

IFS/Arpège Memorandum

From: Peter Lean (ECMWF)

To: (ECMWF) DR, RD Division & Section Heads

To: (Météo-France) Arpège diffusion list

To: (ALADIN) Piet Termonia

To: (HIRLAM) Daniel Santos-Muñoz

File: RD17-xxx

Subject: Minutes of the IFS/Arpège coordination meeting held on 12 June 2017.

Participants:

Météo-France: Claude Fischer, Ryad El Khatib (physically in Reading) François Bouyssel, Karim Yessad, Stéphane Martinez (remote by video-conference)

ECMWF: Stephen English, Deborah Salmond, Olivier Marsden, Peter Lean, Tomas Wilhelmsson (only morning)

ALADIN: Daan Degrauwe (remote by video-conference)

HIRLAM: Daniel Santos-Muñoz (remote by video-conference, only morning)

1. Adoption of Agenda

adopted

2. Approval of Minutes of meeting of 18 October 2016

approved

3. Review of list of actions from last meeting

1. Deborah to start listing those features from FORTRAN 2008 that should stay out of the IFS/Arpège codes (blacklist approach), after CY45. Then circulate this list in MF and with Aladin/Hirlam

Agreed that all Fortran 2008 features would be allowed with the exception of co-arrays. Other specific features could be added to the Exceptions list later, after a due investigation about how crucial or general a reported portability issue would affect platforms and compilers. ACTION CLOSED.

There was further discussion about how it would be useful to provide notes on experiences using different compiler versions for each cycle. Daan suggested it would be useful to record problems experienced with specific compilers and to note which compilers were used for each cycle by EC and MF. It was decided that the cycle release notes may be the best place for this as wiki pages tend not to stay up to date.

ACTION on Tomas W and Ryad to record details of which versions of the compilers were used in the release notes for each cycle. Possibly start this with CY45 FLUBs.

2. Action on Sylvie or Deborah: send the technical and scientific note about the results with the mass-preserving option in IFS to MF, when ready.

Sylvie sent her preliminary note on the subject. If further documentation becomes available then it will also be sent to Claude. ACTION CLOSED.

3. Action on Claude: to provide EC a description of MF's tests and results with LSPRT in assimilation. Further check with Etienne and other MF assimilation staff which fixes from testing Arpège/Arome VAR could enter CY45 (in Toulouse or via EC/Deborah).

- i) Note on LSPRT testing: Gerald Desroziers, Loïk Berre and Etienne Arbogast prepared a note describing problems experienced using LSPRT=.true. in the assimilation. Issues related to the conversion between temperature and virtual temperature impacted not just the model but also the assimilation. In the upcoming esuite LSPRT is set to false for 4D-Var in Arpège and AROME. Claude noted that they are keen to get further feedback on this issue from EC. ACTION CLOSED.

New ACTION on Steve to ask Sébastien Massart to send details of experiences related to LSPRT and the humidity fields in the assimilation at 45r1.

- ii) Fixes for data assimilation have gone into pre-CY45 (v6). A list of these fixes will be provided to EC within the next week. ACTION CLOSED.

4. Action on Steve: to provide MF with a feedback, possibly a work plan description, about the “alpha control vector” work planned at EC. Note: there was a visit by Sébastien Massart to Toulouse in May.

Sébastien had a successful visit to Toulouse in May. Claude noted that MF were keen for this exchange of information to continue and to collaborate, particularly on methods for localisation in hybrid data assimilation. Claude noted that MF would prefer that this work be conducted within the OOPS framework. The exchange of information should continue and the scientific collaboration is highly encouraged, but no specific action is required. ACTION CLOSED.

5. MF to provide EC trace-back information about the SRTM problem/crash encountered in some Arpège PEARP members (this is under investigation in GMAP)

Claude reported that under certain situations, an out of bounds memory access can occur resulting in the crashes. It is hoped that this will be resolved within the next few days. ACTION STILL OPEN.

6. MF (Claude and Yann Michel) to contact Elias Holm for informing about the suggested implementation of a specific REDNMC_Q namelist coefficient in CY45T1.

This has not yet been done. ACTION OPEN.

4. MF information about progress and plans of E-suites and cycles

François Bouyssel presented Météo France's plans for the upcoming E-suites and cycles.

