

# Newsletter n°8

3 April 2026



**11th ACCORD  
Assembly.**

**Brussels,  
December 2025**

**ACCORD MG  
Phase 2 kick-off  
meeting.**

**Toulouse,  
February 2026**



*Credit: Anne-Lise Dhomps*

**ACC  RD**

A Consortium for CO<sub>n</sub>vection-scale modelling  
Research and Development

## CONTACT US

Claude Fischer, ACCORD Programme Manager

@ pm@accord-nwp.org

Anne-Lise Dhomps, Consortium Scientific Secretariat

@ css@accord-nwp.org

<http://www.accord-nwp.org/>

## Contents

<b>Edito</b> .....	3
<i>Claude Fischer, PM</i>	
<b>Personal reflections from the EPS Working Week at GeoSphere in Vienna</b> .....	4
<i>Henrik Feddersen</i>	
<b>Optimizing thinning parameters for radar reflectivity assimilation in ALARO cy46</b> .....	6
<i>Laura Gréta Magyar, and Antonín Bucánek</i>	
<b>Back-phasing the Mode-S EHS thinning and adaptive weighting from AROME cy49 to cy46</b> .....	21
<i>Péter Elek, Vivien Pourret</i>	
<b>All-sky IR assimilation of SEVIRI and FCI in AROME-Austria</b> .....	26
<i>Adhithyan Neduncheran, Florian Meier</i>	
<b>3DnVar with flow-dependent background error covariances</b> .....	32
<i>Chahrazed Bouzerma, Idir Dehmous</i>	
<b>Large eddy simulations with ALADIN input</b> .....	43
<i>Petra. Smerkol, Nika Kastelec, Jure Cedilnik</i>	
<b>The implementation of the sweep interpolation in the semi-Lagrangian scheme of the ACCORD system</b> .....	51
<i>Natalia Szopa, Petra Smolíková</i>	
<b>Visualization of norms</b> .....	55
<i>Harold Petithomme</i>	
<b>Patricia around ACCORD 26 beautiful countries (4)</b> .....	60
<i>Patricia Pottier</i>	
<b>Previous editions</b> .....	62
<i>Anne-Lise Dhomps, CSS</i>	

## Edito

Claude Fischer, ACCORD Programme Manager

The 8th ACCORD Newsletter is the first one under phase 2 of the consortium. Quite some management activity has been undertaken since January in order to get ACCORD up and rolling with the new compositions of its several bodies. PAC already had its first meeting devoted to the topic of open sourcing part of the codes. It is intended that STAC meets for its first phase 2 meeting in spring. The first phase 2 Assembly meeting is scheduled for the 3rd of July.

At the scientific and technical planning level, the new Management Group (MG) has started its regular weekly online meetings (usually on Friday mornings), complemented by an in-person kick-off two-day meeting in Toulouse. The new MG has among other things taken up the continued preparation and the monitoring of the ACCORD-funded actions from the Detailed Action Plan (DAP-2026). In the next weeks and months, the MG will organize the updates of the DAP with the additional actions to take place this year and start preparing the draft RWP-2027. The new MG members will thus take their full role in organizing the R&D activities, building their links with the teams, the scientists and the LTMs. A caveat regarding the management tasks is the missing two Area Leader positions (Physics, MQA). The Programme Management is currently exploring solutions for filling these two empty MG functions. In the meanwhile, activity in connection with the physics parameterizations and meteorological quality assurance will be organized by the existing MG in close collaboration with a few key scientists and with an emphasis on across-ACCORD priority directions.

This Newsletter 8 has several contributions related to data assimilation (4) and algorithmic or modelling-related topics (3). Our Area Leader for EPS in phase 1, Henrik Feddersen, is contributing with a summary of the Working Week hosted by GeoSphere Austria in October 2025. Patricia is sending her own greetings by a personal contribution at the end of the newsletter.

The Newsletter remains an essential tool for communicating results across ACCORD, in a fairly brief (yet clearly structured) manner enabling scientists and teams to rather promptly present their achievements. For future newsletter contributors, please check the editorial information important to read and available on the ACCORD website.

# Personal reflections from the EPS Working Week at GeoSphere in Vienna

Henrik Feddersen (Danish Meteorological Institute)

The EPS working week that took place at GeoSphere in Vienna, 20-24 October 2025, was the last in phase 1 of ACCORD, and very appropriately also the most popular one in terms of in-person participants (19) since the first ACCORD EPS working week in 2022. This may partly be due to the location in Vienna which is within easy reach for most, but I like to think it is also because we have found a form that makes it attractive to participate physically in the working weeks. During phase 1 we have experimented a little with the form and frequency of the working weeks, but it seems like one EPS working week per year fits most. With a relatively small group of active EPS people we try to include all aspects of ensemble forecasting in the working weeks and divide the week into mainly presentations during the first couple of days and mainly practical work during the rest of the week. The working weeks are always hybrid, so online participants can follow and give presentations and to some extent participate in discussions with in-person participants.

For the working week in Vienna there were relatively many presentations which left less time for practical work, but with several new people in the group, everyone was given the opportunity to present their work. With new people joining the group the aim of a working week is not only to get some practical work done, but also to establish new contacts and get to collaborate closer together.

The full programme for the working week, including presentations can (as always) be found on the ACCORD wiki, see

[https://opensource.umr-cnrm.fr/projects/accord/wiki/Vienna\\_20-24\\_October\\_2025](https://opensource.umr-cnrm.fr/projects/accord/wiki/Vienna_20-24_October_2025).

The presentations included both overview talks and more detailed presentations, many about model perturbations, but also presentations about the latest harp updates, use of URANIE to study parameter sensitivities in MUSC and a new method (cascading ensembles) to increase ensemble spread.

As area leader for EPS I am pleased to note that progress for EPS has actually been made since the start of ACCORD phase 1: More and more weather services in ACCORD have access to limited-area, high-resolution ensembles – in many cases provided by regional collaborative efforts. A lot of effort has been devoted to improving the modelling of uncertainty of physics parameterisations through the introduction of SPP and other types of parameter perturbations. Finetuning SPP has proven to be quite demanding, but hopefully the effort that has been put into the use of the "verification, validation and uncertainty quantification" (VVUQ) tool URANIE can make it less demanding in the future. Introduction of flow-dependent SPP is another interesting approach that might improve uncertainty modelling.

But challenges still remain: Ensembles still tend to be underdispersive, both initial condition and model perturbations might need adaptation for ensembles at hectometric scale, and what role will AI/ML models play in relation to EPS? I hope the development of a common scripting system – with EPS capabilities – will improve collaboration across CSCs and help us meet these challenges in the next phase of ACCORD.

The EPS working week in Vienna was my last EPS working week as area leader. Looking back over the last five years, the working weeks have always been enjoyable and highlights of the ACCORD collaboration, and I strongly encourage everyone in ACCORD working on any aspect of EPS to join the EPS working weeks in person. Inger-Lise Frogner (MET-Norway) has taken over as area leader for EPS from 2026, so the area is in good hands, and I am confident that EPS will continue to develop in the next phase of ACCORD. She will be in charge of the next EPS working week which will be hosted by DMI in autumn 2026, where I – on behalf of DMI – hope to welcome both old and new EPS friends.

# Optimizing thinning parameters for radar reflectivity assimilation in ALARO cy46

Laura Gréta Magyar<sup>1</sup> and Antonín Bučánek<sup>2</sup>

<sup>1</sup>HungaroMet Nonprofit Zrt.

