SPECIAL PROJECT FINAL REPORT

All the following mandatory information needs to be provided.

Project Title:	Forecasting at sub-kilometre resolution with the HARMONIE-AROME model
Computer Project Account:	spiecla2
Start Year - End Year :	2023 - 2025
Principal Investigator(s)	Colm Clancy
Affiliation/Address:	Met Éireann Glasnevin Hill Dublin 9 Ireland
Other Researchers (Name/Affiliation):	James Fannon Eoghan Harney Tatjana Kokina Eoin Whelan Met Éireann

The following should cover the entire project duration.

July 2025

Summary of project objectives

The objective of this project is to explore configurations for running the HARMONIE-AROME limited-area NWP model at hectometric resolutions, and to assess the potential benefits for forecasting over Ireland.

Summary of problems encountered

The HPC platform can often be unstable, with unexplained crashes requiring task requeues. Retrieval from ECFS can be quite slow too.

Experience with the Special Project framework

My experience was very positive. The application and general administration procedures are nice and straightforward, minimal overhead.

Summary of results

The work in this project focussed on model dynamics and boundary options at high resolution, both in the horizontal (mainly at 750m grid-spacing) and with an increased number of vertical levels.

In terms of the limited-area domain and boundaries:

- When looking at different nesting options, the age of the initiating global boundary files tends to dominate performance. Using IFSHRES boundary conditions generally gave better scores than nesting within an intermediate HARMONIE-AROME, and this approach is currently taken in the operational 750m HECTOR system
- Some coupling of hydrometeor and cloud species can be advantageous, however an optimal configuration remains elusive: full coupling of all throughout the simulation may not be ideal
- Results can be rather sensitive to the size of the domain, potentially even masking any effect in scores from increased resolution.

The model time scheme was tested for larger, more complex domains. HARMONIE-AROME uses a single-stage scheme with a SETTLS treatment of nonlinear terms, whereas other ACCORD model configurations use a predictor-corrector scheme which is, in theory, more stable. Here we found that:

- While generally more stable, more work is likely required for use on the large pan-European operational UWC-West DINI domain
- Use of the predictor–corrector scheme seems to trigger an upper-level noise pattern, which may be leading to the reduced performance with DINI. Extra diffusion is a potential, though unsatisfactory, means to damp this.

Earlier work on the boundaries was described in detail in the 2024 report. The updated results and the time scheme testing were described in an internal note, a version of which is appended to this report.

List of publications/reports from the project with complete references

Clancy C, Fannon J, Harney E, Kokina T, Whelan E.
High-resolution and Dynamics Experiments at Met Éireann.
ACCORD Newsletter No. 5, March 2024

Clancy C, Fannon J, Harney E, Whelan E.
Boundary and Dynamics Options for HARMONIE-AROME at Hectometric Scale
Internal NWP Note available from Met Éireann, 2024

Clancy C.
Experiments with time-stepping, domains, and resolution.
Internal NWP Note available from Met Éireann, 2025

Future plans

This work guides the ongoing development of the 750m HECTOR system now running operationally at Met Éireann. Preliminary work has been carried out on expanding this to an ensemble system, and work will likely continue in this direction.

Domains and Boundaries

Previous testing suggested benefits to coupling hydrometeors and cloud species at the boundaries, as well as having as large a domain as possble; this was described in previous progress reports. Here we look into this in a bit more detail, using the HARMONIE-AROME Cy46h1.1 tag.

Domain size

On the 1st of September 2020, heavy rain in the west led to localised flooding in Connemara. This has been a useful case to examine issues of boundary spin-up with rain. Figure 1 shows rainfall simulations using different domains and resolutions: the IREPS 2.5 km domain is used as a reference in Fig. 1a, while the 750 m HECTOR domain is used in Fig. 1b. The lower panels then continue at 750 m resolution but increase the domain size. In all of the 750 m simulations, full hydrometeor boundary coupling is used. We see that moving the boundary further away to the west does result in changed rainfall patterns, with more to the west, but we also see spurious patterns on the boundaries.

