19
	K	0.881	1.055	1.20

Table 7: Speed test results directly comparing each interpolation mode with no interpolation. Test carried out for 54L ATMS coefficients with a 54L profile. Timings are ms per profile.

Interp mode	Model	No interp	Interp	Interp:No interp
1	TL	0.509	0.621	1.22
	AD	0.645	0.821	1.27
	K	0.880	2.045	2.32
3	TL	0.509	0.598	1.17
	AD	0.645	0.776	1.20
	K	0.880	1.658	1.88
4	TL	0.509	0.676	1.33
	AD	0.645	0.876	1.36
	K	0.880	2.091	2.38
5	TL	0.509	0.656	1.29
	AD	0.645	0.834	1.29
	K	0.880	1.720	1.95

Table 8: As for Table 7, but with pressure as an active variable in the TL/AD/K (i.e. lgradp=true).

Test case	Model	v11.3 peak memory (MB)	v12.1 peak memory (MB)	v12.1:v11.3
	Direct	1.666	1.745	1.05
1	TL	1.665	1.745	1.05
MW, no interp	AD	1.665	1.745	1.05
	K	2.301	2.437	1.06
	Direct	1.665	1.745	1.05
2	TL	1.836	1.984	1.08
MW, interp	AD	1.854	2.002	1.08
	K	2.832	2.968	1.05
	Direct	1.088	1.168	1.07
3	TL	1.343	1.461	1.09
IR v7 clear	AD	1.349	1.468	1.09
	K	1.718	1.832	1.07
	Direct	1.224	1.292	1.06
4	TL	1.446	1.559	1.08
Visible clear	AD	1.444	1.562	1.08
	K	1.699	1.818	1.07
	Direct	620.2	708.5	1.14
5	TL	620.2	708.5	1.14
All gas clear	AD	620.2	708.5	1.14
	K	620.2	708.5	1.14
	Direct	1.091	1.505	1.38
6	TL	1.28	1.727	1.35
IR aerosol	AD	1.283	1.731	1.35
	K	1.651	2.143	1.30
	Direct	1.503	1.935	1.29
7	TL	2.033	2.459	1.21
IR cloud	AD	2.022	2.458	1.22
	K	3.55	4.018	1.13
	Direct	684.5	708.4	1.03
8	TL	684.5	708.4	1.03
PC	AD	684.5	708.4	1.03
	K	684.5	708.4	1.03
	Direct	659.2	708.4	1.07
9	TL	659.2	708.4	1.07
PC, rec. rads	AD	659.2	708.4	1.07
	K	659.2	708.4	1.07

Table 9: Memory test results.

Test case	Model	v11.3 peak memory (MB)	v12.1 peak memory (MB)	v12.1:v11.3
	Direct	27.38	17.5	0.64
5	TL	27.38	17.5	0.64
All gas clear	AD	27.38	17.5	0.64
	K	50.15	46.93	0.94
	Direct	111.7	105.1	0.94
8	TL	111.7	105.1	0.94
PC	AD	111.7	105.1	0.94
	K	142.1	137.4	0.97
	Direct	82.3	68.85	0.84
9	TL	82.3	68.85	0.84
PC, rec. rads	AD	82.3	68.85	0.84
	K	92.73	88.03	0.95

Table 10: IASI memory test results with coefficient extraction.