The current Arpège e-suite CY42_op2 contains a number of developments:

- ⇒ Substantial increase in the overall numbers of observations used due to:
 - more channels used for some instruments (e.g. MHS and GMI),
 - reduced thinning (in IASI and GEO radiances)
 - revised whitelist.
- ⇒ Upgrades to the convection scheme and surface exchange (new schemes PCMT and SURFEX).
- ⇒ SST coupling with 1d ocean model.
- ⇒ MERCATOR SST product is now used in favour of OSTIA SSTs as it contains more information at fine scales.

It was noted that the e-suite is ongoing and likely to be delayed until September due to mixed results in forecast impact at short lead times as verified against ECMWF analysis. However, the impact at longer lead times is largely positive.

The next e-suite will be CY43T2 and is likely to start in autumn 2017 with an aim of operational implementation in mid-2018. This includes a number of changes:

- ⇒ Migration to the VORTEX python toolbox for Arpège, EDA and Arome.
- ⇒ Use of GRIB2 for output post-processing files
- ⇒ Higher resolution in Arpège: TL1798c2.2L105.
- ⇒ Physics:
 - Tunings in convection scheme
 - Prognostic graupel in microphysics
- ⇒ Introduce VarBC for aircraft data.
- ⇒ EPS to be run x4 per day (instead of current x2)
- ⇒ Single precision for some forecast tasks (to be confirmed)

MF's new HPC is planned for 2019.

There was a discussion about the extent to which OOPS would be included in the benchmarking for the next HPC at both MF and EC. Steve and Olivier reported that no OOPS benchmark is in RAPS 16, but there should be something in RAPS17 (likely a single minimisation, not full 4D-Var). RAPS17 is likely to be ready by the end of the summer.

5. EC information about progress and plans of E-suites and cycles

Steve presented ECMWF's latest progress and plans. He reported that there had been no operational issues associated with the cubic octahedral grid. The problems associated with the stability of the TL had turned out to be related to the cumulus downdraught linearized physics and not related to the cubic grid as had previously been suspected. As the TL stability issues have now been fixed, experimentation on using the cubic octahedral grids in the inner loop could resume. However,

computational affordability was likely to prohibit inner loops using the cubic octahedral grid from going into operations.

The current e-suite is CY43R3. Highlights include:

- ⇒ Wavelet filtering of EDA humidity variances leads to around 3% reduction in the standard deviation of first guess departures for tropical wind observations (through the tracer effect).
- ⇒ Introduction of new radiation scheme ECRAD: Currently, the new scheme replicates the functionality of the old scheme, but runs more efficiently resulting in considerable cost savings. Experimentation is under way using the new science enabled by the new scheme.
- ⇒ Improved tropical cyclone initialisation (which mainly resolve the problems experienced at 43r1).
- ⇒ Use of new microwave humidity sounding data (SAPHIR, GMI)
- ⇒ Microphysics changes which avoid unrealistic changes in precipitation intensity over land.

The next e-suite will be CY45R1 and has a target of operational implementation in Jan/Feb 2018:

- ⇒ HRES partial coupling with NEMO:
 - In the tropics: full coupled
 - Extra-tropics: only partial coupling (increments added to OSTIA SSTs up to day 4) as the 0.25 degree ocean model is not capable of resolving the boundary currents (especially the Gulf Stream).
 - Gives strong improvement in tropical forecast scores, including a reduction in error for tropical cyclones.
 - Computational expense: originally costs an extra 5 minutes to run which was deemed unaffordable. However, optimisations in the coupling along with others in the physics led to a 2 minute saving meaning that the extra 3 minute required is now considered affordable given the strong forecast benefit.
- ⇒ Rayleigh Friction (LRFRIC)
- ⇒ Upgrade RTTOV to v12:
 - Overall the impact looks good (especially verified against humidity sensitive radiances and AMVs)
 - However, a degradation is evident when looking at radiances from the temperature sounding instruments. It is thought that there is a problem with the coefficient files for CrIS. It is hoped that using the old CrIS coefficient files will solve these issues otherwise RTTOV 12 will have to be pulled from the cycle.
- ⇒ OCEAN5 sea ice will replace OSTIA.

Work continues on overlapping 12h window 4D-Var with implementation potentially in CY45R2. The main development to enable this was to separate the definitions of the first guess from the background. This development also enables a new “soft re-centring” option in the EDA.

6. Updated status of build of CY45; validation of assimilation at MF in CY43 and CY44

Claude reported difficulties in the validation of Arpège /Arome data assimilation.