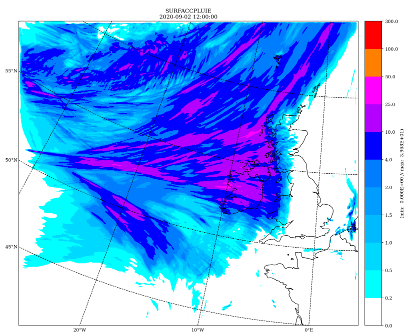

(a) Reference 2.5 km on IREPS domain

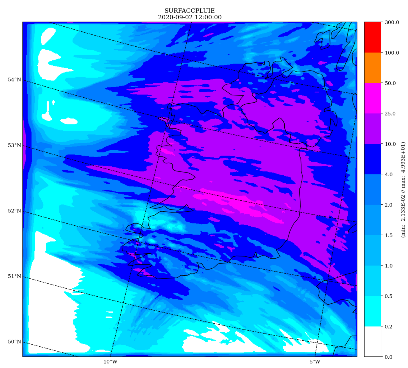

(b) 750 m, 800×800 domain

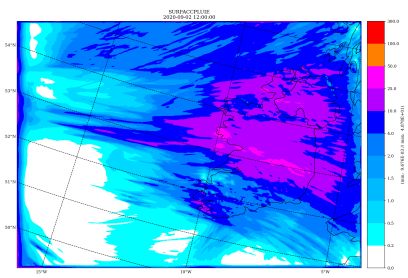

(c) 750 m, 1200×864 domain

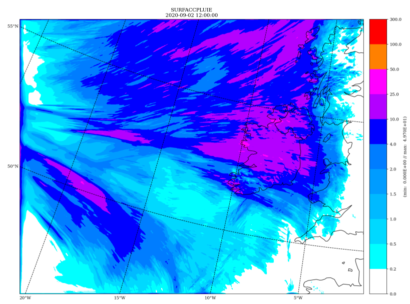

(d) 750 m, 2000×1600 domain

Figure 1: Forecasts of 36-hour accumulated rainfall starting from the 1st of September 2020.

We might assume that increasing the domain size would be generally a good thing to do, as long as it's affordable. Figure 2 shows verification of experiments at both 2 km and 750 m horizontal resolution. Both "large" and "small" geographic regions are tested as domains, and these are described in Table 1.

The scores in Fig. 2 suggest that things might not be as simple as "the higher the resolution, and the larger the domain, the better". Looking for specific differences, Fig. 3 shows sample cloud forecasts from the four experiments. We see that these can be sensitive to boundary proximity, as well as resolution. In the relatively-short 10-day testing period, such cases could cause wide variations in temperature scores.

Perhaps these results merely emphasis the need for longer testing periods in order to obtain more robust statistics. It is still something to bear in mind when comparing experiments: we still struggle to definitively show improvements with increased resolution. On the other hand, we often use smaller domains for cheaper testing: more care might be needed with this approach.

Experiment	Grid-size	Grid-points	Description
2kmLarge	2 km	720×648	IREPS domain area
2kmIRL	2 km	300×300	HECTOR domain area
750mIRL	750 m	800×800	HECTOR domain area
750Large	750 m	1200×1200	Increase HECTOR domain

Table 1: Details of domains used for the experiments in Figure 2.

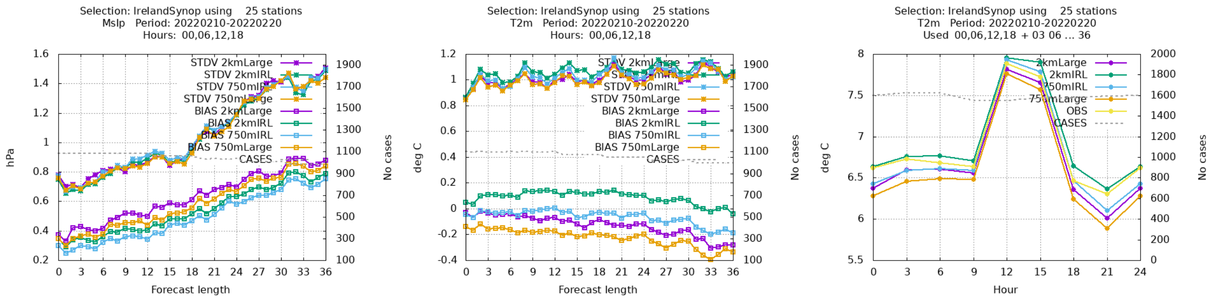

Figure 2: Verification of the experiments described in Table 1, using Irish synoptic stations. The period covered is 10-20 of February 2022. Four long forecasts per day were run, with 3-hour cycling in between using large-scale mixing and surface assimilation only.

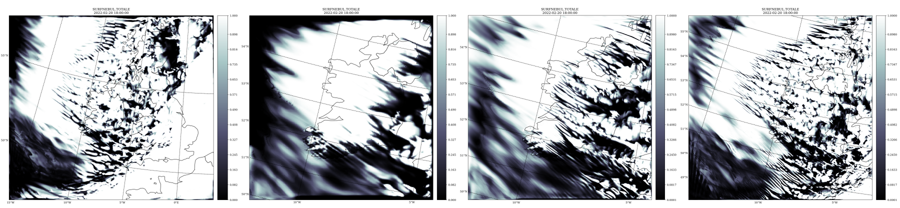

Figure 3: Forecast af 18 hours of total cloud cover from forecasts starting at 00z on the 20th of February 2022. The experiments are those in Table 1; left to right: 2kmLarge, 2kmIRL, 750mIRL, 750mLarge.

Boundary coupling options

As mentioned above, previous work suggested full coupling of clouds and hydrometeors from the global model, at the expense of some artifacts on the inflow boundary. However, more recently it was noticed that the options in the code are a bit more subtle than was realised.

For the various parameters, we have the `NCOUPLING` option. Setting this to 1 couples the field from the boundaries. However, a setting of -1 (the default for cloud and hydrometeors in HARMONIE-AROME), does not mean no coupling; rather it means "couple" to a fixed value, set by `REFVALC`, which is generally 0. Setting `NCOUPLING`=0 truly means uncoupled, so that the model is free to evolve in the boundary coupling zone. There is also `NREQIN` which similarly determines the initial coupling.