<sup>2</sup>Czech Hydrometeorological Institute

## 1 Introduction

---

The assimilation of high-density radar reflectivity observations offers great potential for improving convective-scale numerical weather prediction, but it also poses significant challenges for data assimilation systems. In this article, we present work carried out during an ACCORD/LACE stay at the Czech Hydrometeorological Institute (CHMI), focusing on optimizing the thinning distance for radar reflectivity assimilation in ALARO cy46. Using an existing method based on observation-error correlations, we first derived a theoretically optimal thinning distance and then conducted sensitivity experiments to assess the impact of different thinning choices on verification scores. Proper tuning of the thinning distance is particularly important in the 3D-Var implementation of the ACCORD CSC, where we expect that two observations do not have correlated errors, and therefore the observation-error covariance matrix  $R$  is modelled as diagonal. An appropriate thinning strategy helps ensure that this assumption is respected, prevents the introduction of spurious correlations, and supports a stable, efficient assimilation that translates into measurable forecast improvements.

### 1.1 Thinning procedure for reflectivity in ALARO

In ALARO cycle 46, radar reflectivity data are assimilated using the so-called "1D+3D-Var" approach, as described by Wattrelot et al. (2014) [10]. In this framework, radar reflectivity is not assimilated directly; instead, pseudo-observed relative humidity columns are derived from the observed reflectivity columns and are used in the 3D-Var system. Figure 1 illustrates the preprocessing steps applied prior to the 3D-Var assimilation in our configuration.

During the assimilation cycle, radar observations can be filtered at two stages. The first opportunity arises immediately after reading the reflectivity data from the HDF5 files, during their conversion from polar to Cartesian coordinates within the BATOR module. At this step, the effective sampling distance can be adjusted using the `HODIM%Sample` namelist parameter.

A second filtering step is applied in the screening module following the Bayesian inversion. At this stage, a horizontal thinning of the relative humidity pseudo-observations is performed by defining a square grid using the `RMIND_RADAR` namelist parameter. Only one column per grid cell is retained, with preference given to columns located closer to the radar and those containing a larger number of elevation angles. This process is then repeated using an arbitrarily chosen grid size (`RFIND_RADAR`) and an optional horizontal shift (`XYSHIFT_THIBOX`) relative to the initial grid.

Figure 1 highlights a third option for reducing the density of radar observations. This procedure, known as superobbing [3], merges closely spaced observation points into a single representative super-observation whose value best reflects the characteristics of the combined measurements. The HOOF tool (developed by Peter

Smerkol [9]), used through the BATOR module to homogenize OPERA HDF5 ODIM files for assimilation in ALARO, includes the capability to perform superobbing. Comparing this method with the “traditional” filtering approach described above would be highly informative, but such an analysis lies beyond the scope of the present work.

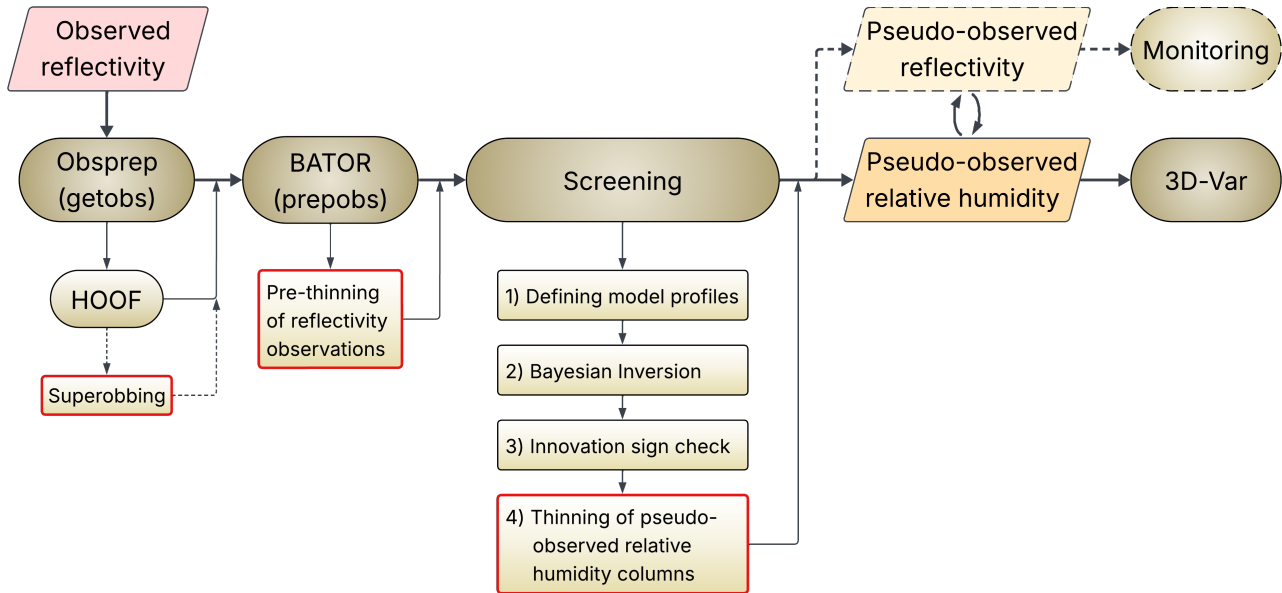


Figure 1: Flowchart of preprocessing steps for radar reflectivity observations using Bayesian inversion prior to 3D-Var assimilation

## 2 Methodology

### 2.1 Observation-error correlations

To select a suitable thinning distance, we adopted the criterion of Liu and Rabier (2003) [5], according to which observational sampling becomes counterproductive when error correlations between adjacent observations rise above about 0.2, leading to reduced analysis and forecast quality.

To compute the observation-error correlations, we used the ObsTool software (written in R) developed by Patrik Benáček at CHMI, which implements the method described by Desroziers et al. (2005) [2]. The tool uses first-guess and analysis departures to diagnose the error correlations of the observation-error covariance matrix  $R$  in observation space. For each observation type, pairs of first-guess and analysis departures are selected within a defined time window and subsequently grouped according to their spatial separation using predefined distance bins. This allows the observation-error correlations to be estimated as a function of distance. The results obtained from running ObsTool for our experimental period are reported in Section 2.3.

### 2.2 Operational model setup

The model configuration used in this study, along with the operational namelist settings relevant to thinning, is outlined below.

- **Model:** ALARO NH-v1B
- **Local CHMI cycle:** 46t1as\_op1
- **Domain (ALARO/CZ):**  $\Delta x \approx 2.3$  km; 1069x853 grid points; 87 vertical levels



Table 1: Thinning-related namelist settings used in all experiments and in the operational configuration

	Oper.	a52	a53	a54	a55	a56
HODIM%Sample	5000	5000	2500	5000	5000	15000
RMIND_RADAR	16275	1000	1000	10000	20000	16275
RFIND_RADAR	16275	1000	1000	10000	20000	16275
XYSHIFT_THIBOX(13)	0.5	0	0	0.5	0.5	0.5

configured to obtain the CCMA ODB datasets required to run ObsTool, and therefore, only the 3-hour assimilation cycle was performed in these two cases. To obtain the most reliable estimates of the observation-error correlations, we needed to “turn off” the filtering. This was achieved by substantially reducing the thinning distance; consequently, RMIND\_RADAR and RFIND\_RADAR were set to 1000 m, and the horizontal shift of the second grid was disabled (XYSHIFT\_THIBOX(13) = 0). In a52, the sampling distance in BATOR (HODIM%Sample) was kept at 5000 m, whereas in a53 the aim was to reduce thinning even further to allow as many observations as possible to pass through. We therefore set the BATOR sampling distance to 1000 m, which corresponds to the resolution of the OPERA radar data. However, the experiment exceeded the available memory despite using the maximum number of compute nodes allowed by the system. In practice, all 12 nodes were fully utilised, each equipped with 384 GB of memory, meaning that more than 4.5 TB of total memory was still insufficient to process the resulting number of observations. Consequently, the sampling distance had to be increased to 2500 m to keep the experiment within memory limits. Because the computational cost increased substantially, only four assimilation cycles could be completed, covering the first 12 hours of the experimental period. The upper panels of Figures 3 and 4 display the observation-error correlations of R and the corresponding number of assimilated observations as a function of distance for experiment a52, while the results obtained from the shorter a53 period are shown in the lower panels. A detailed analysis of these plots and the associated results is provided in Section 3.1.