Through collaboration with EC there was a recent breakthrough allowing CY43T2 assimilation to run with OpenMP parallel screening activated. The results look acceptable, but more time is required to complete the scientific validation of Arpège 4D-VAR and screening as well as Arome 3D-VAR. Furthermore, several instances of array bounds violations have been detected with new arrays being passed through argument lists in Arpège/Arome specific options of the observation operators (APACHE, ACHMT).

Some substantial code changes are needed in CANARI (Patrick Moll) for compatibility with calls to observation handling code that was changed by the recent re-factoring effort of CY43.

CY44 is built, but will not be fully validated for data assimilation at MF. The next scientific validation will be for CY45 and CY45T1.

There have been exchanges of code in the pre-phasing for CY45 (currently at v6). Further changes in v7 are expected in the next few days including changes from Olivier. Technically, it builds and runs, but no scientific validation has been done yet.

Issues obtaining bit-identical results between CY44 and CY45. MF are not able to do this yet for running forecast models with physics turned on. This is not yet understood, but there is a suspicion that it could be related to the compiler (e.g. perhaps de-optimisation of new code leads to slightly different results?).

It is hoped that CY45 can be declared by June 19th, however if the bit-reproducibility issues are not resolved or understood then it may be delayed with a hard deadline of June 26th.

Deborah gave an update on CY45 progress at ECMWF. CY45 minimisation does not bit-compare with CY44 due to an issue with compiler optimisation in the diffusion TL which leads to small differences when optimisation is above -O0. A long running experiment is in progress to confirm that the change has no meteorological impact.

Further testing has been done at CY45 comparing IFS with OOPS in a minimisation with/without physics and with/without VarBC.

7. HIRLAM comments

None.

8. ALADIN comments

Daan asked when the next opportunity would be for LAM partners to provide contributions. Claude said it would be CY45T1 (around September).

9. Specific issues:

9.1. OOPS Progress

Deborah presented an update on progress with OOPS

VarBC: Progress has been made implementing VarBC in OOPS. There is more work required to split the VarBC class to reduce memory usage. Scientific validation is still required to understand if the use of a different preconditioner in OOPS is a problem or not. Roel will remain at ECMWF until the end of the summer.

Duplicate routine removal: Work is required to remove several duplicate routines which have been created to allow OOPS to run. Maintaining two copies of the code is dangerous with a high risk of code divergence. To achieve this the routines will need to be adapted so that they can work with the two different flavours of the trajectory object. This could be achieved either using optional arguments, or the abstract class features of Fortran 2008. Olivier will discuss these options at the next technical video-conference.

Jo-table: Peter has completed work to add the diagnostic Jo-table to OOPS. This involved packaging the existing diagnostic Jo table data from the YOMCOSJO module into a new object passed by argument. The Jo-table summarizes Jo categorised by observation type and observed variable (information which is not available to the OOPS C++ layer). A small change to the C++ is required to call a new printJo() method to perform this categorisation. This change will be included in CY45R1.

Namelist splitting: Some options in IFS namelists change in each outer loop. New work is required to split some namelists to extract those options which change into configuration files which can be used by OOPS.

NEXRAD observation operator: Phillippe Lopez is making good progress removing the observation operator for the rainfall composite products from the model physics for compatibility with OOPS. This will put this operator in line with all other operators. Similar work was undertaken several years ago by Alan Geer for the all-sky operators.

OOPS layer changes: at CY45 major changes were required in the OOPS abstract layer, involving passing the model object to areas where conceptually it might at first glance seem that it was not appropriate to do so. For example, in IFS there doesn't seem to be a way of creating a model state

without using knowledge contained in the model object. Further discussion is needed to see if this can be avoided, but no solutions are readily apparent.

Handling of time: Olivier described changes he plans to make in the handling of time variables in IFS/OOPS potentially use new date-time objects. The work has not yet started, but Claude was interested to hear more on these plans (perhaps at the next video-conference). It is likely also that exchange with Karim at MF would be important in order to encompass many of the various uses of time within IFS/Arpège and the LAMs applications.