The Connemara case is again used to try different sets of options given in Table 2, with results shown in Fig. 4. These otherwise use the 750 m HECTOR configuration and can be compared to the full-coupled case in Fig. 1b, where we set everything to 1. A wider range of combinations had been tried. In general it was found that coupling ice after the initial time had some very large boundary effects. Balancing the need for reducing dry spin-up regions and avoiding spurious boundary artefacts, the configuration C2 in 4c seems to be the best. This work remains ongoing, however; there may be other factors such as physics affecting the boundary rain.

	C0	C1	C2	C3
Rain	(0,0)	(0,0)	(0,0)	(1,1)
Ice	(0,0)	(1,1)	(1,0)	(1,0)
Liquid	(0,0)	(1,1)	(1,1)	(1,1)

Table 2: Coupling options (NREQIN,NCOUPLING) for the experiments shown in Fig. 4. Note that all use (0,0) for snow and graupel (although the latter is not in the IFS LBC).

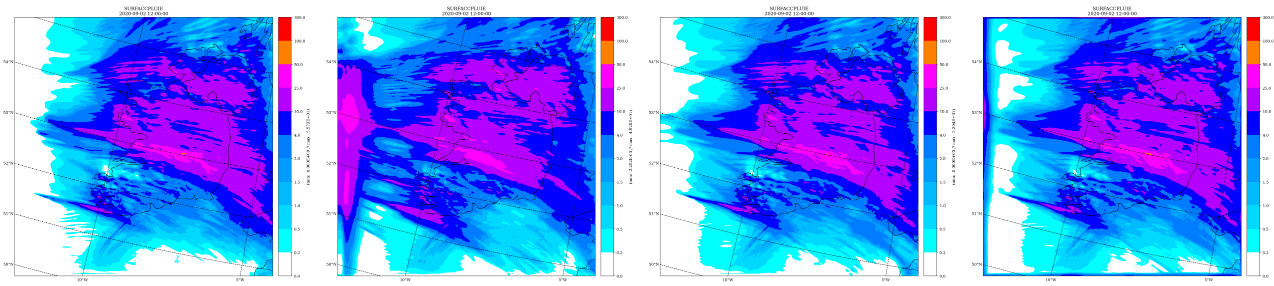

(a) C0 (b) C1 (c) C2 (d) C3

Figure 4: Forecasts of 36-hour accumulated rainfall starting from the 1st of September 2020, on the 750 m HECTOR domain. Each experiment uses different coupling options as detailed in Table 2

This alternative configuration was tested against the fully-coupled HECTOR configuration for two test periods: 10th-20th of February 2022 and 10th-25th of June 2023. Verification of standard surface parameters showed essentially no difference. Three-hourly rainfall at synoptic and CAMP stations is verified in Fig. 5. The slight differences mainly occur in the early hours. The spurious boundary rain in the fully-coupled case will give us more rain, of course, but this may or may not always be a good thing.

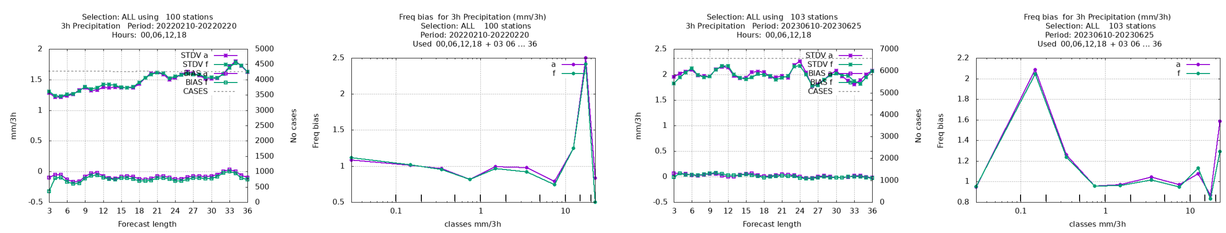

(a) 10th-20th of February 2022 (b) 10th-25th of June 2023

Figure 5: Verification of 3-hourly rainfall for experiments at 750 m comparing the fully-coupled HECTOR configuration (a, purple) with the C2 coupling configuration from Table 2 (f, green).

Time scheme

The time scheme in the ACCORD systems is a general iterative centered implicit (ICI) method, consisting of a first step followed by NSITER additional steps (e.g. Bénard et al, 2010). This is designed for improved stability with increasing iterations, but at the cost associated with the extra work. Both AROME and ALARO use a predictor–corrector (PC) method, so that NSITER=1. In contrast, HARMONIE-AROME uses a single step, NSITER=0, but with the nonlinear terms treated with the SETTLS scheme (Hortal, 2002).

Table 3 highlights some of the main namelist differences between the two approaches. In addition to the options described already, the so-called cheap PC does not recompute semi-Lagrangian (SL) trajectories. LNESC is a non-extrapolating option, which can't have the same value as LSETTLS. It should be noted that all three CSC (Canonical System Configurations) use SETTLS for the SL trajectory calculations (LSETTLST=TRUE).

Setting	AROME/ALARO	HARMONIE-AROME
LPC_FULL	TRUE	FALSE
LPC_CHEAP	TRUE	FALSE (N/A)
NSITER	1	0
LSETTLS	FALSE	TRUE
LNESC	TRUE	FALSE
SITRA	100	70

Table 3: Main time scheme settings among the CSC.