Building on these findings, we conducted a sensitivity study consisting of three additional experiments, a54, a55, and a56, each employing different thinning configurations. These experiments were designed in light of the ObsTool diagnostics, which provided a guideline for a theoretically optimal thinning distance. In contrast to a52 and a53, both the assimilation and the production were run for these experiments, and verification was carried out to assess the impact of the different thinning setups on forecast quality. In this configuration, the system produces a +48-hour forecast at each 00 and 12 UTC. The corresponding results are presented in Section 3.2.

## 3 Results

### 3.1 Optimal thinning based on observation-error correlations

What we observed from the ObsTool diagnostics is that extending the time period – and thus increasing the total number of assimilated observations – leads to smoother and generally lower diagnosed correlations. This behaviour became evident when we began running experiment a52: as more assimilation cycles were completed, the diagnosed correlation curves progressively smoothed out and their magnitudes decreased. Additionally, the increased observational density in the higher-resolution setup (a53) also shaped the diagnostics. Figures 3 and 4 show that a53 configuration produces smoother and more internally consistent diagnostics; however, it also causes the observation-error correlations to decrease more rapidly with distance. In a53 setup, the correlations fall below 0.2 at approximately 7–10 km, whereas in the more robust a52 experiment this range shifts to about 10–15 km. Although a longer time period is desirable for improved statistical reliability, these results nonetheless provide a useful indication of a theoretically optimal thinning distance.