FcKit: changes were made in IFS to use the EcKit/FcKit logger so that OOPS and IFS log output could go to the same output stream. This touched on a wider issue related to the use of external libraries (such as FcKit) which do not yet have frozen APIs. Ryad and Olivier noted that using external libraries without a stable API presented new challenges for those building new cycles. Fortunately up until now, the external libraries used by IFS had all been rather stable and so there was rarely a need to install new versions for each cycle. The increasing use of the various libraries under ifs-support (several of which are not yet mature and so do not yet have stable APIs) has increased the complexity of the system.

A particular issue related to MF developments on top of grib_api needs to be resolved. ACTION on MF (Claude and Philippe) to contact EC (Deborah for IFS/Arpège and Shahram Najm for grib_api) and co-ordinate to ensure that MF contributions enter the main grib_api/eccodes package.

It was further agreed that EC and MF should communicate well about the target support libraries and their versions, required for compiling a new IFS/Arpège version. This probably should be done well ahead of the build of a new release, and should include any code not jointly released in the Reading or Toulouse NWP main trunks (eg OOPS/C++, ECKIT/FCKIT etc.).

Claude presented MF progress with OOPS

- ⇒ Prototypes are now working with CY43 code + branches
- ⇒ 3DVar is running with all observation types.
- ⇒ Results are not bit-identical between Arpège and OOPS, but they look reasonable.
- ⇒ Work is underway porting the ensemble DA prototypes to the CY43 (it is hoped to include these in the next benchmark suite by the end of the year)
- ⇒ Ryad has done work with the FullPos object and a first proposal for OOPS-IFS interfaces is ready with CY45

9.2. Single precision IFS

Experiments at ECMWF using single precision in the forecast show that it runs around 40% faster. More work is required to do a full scientific validation and prove that it doesn't degrade the forecasts.

There are no plans to implement single precision in the data assimilation inner loop, but discussions are on going whether it may be possible to implement in initial, and possibly 4DV outer loop, trajectories.

There is no specific timescale for when single precision forecasts will be implemented in an operational suite at ECMWF. However, the savings from single precision were one of the assumptions made in plans to meet the 2025 Strategy and they have also been assumed in the business case for the next HPC, so it is safe to assume that it will become operational at some point.

9.3. ATLAS, ESCAPE/dwarves: link with evolution of IFS

Atlas: is included in the IFS source code behind `#ifdefs`. Deborah reported that initial tests showed that more optimisation is required and so it is not active in 45r1 and remains in the domain of research.

Filip Vana has had some success using Atlas in a research experiment where he ran IFS entirely in grid-point space (without using any spectral transforms). This would not have been possible without Atlas.

ESCAPE/dwarves: Claude asked whether there were plans for developments/optimisations made using the dwarves to eventually go back into IFS. Deborah confirmed that this had always been the plan. However, Olivier reported that sometimes the dwarves were found to not be representative of the full IFS. He cited an example where an optimisation had been very effective in a dwarf, but when used in IFS actually caused a slow-down (this case was with the cloud scheme, possibly due to differences in cache usage).

For the time being, there was no concrete plan to implement codes or changes coming from specific dwarves back into the IFS.

9.4. development status of ECRAD

Steve confirmed that ECRAD is in CY43R3 (current e-suite). The new code replicates the functionality of the old radiation scheme, but runs faster. Robin is doing scientific validation of the new features enabled by ECRAD (such as representation of some 3D radiative effects). Robin is also keen to provide a standalone release of this software outside of the IFS repository.

9.5. status of ODB, ODB1 versus ODB2, COPE

Peter gave an update on the latest progress with COPE. This mainly relates to porting the BUFR pre-screening tasks (non-background dependent) used at ECMWF to run quasi-continuously in the SAPP system. Further progress has been made to allow IFS to be initialized by odb2 files (work that was enabled by the new “ifsobs” interface layer). As odb_api is not a parallel database library, the parallel IO and redistribution of the odb2 data needs to be handled explicitly in the ifsobs layer.

At CY45R1, the interface to the ODB column data in the observation operator will be upgraded. After 45r1, to access data from an ODB column in IFS the following syntax will be used:

VAL=ROBODY%BODY%OBSVALUE(IROW)

i.e. <view_name>%<table_name>%<column_name> (row)

As this involves many hundreds of lines of code changes, it is anticipated that the merge at CY46 may be difficult and so a python script is available to automatically upgrade code changes from CY43Tx to the new interface (this will be run by ECMWF during early pre-phasing).