The HARMONIE-AROME settings result in a less stable scheme. This is dealt with by a lower value of SITRA, the so-called cold reference temperature for the semi-implicit: lower values are more stable, but reduce accuracy. In addition a shorter time step is used by HARMONIE-AROME, and the extra expense from this generally balances out the extra computations required by the cheap PC in AROME/ALARO.

In operations in HIRLAM countries, this configuration has run successfully for years at the default resolution of 2.5 km in the horizontal and 65 vertical levels (Gleeson et al., 2024). However as these increase, the performance and viability of the SETTLS option needs to be addressed, especially since dynamics developments are mostly happening within the AROME and ALARO communities.

As part of a previous Special Project, spieclan in 2021, HARMONIE-AROME stability was examined for a number of domains down to 500 m resolution. As well as the reduction in horizontal grid size, the vertical resolution was increased to the MF_90 levels. This adds extra layers close to the surface and can induce more stability issues than the horizontal changes.

Of the domains tested in that work, Greenland proved the most challenging (unsurprisingly), and the only fully stable run was with the PC scheme. For the less extreme Irish domain, the default SETTLS was sufficient, and gave near-identical verification scores; see Figure 26 in that Note. Based on this, the 750 m HECTOR suite has used SETTLS .

The DINI domain used by UWC-West has a horizontal grid-spacing of 2 km, and the MF_90 vertical levels. It covers a much larger area than the previous IREPS, and includes complex regions such as Greenland and the Alps and thus, given the extra vertical resolution, tests the time scheme's stability to a greater extent. In order to explore a more stable configuration, an AROME/ALARO-style PC scheme was tested in a number of cases. However, plots of winds near the tropopause, at around model level 20 (roughly 200-250 hPa) showed often noisey patterns; an example is shown in Fig. 6.

In the following sections this behaviour will be explored, along with the results of testing with time-scheme options in HARMONIE-AROME. Experiments have been carried out with various model cycles and code versions, but this behaviour has been repeatable.

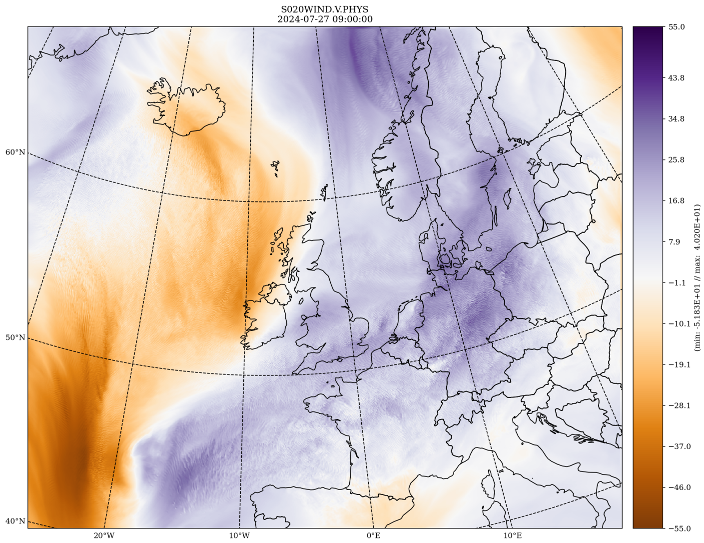

Figure 6: Sample snapshot of meridional wind at level 20 on the DINI domain, when the SETTLS scheme is directly replaced with the PC. This is a 24-hour forecast starting from 09z on the 26th of July 2024.

Single HARMONIE-AROME experiments

The 09z run on the 26th of July 2024 has been a case of interest, as the perturbed members of the DINI-EPS system showed some instabilities. This was examined in a number of single HARMONIE-AROME experiments, initially to address the stability. To begin with, the full operational DINI20A domain was used with the uwcwdr branch (essentially the operational 43h2.2 version). As was expected, the use of the predictor–corrector and AROME-like settings was more stable, but the upper noise mentioned above became apparent. To explore this noise with quicker, cheaper runs, a smaller 600×600 DINI20Ase domain over central Europe domain was created, and the 46h1.1 tag was used for these tests. Recreating the noise seemed to require a domain "big enough" to contain enough of the the upper jet.

An extensive series of tests exploring parameter configurations was carried out. The set of experiments detailed in Table 4 will be shown here, as these cleanly summarise the testing. Various other tests ruled out various other settings as direct causes/solutions: time-step, off-centring, grid truncation, upper boundary nesting `LUNBC`. The different diffusion settings in AROME affect the forecasts at the upper levels, but again do not fundamentally change the appearance of the noise.

Wind-speed forecasts at model level 20, around 200 hPa, are shown on the left in Fig. 7 for the experiments in Table 4. Noise patterns can be seen to appear in some, particularly over Germany. The right-hand panels in Fig. 7 show cross-sections of vertical divergence through this, from 10-15E longitude, at 50N latitude. Various kinetic energy spectra are then shown in Figs. 8 and 9.