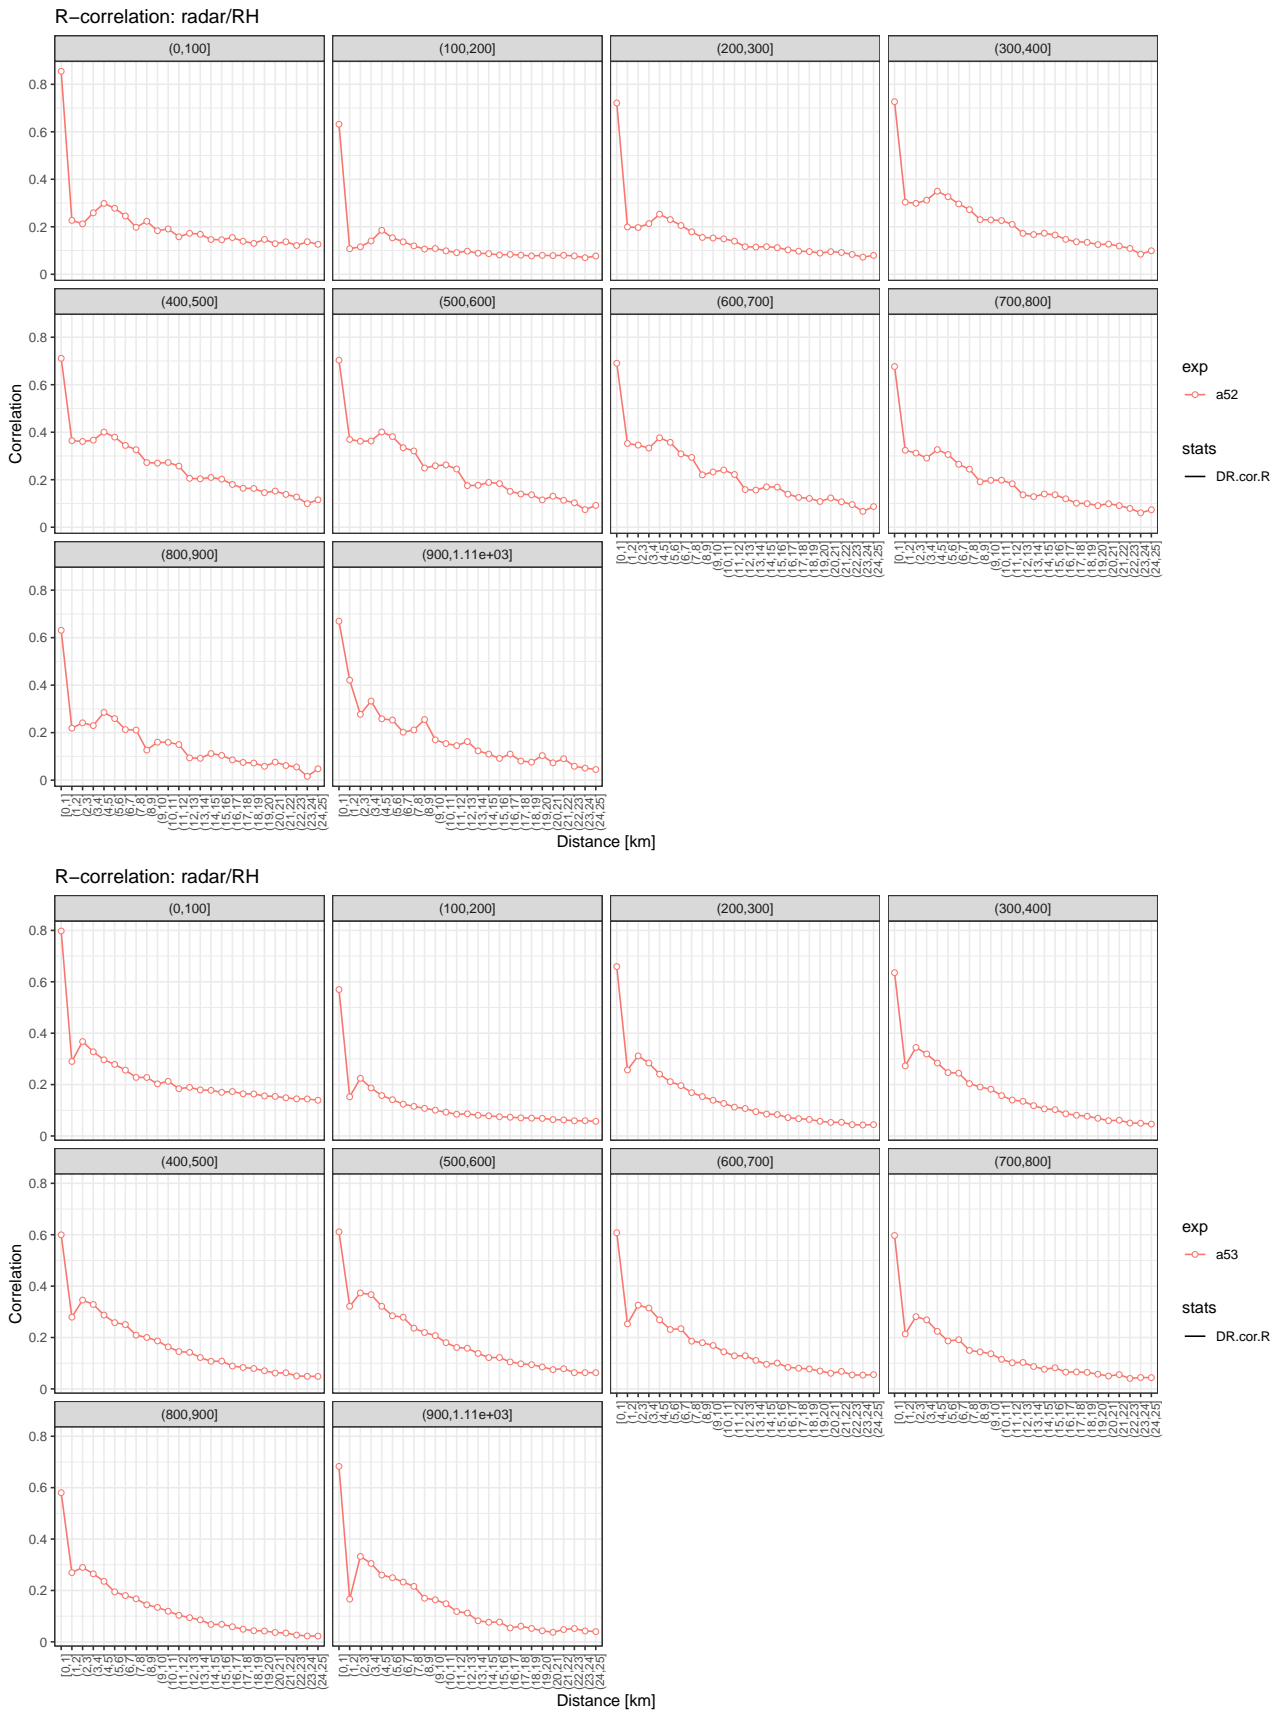


Figure 3: Diagnosed observation-error correlations of R as a function of distance for predefined vertical intervals in the a52 (upper) and a53 (lower) experiments