10. Content and timing of cycles

MF will upgrade their HPC in spring 2019. François suggested that a best case scenario would see the first cycle on the new HPC being CY46T1, but with a worst case of CY45T1. Regarding the timing of the OOPS implementation, two scenarios were discussed:

- i) “Early OOPS”: move to next HPC with OOPS
 - a. Include OOPS in CY46, or have a quick CY47 soon after CY46 (Oct 2018)
- ii) “Late OOPS”: move to next HPC without OOPS
 - a. Implement OOPS some time after CY47.

This choice needs to be revisited for further discussion when the progress with OOPS is further along. ACTION on all.

ECMWF will face a similar decision regarding the timing of OOPS and the move to the next HPC. At present the preferred cycle for OOPS is 46r1 but there will be further discussion at the OOPS Project Board and also the RD management meeting, and furthermore with FD, before a consensus can be reached.

Joint cycle	ECMWF	MF	Start of phasing	Declaration	Misc. / Oper plans
CY43			September 2015	February 2016	Declared 25 Feb.
	CY43R1		March 2016	June 2016	Scientific changes
	CY43R2		May 2016	?	Re-factoring for OOPS
	CY43R3		October 2016	November 2016	Model + DA
		CY43T1	April 2016	June 2016	Including Aladin and Hirlam
		CY43T2	October 2016	mid-November	Wrap-up of bugfixes from [CY40-CY42], as well as MF E-suite

					changes from CY42_op1/op2
CY44			mid- November 2016	End of February 2017	The build process of this cycle might be in multiple steps to accommodate necessary input for OOPS in IFS. Tbd in forthcoming video-conferences.
	<i>CY44R1</i>				<i>Dropped</i>
		<i>CY44T1</i>		<i>Cancelled in order to build the technical OOPS cycle CY45</i>	<i>Dropped</i>
CY45			March 2017	June 26 th (at latest)	MODEL object re- factoring
		CY45T1	2nd October 2017	End of November 2017	Including Aladin and Hirlam
	CY45R1		May 31 st 2017	August 2017	Science tbc Operational Jan/Fab `18
	CY45R2		July 2017	October 2017	12h overlapping DA (op. Jul `18)
CY46			Start Jan 15 th , 2018	End of March 2018	
		CY46T1	May-June or Sept-Oct 2018		Timing depending on schedule of CY47
	CY46R1				OOPS?
CY47			Autumn 2018 or winter/spring 2019		Early or late scenario. To be further discussed.
	CY47R1				New HPC technical (Q2 2020)

11. AOB

Testing strategy: Claude was interested in whether or not there was an appetite for setting up an automated test harness for IFS/ Arpège. Deborah pointed out that the overhead in maintaining an appropriate test framework may become a full time job. Peter reported positive experiences using a unit test framework in the ifsobs developments and noted how it might be useful to catch changes which break non-standard configurations early on (e.g. forecast sensitivity suite).

Claude asked about the potential plans to remove (or alternatively prune) the “classical” IFS CNT* control call tree once OOPS binaries start to be operational. EC believe that the control level structure will have to be maintained for a while, as various configurations will still require the CNT* call tree (surface analysis, perhaps singular vectors if not switched to OOPS before 2019 etc.).

12. Next meetings

Next technical video-conferences:

⇒ Tuesday 26 September, 2017 14h30 CET / 1.30pm UK

Next Coordination video conferences:

⇒ Thursday 28 September, 14h30 CET / 1.30pm UK

Next physical Coordination Meeting:

⇒ Monday 19 March 2018, meeting to take place in Toulouse

List of actions

1. **MF** to provide EC trace-back information about the SRTM problem/crash encountered in some Arpège PEARP members (this is under investigation in GMAP)
2. **MF (Claude and Yann Michel)** to contact Elias Holm for informing about the suggested implementation of a specific REDNMC_Q namelist coefficient in CY45T1.
3. **Tomas and Ryad** to record details of which versions of the compilers were used in the release notes for each cycle. Possibly start this with CY45 FLUBs.
4. **Steve** to ask Sébastien Massart to send details of experiences related to LSPRT and the humidity fields in the assimilation at 45r1.
5. **MF (Claude and Philippe)** to contact EC (Deborah and Shahram Najm) and co-ordinate to ensure that MF grib_api developments are included in the main eccodes package.
6. **All** to revisit decision on timing of OOPS implementation with regards porting to new HPC. As a side-effect probably, re-assess timing scenarios for CY47 (Early/late).