The noise00 uses the default SETTLS scheme in HARMONIE-AROME, while noise01 switches SETTLS for a non-extrapolating option with `LNESC`. This may not be an advisable scheme, although

	noise00	noise01	noise10	noise14	noise28
NSITER	0	0	1	2	1
LPC_FULL	FALSE	FALSE	TRUE	TRUE	TRUE
LPC_CHEAP	FALSE	FALSE	TRUE	TRUE	TRUE
LNESC	FALSE	TRUE	TRUE	TRUE	TRUE
LSETTLS	TRUE	FALSE	FALSE	FALSE	FALSE
NPROFILEHD	3	3	3	3	2
RPROFHDBT	10000	10000	10000	10000	30000
RPROFHDTP	100	100	100	100	100
RPROFHDEX	1	1	1	1	4
RPROFHDMX	10000	10000	10000	10000	500

Table 4: Details of experiments used in Figs. 7 to 9. The settings in noise00 are HARMONIE-AROME defaults. Note that some of the RPROFHD parameters are only relevant with NPROFILEHD=3.

it remained stable in this case while the SETTLS noise00 gave SL trajectory warnings. However, upper-level noise can be seen to appear. This increases when we move to the more AROME-like PC in noise10, but disappears if we add an extra stage with NSITER=2 in noise14. The extra NSITER brings more cost; ideally we would prefer the cheap PC which is comparable to SETTLS.

We can then try to mitigate the noise by altering the diffusion (see options in Yessad, 2016). The strength of the horizontal spectral diffusion is controlled by various RDAMP settings for each parameter. Additionally, a vertical profile is applied to increase the effect at higher levels. There are a few options governed by NPROFILEHD and related settings. The noise28 experiment (Table 4) used the PC scheme of noise10 but with a modified profile designed to begin to take greater effect close to 200 hPa. It successfully damps the noise, although with some effect on the energy spectra (Fig. 9).

To summarise these tests:

- the default SETTLS scheme struggles with stability, but gives a smooth solution
- the PC scheme shows noise patterns at upper levels, seemingly triggered by the non-extrapolating aspect
- increasing the iterations removes this at a cost
- alternatively, we can modify the diffusion to address it

This is of course just a single case. The noise may not be a concern, rather related to the stable region around the tropopause. Increasing diffusion may have unwanted consequences elsewhere. The next section will look at longer experiments with this.

There are other options along with the spectral horizontal diffusion. A sponge option was tested but seemed to be too drastic in scrubbing all detail. An upper absorption layer has also been introduced (activated by LGWFRIC) in more recent ACCORD model versions. However, it is not available in HARMONIE-AROME Cy46.

Longer cycling experiments

Longer cycling tests were next carried out, in parallel to Cy46 testing by the UWC-West Development Team. In all of these, the "default" is considered now to be the UWC-West default configuration, i.e. MF_90 levels, quadratic grid, 50 s timestep at 2 km resolution, and `XIDT`=0.14. Testing on the full DINI20A domain used 3DVAR in general, with 48-hour forecasts at 00z and 12z and 3-hour cycling in between.

The settings to switch to the PC were given in Table 3. An initial range of experiments for 11 days was carried out on different domains. When run on a smaller domain covering Ireland and the UK, the PC ran successfully with 60 s timestep, and the verification scores were very similar to those of the default SETTLS, which is consistent with previously-mentioned work on small higher-resolution domains. In Fig.10, the experiments use the full DINI20A domain. Now things are quite different. The PC at 60 s remained stable but with many SMILAG trajectory warnings, and so a 50 s version was added. Even still, we get quite a spread in scores compared to the default.

The effect of domain size was mentioned earlier, and here again we have results which need careful consideration. Figure 11 shows forecast snapshots with the PC of upper-level zonal wind. On the small domain the solution is quite smooth, but on the full DINI domain we see the noisey behaviour discussed previously. As mentioned before, this only seems to manifest for reasonably-large domains (enough to capture certain aspects of the jet?). In general the default HARMONIE-AROME gave smoother MSLP forecasts, with the PC simulations seemingly affected by the upper noise.

For the rest of this section we compare three month-long (December 2024) experiments on the DINI domain, with details given below in Table 5. As mentioned, the PC at 60 s struggled with stability warnings on this large domain, so the experiment labelled dddc here matches the 50 s of the default SETTLS (aaac). The final experiment, dddd, uses the enhanced vertical profile already seen in the noise28 experiment of Table 4.

Name	Scheme	Diffusion
aaac	Default	Default
dddc	PC, 50 s	Default
dddd	PC, 50 s	`NPROFILEHD=2, RPROFHDBT=30000, RPROFHDTP=100, RPROFHDEX=4, RPROFHDMX=500`

Table 5: Details of month-long experiments on the DINI20A domain.

We first show kinetic energy spectra in Fig. 12 at a number of levels for these experiments towards the end of the period. The two PC experiments both have the upper level "feature". The extra diffusion helps somewhat, but not completely.

Verification is shown for surface parameters in Fig. 13. Overall it seems like the default with SETTLS (aaac, red) gives generally the best scores. While the PC (dddc, green) generally degrades these, modifying the diffusion profile (dddd, yellow), brings the results back a bit closer to the default: this suggests the upper-level behaviour with the PC is indeed causing problems on this large domain. For 2 m temperature, on the other hand, the unmodified PC gives generally warmer forecasts, which may be a good thing considering our general cold bias. The scores for this parameter can be particularly region-specific, with the overall averaged bias often dominated by stations at high altitudes; Fig. 13c shows the results for the other three UWC-West countries.