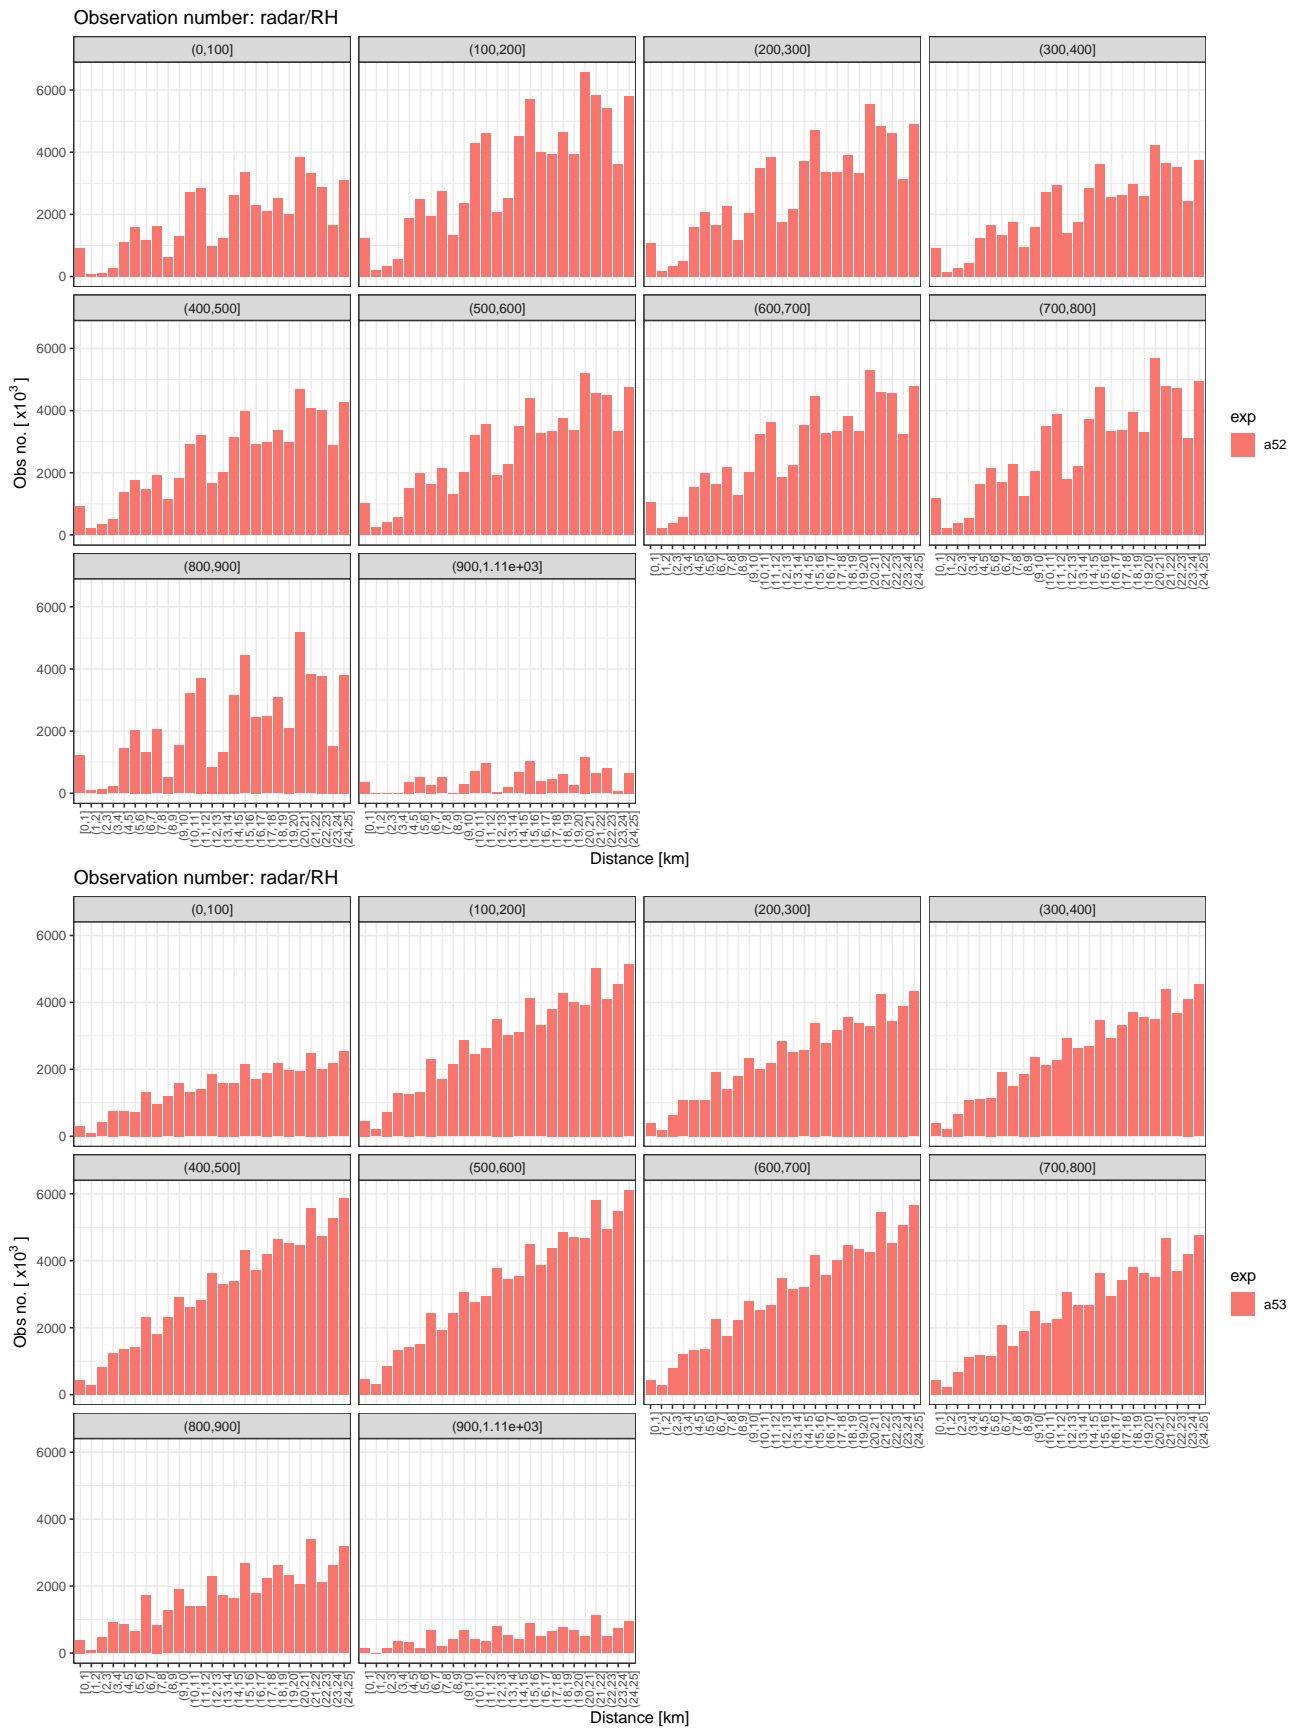


Figure 4: Number of relative humidity pseudo-observations as a function of distance for predefined vertical intervals in the a52 (upper) and a53 (lower) experiments used in the ObsTool diagnostics. (Note that experiment a52 covers a 10-day period (80 assimilation cycles), whereas experiment a53 spanned only 4 assimilation cycles (12 hours).)

### 3.2 Sensitivity studies

#### 3.2.1 Effects of thinning on observation density and assimilation metrics

To illustrate how different thinning configurations affect the spatial density of observations, we begin by introducing Figure 5, which compares experiments a54, a55, and a56 with the operational reference (see Table 1 for the thinning configurations used in each experiment). The plots show the western part of Czechia and provide a direct visual comparison of how each thinning setting modifies the density and spatial distribution of assimilated observations.

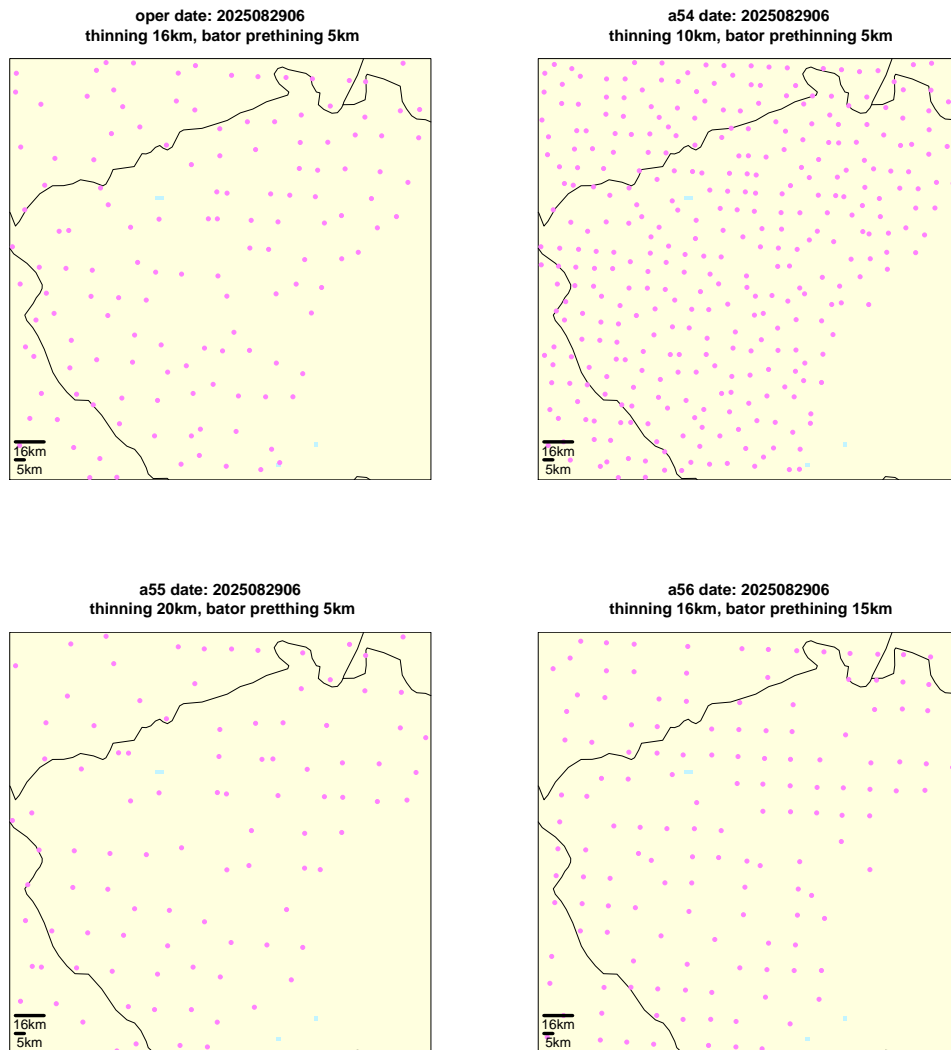


Figure 5: Observation density over western Czechia for the operational reference and thinning experiments a54, a55, and a56

As Figure 5 illustrates, even when a specific thinning distance is prescribed during screening, it is not possible to prevent observations from falling closer than the target spacing. This behaviour follows directly from the characteristics of the thinning procedure described in Section 1.1. Figure 6 further highlights this effect by showing histograms of the minimal intervals between observations for experiments a52, a54, a55, and a56. In the cases of a54 and a55, the largest counts in the distributions cluster around the respective target thinning distances, yet more than half of the observation pairs still appear below these. In contrast, the distribution for experiment a56 exhibits a slightly different pattern, with observation separations shifted more towards higher values.