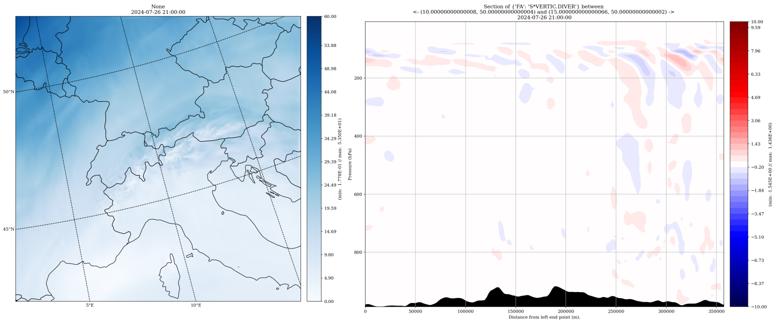

(a) noise00: default settings

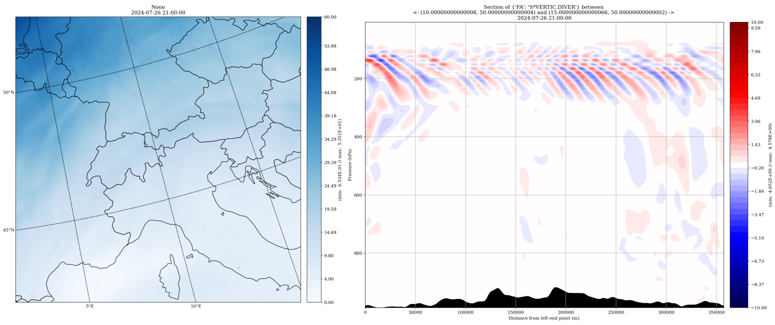

(b) noise01: single-stage non-extrapolating

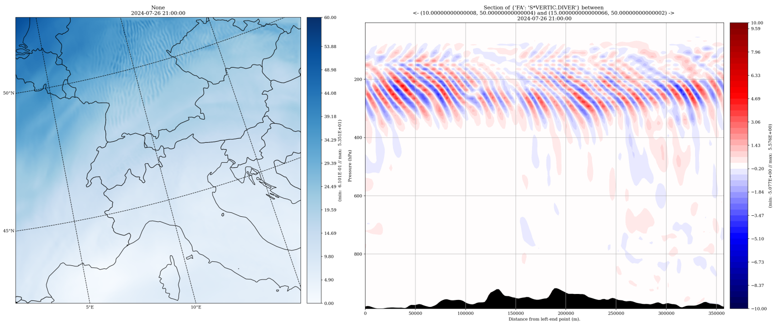

(c) noise10: predictor–corrector

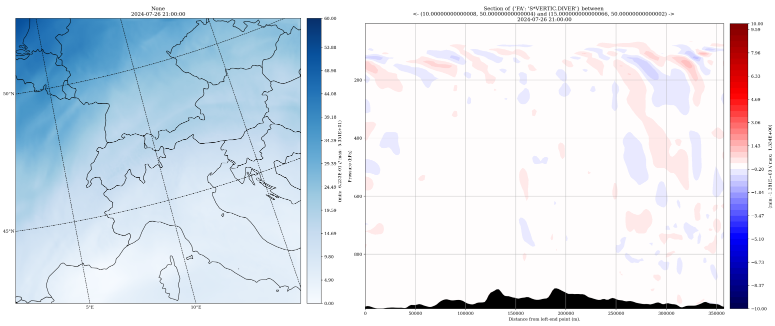

(d) noise14: additional step with `NSITER=2`

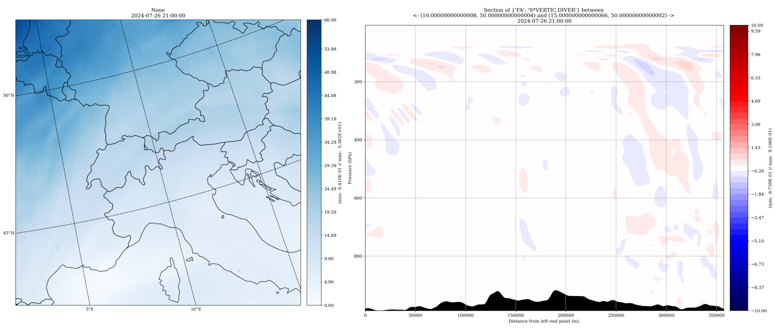

(e) noise28: predictor–corrector with adjusted vertical profile of diffusion

Figure 7: 12-hour forecasts of (left) wind speed at model level 20 (around 200 hPa), and (right) cross-sections of vertical divergence from 10-15 longitude at 50 latitude The experiments are described in Table 4, and the forecasts begin at 09z on the 26th July 2024.