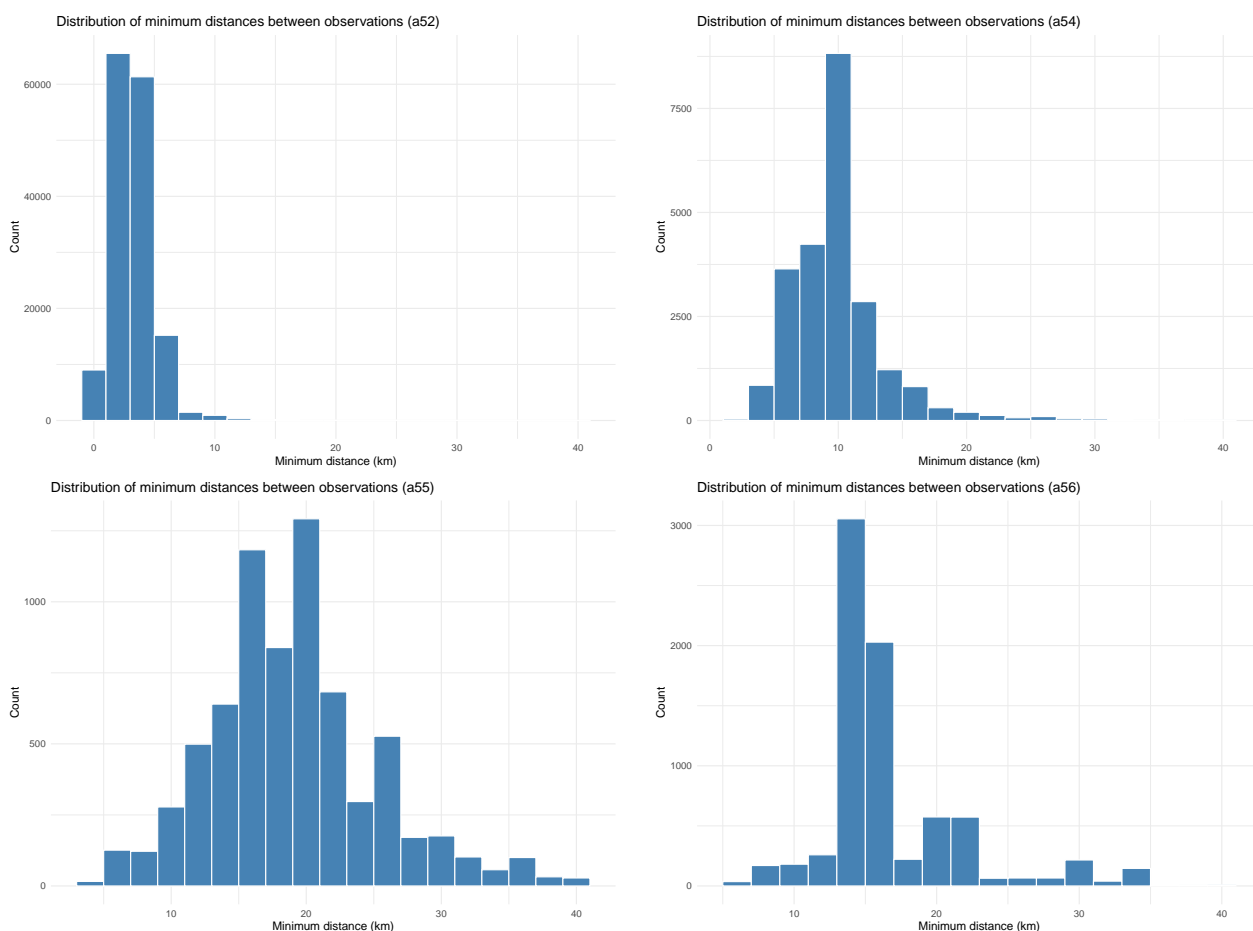


Figure 6: Histograms of minimal distances between observations for experiments a52, a54, a55, and a56. Here experiment a52 serves as a reference case in which the filtering was "turned off".

As a complementary quantitative perspective, Table 2 summarises the number of observations retained after screening, together with the subset of pseudo-observed relative humidity that was actually assimilated, for all experiments. This comparison provides a clearer view of how the different thinning configurations translate into effective observation usage within the assimilation cycle. Table 2 also reports the number of Lanczos iterations reached during the minimisations. Since this iteration count is limited to 50 in ALARO/CZ, it is informative to examine whether any of the experiments approach or reach this threshold. In addition, the table lists the final values of the cost function and the running times of both the screening and the minimisation processes, providing a more complete picture of the computational behaviour across the different thinning configurations.

Table 2: Summary of active observations, minimisation behaviour, and computational metrics for the different thinning configurations. RFL indicates reflectivity observations; H indicates pseudo-observed relative humidity; JO/n indicates that the cost-function value has been normalised by the number of assimilated observations.

	Data Count Screening		Data Count Minimization	No. of iterations performed	Runtime (sec)		Cost function (JO/n)
	H	RFL	H		Screening	Minimization	
a52	8682	382573	8649	36	143.6	129.9	0.04
a53	677570	1987678	631250	50	241.7	2362.7	0.02
a54	2816	335112	989	30	139.4	120.2	0.05
a55	2863	335112	504	30	140.4	121.3	0.06
a56	290	37082	232	30	101.5	108.5	0.05

From the table, it is evident that experiment a53, which has by far the highest number of assimilated observations, reaches the maximum allowed number of Lanczos iterations. In parallel, both the screening and minimisation runtimes increase substantially compared to the other experiments, reflecting the greater computational demand associated with processing such a large observation set. The table also highlights that increasing the sampling interval in BATOR (experiment a56) has a more pronounced impact on observation selection during the screening step, which in turn results in somewhat more favourable computational costs.

### 3.2.2 Impact of thinning distance configurations on forecast and verification scores

As we mentioned in Section 2.3, we conducted a verification for experiments a54, a55, and a56 to evaluate how the different thinning configurations affected the resulting forecasts. The verification was performed at 6-hour intervals by computing departures of the forecasted meteorological parameters from their corresponding SYNOP observations at the observation locations, and these differences were then used to derive statistical scores.

The results indicate that, of all meteorological quantities examined, only the precipitation scores showed noticeable deviations from the operational reference, suggesting that changes in thinning distance influenced mainly the precipitation forecast, while other fields remained essentially unchanged. Figure 7 illustrates the evolution of the mean error (bias) and the root mean square error (RMSE) as a function of forecast lead time for all experiments, shown for precipitation in the forecasts initiated at 00 UTC and 12 UTC, respectively.