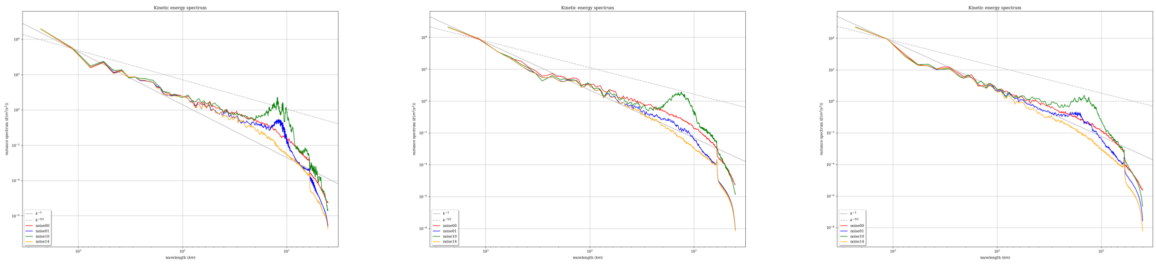

Figure 8: Kinetic energy spectra at model level 20 for experiments from Table 4, shown at forecast lead-times of (left to right) 6, 12, 24 hours.

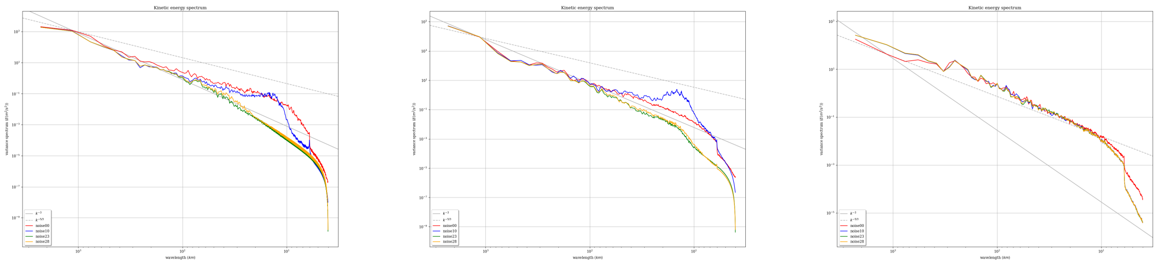

Figure 9: Kinetic energy spectra after 24-hours at (left to right) model levels 10, 20 and 60 for experiments from Table 4. Note, noise23 has a slightly different diffusion profile to noise28.

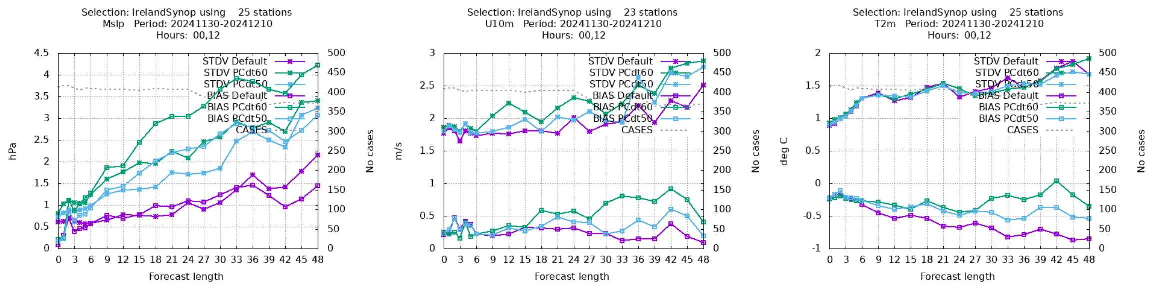

Figure 10: Verification of MSLP, 10 m wind speed and 2 m temperature using Irish synoptic stations, comparing various default and PC experiments. All have two 48-hour forecasts per day for the period of the 30th of November to 10th of December 2024. Experiments are on the full DINI20A domain: default (purple), PC with timesteps of 60 s (green) and 50 s (cyan).

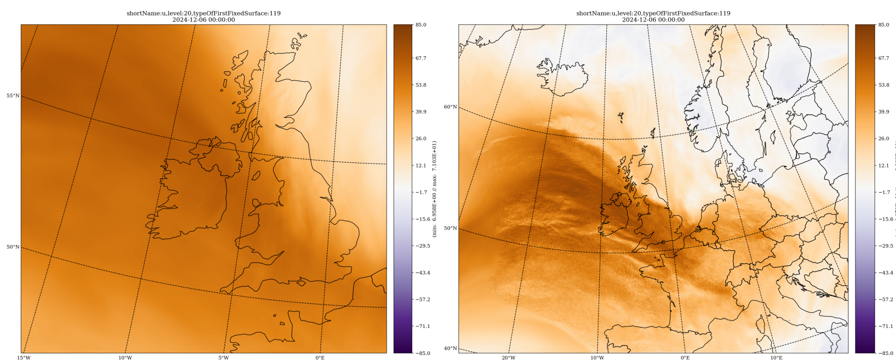

Figure 11: Forecast at 24 hours from 00z on the 5th of December 2024 of zonal wind at model level 20 forecast using the PC at 60 s. Left, on the IRELAND20 domain; right, DINI20A.