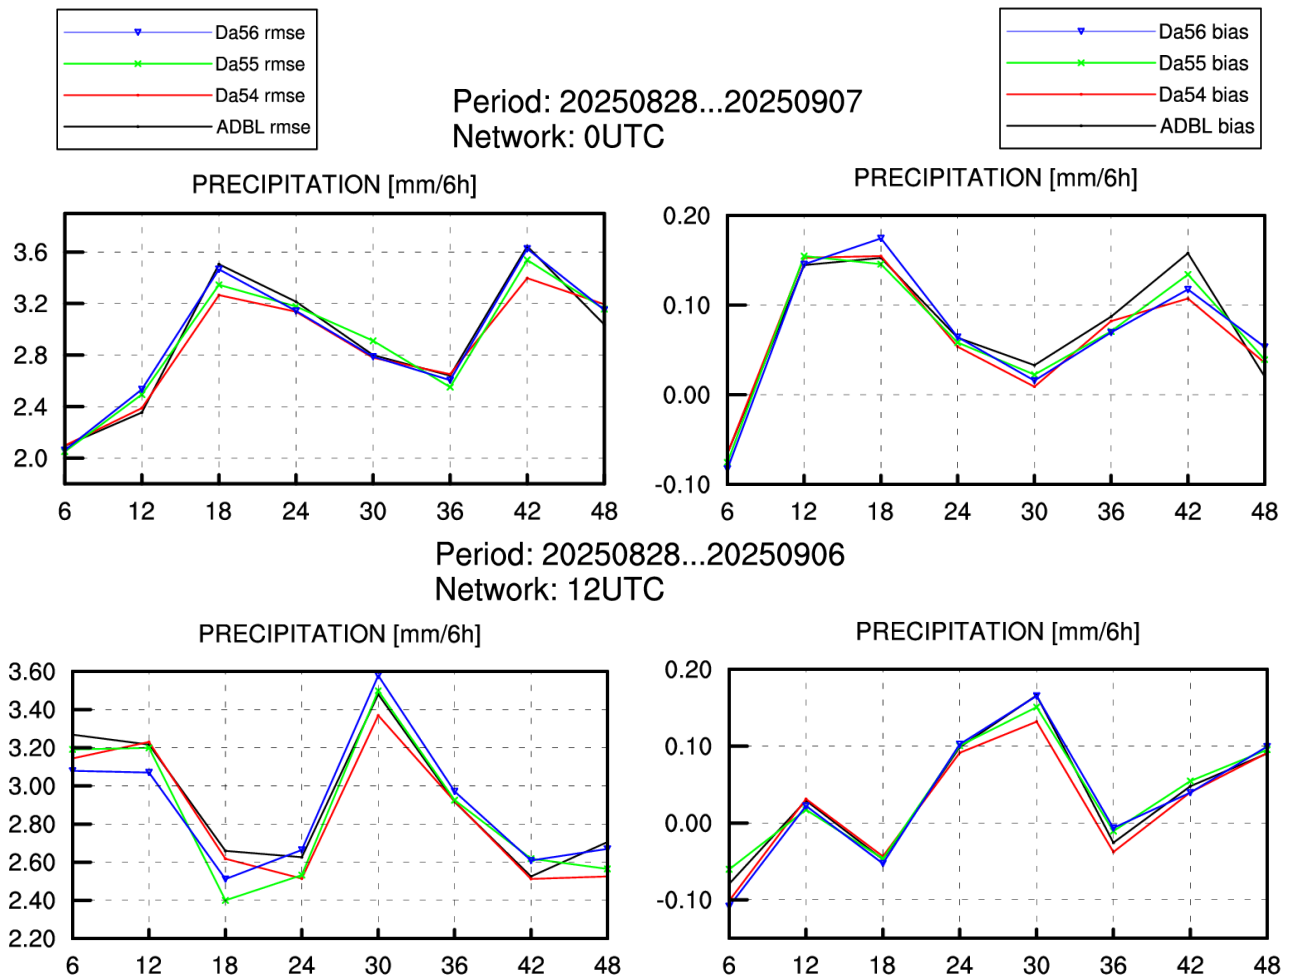


Figure 7: Evolution of RMSE (left) and bias (right) as a function of forecast lead time for all experiments for precipitation in the 00 UTC (upper) and 12 UTC (lower) production cycles, with ADBL denoting the operational reference.

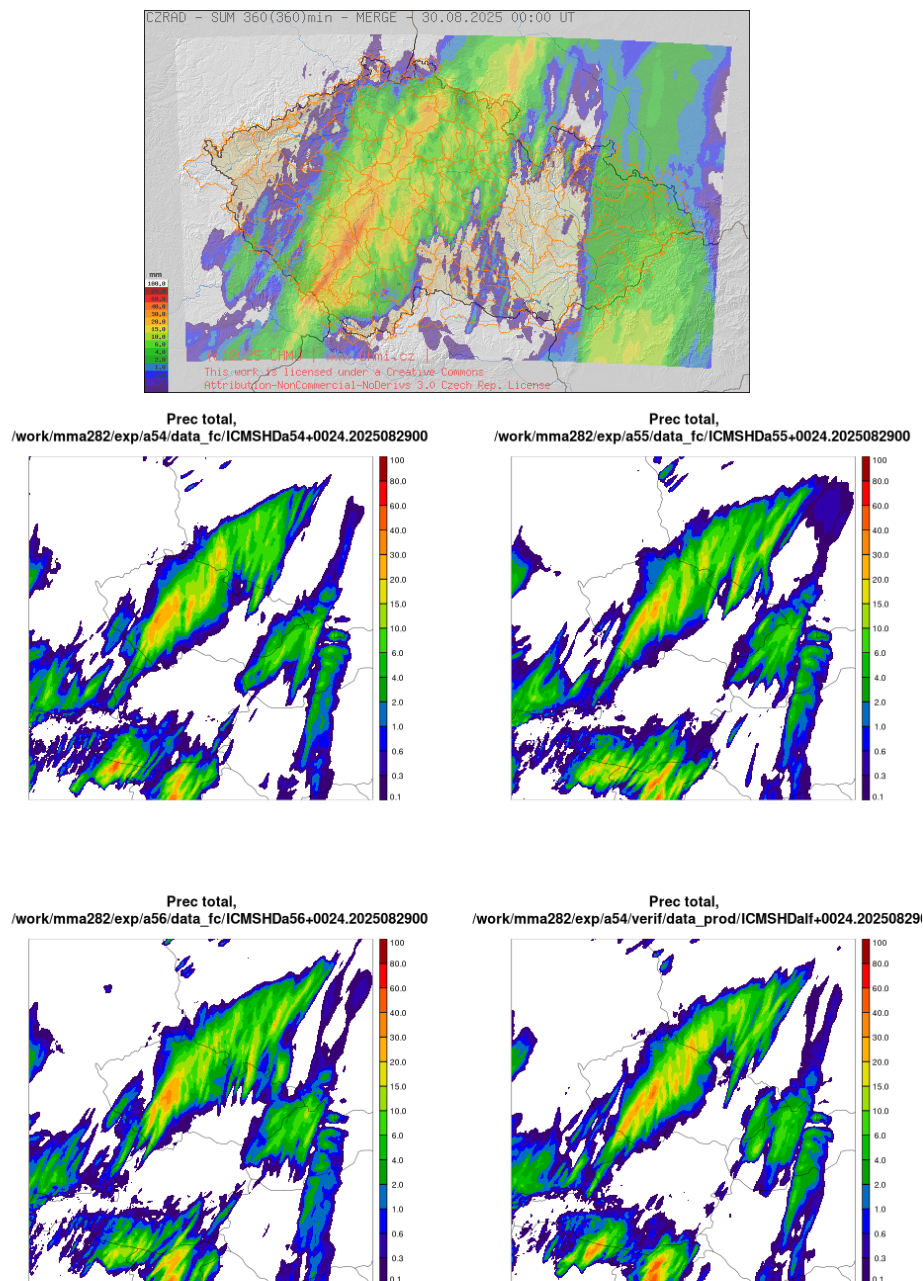


Figure 8: 24-hour forecast of total precipitation from the 00 UTC 29 August 2025 runs for experiments a54 (upper left), a55 (upper right), and a56 (lower left) with varying thinning distances, compared with the operational configuration (lower right) and the combined radar reflectivity and rain gauge precipitation estimate (top). A 6-h accumulation period is applied to all fields.

Based on these plots, the precipitation verification shows either a neutral or slightly positive impact across all three experiments. Only in a few cases does experiment a56—which is the configuration where the sampling distance was increased in BATOR—exhibit a minor degradation relative to the operational reference. Among the tested configurations, a54—corresponding to the lowest thinning distance of 10 km—appears to perform the best overall, as its RMSE remains consistently below, or at worst comparable to, that of the operational reference throughout the forecast range.

To visualise the results, Figure 8 presents the total precipitation fields for all experiments alongside the operational run, complemented by the observational reference derived from the combined radar reflectivity and rain gauge product over Czechia. The spatial structures appear to be remarkably similar in all configurations, with only minor variations in precipitation intensity, indicating that the different thinning distances exert minimal influence on the overall forecast performance.