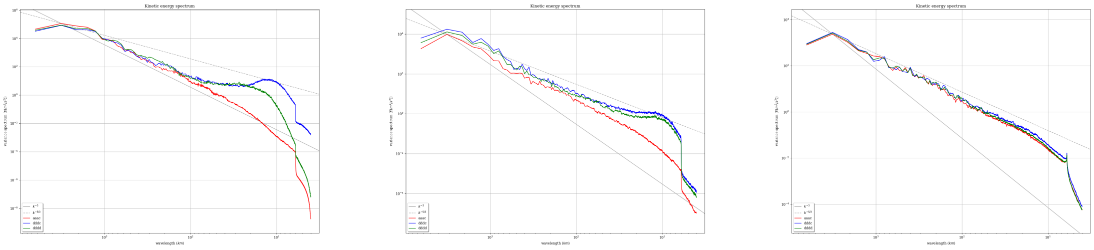

Figure 12: Kinetic energy spectra at levels (left to right) 20, 50, and 80, for the experiments in Table 5. The spectra are for the 24-hour forecast beginning at 12z on the 30th of December 2024.

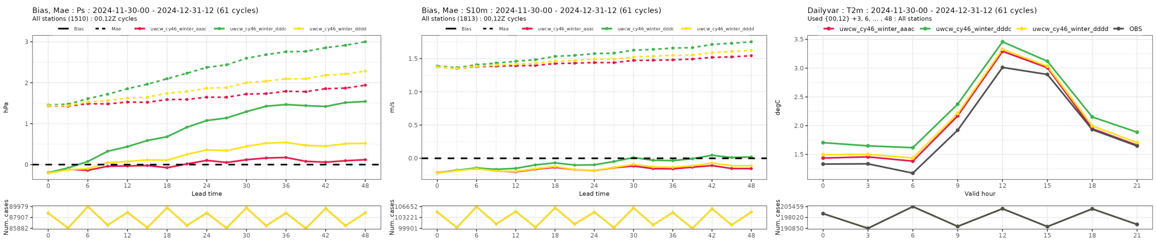

(a) Ps, S10m, T2m for all valid stations within the DINI20A domain.

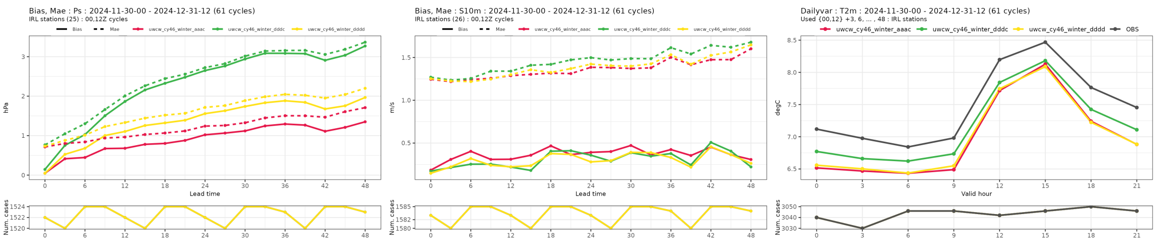

(b) Ps, S10m, T2m for Irish synoptic stations only.

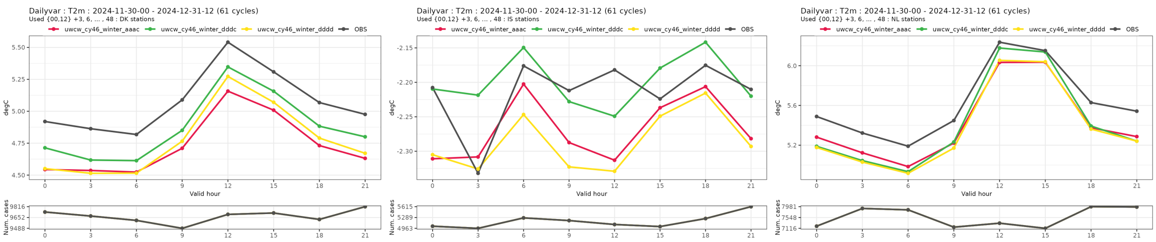

(c) T2m only for stations in (left to right) Denmark, Iceland, and the Netherlands.

Figure 13: Verification of surface pressure (Ps), 10 m wind speed (S10m) and 2 m temperature (T2m) for experiments on the full DINI20A domain. Two 48-hour forecasts per day were run from the 30th of November to the 31st of December 2024. The experiments are described in Table 5: aaac (red), dddc (green), dddd (yellow).

References

- Bénard, P., Vivoda, J., Mašek, J., Smolíková, P., Yessad, K., Smith, C., Brožková, R., Geleyn, J.F. (2010). Dynamical kernel of the Aladin–NH spectral limited-area model: Revised formulation and sensitivity experiments. Q. J. R. Meteorol. Soc., 136, 155–169
- Gleeson, E., Kurzeneva, E., de Rooy, W., Rontu, L., Martín Pérez, D., Clancy, C., Ivarsson, K.-I., Engdahl, B.J., Tijm, S., Nielsen, K.P., et al. (2024). The Cycle 46 Configuration of the HARMONIE-AROME Forecast Model. Meteorology, 3, 354-390
- Hortal, M. (2002). The development and testing of a new two-time-level semi-Lagrangian scheme (SETTLS) in the ECMWF forecast model. Q. J. R. Meteorol. Soc., 128, 1671–1687
- Yessad, K. (2016). Horizontal Diffusion Computations in the Cycle 43 of ARPEGE/IFS. Météo France/CNRM/GMAP/ALGO documentation.