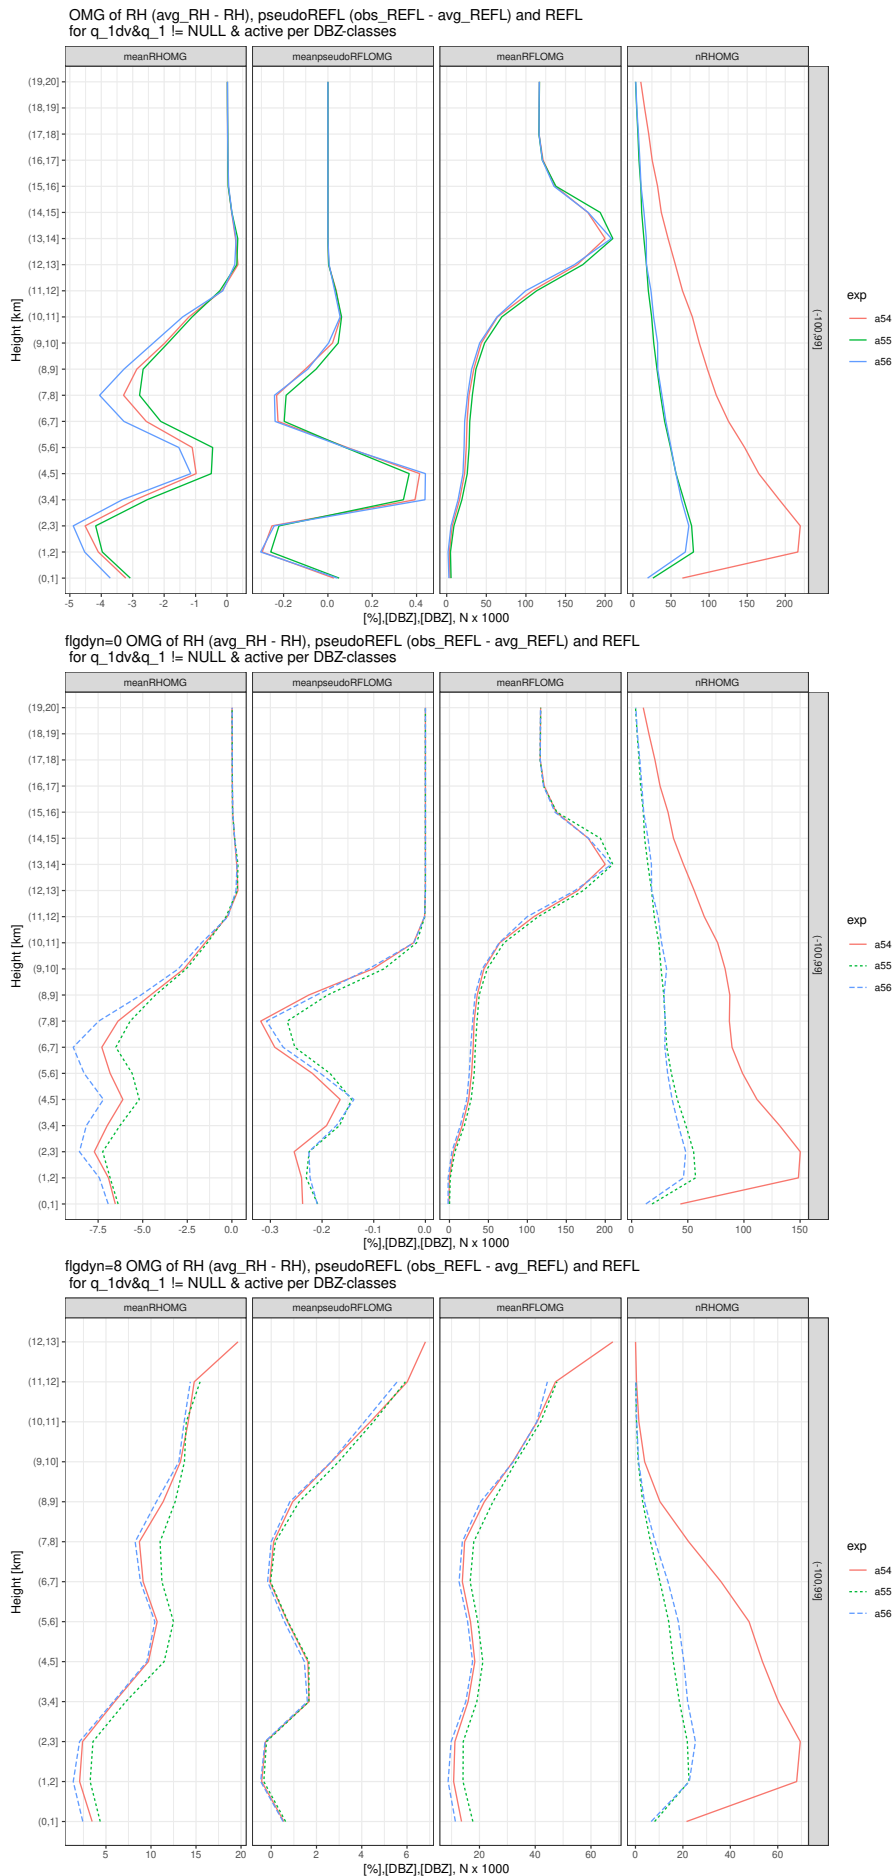


Figure 9: Observation–minus–guess statistics for relative humidity (meanRHOMG), pseudo-observed reflectivity (meanpseudoR-FLOMG), and observed reflectivity (meanRFLOMG) for experiments a54, a55, and a56. The upper, middle, and lower panels show results for all active, dry (fgdyn = 0), and wet (fgdyn = 8) observations, respectively. The rightmost panels display the number of active relative humidity pseudo-observations.

To further assess the impact of the different thinning distances, we computed observation–minus–guess (OMG) statistics for active pseudo-observed relative humidity, pseudo-observed reflectivity, and observed reflectivity. These results are shown in Figure 9. The first-guess departures are also plotted separately for dry ( $\text{flgdyn} = 0$ ) and wet ( $\text{flgdyn} = 8$ ) observations to evaluate their individual contributions to the overall statistics. The rightmost curves in each plot clearly illustrate the influence of the thinning distances on the number of active observations: experiment a54 separates noticeably from the others, reflecting the expected increase in sample size associated with its reduced thinning distance. This warrants some caution, as a larger sample may shift the resulting statistics (something we could already notice when discussing the observation-error correlation diagnostics in Section 3.1). What also immediately stands out is that dry observations are far more numerous than wet ones and therefore seem to drive the overall statistics. In addition, experiment a56 appears to dry the atmosphere the most among all configurations, shifting the mean relative-humidity departures further into negative values.

As a spatial verification metric, Fraction Skill Scores (FSS) [7] were computed over the full evaluation period to compare the precipitation fields across the different experiments. The FSS is computed here against a combined product of radar reflectivity and rain-gauge measurements, which is available only over the area of the Czech Republic. Figures 10 and 11 show the fraction of useful forecasts—an FSS value must exceed a predefined threshold to be considered useful—across all forecast ranges for several precipitation thresholds, using a 1-hour accumulation period, for the three experiments relative to the operational reference.

Both figures indicate that all experiments, as well as the operational run, achieve very similar scores. Unfortunately, it is not possible to assess the FSS at higher precipitation thresholds in the present case, as the percentage of useful forecasts drops sharply from 4mm/1h onward, as illustrated in Figure 11. For light precipitation, however, more than 70% of the forecasts are useful with reasonably good spatial accuracy across the area. This percentage decreases progressively with increasing thresholds, and the displacement error becomes more pronounced as well. The experiments differ most clearly at the 1mm/1h and 2mm/1h thresholds. At 2mm/1h, experiments a54 and a55 show a very slight improvement, while around the 1mm/1h threshold, experiment a56 performs the worst among all configurations.

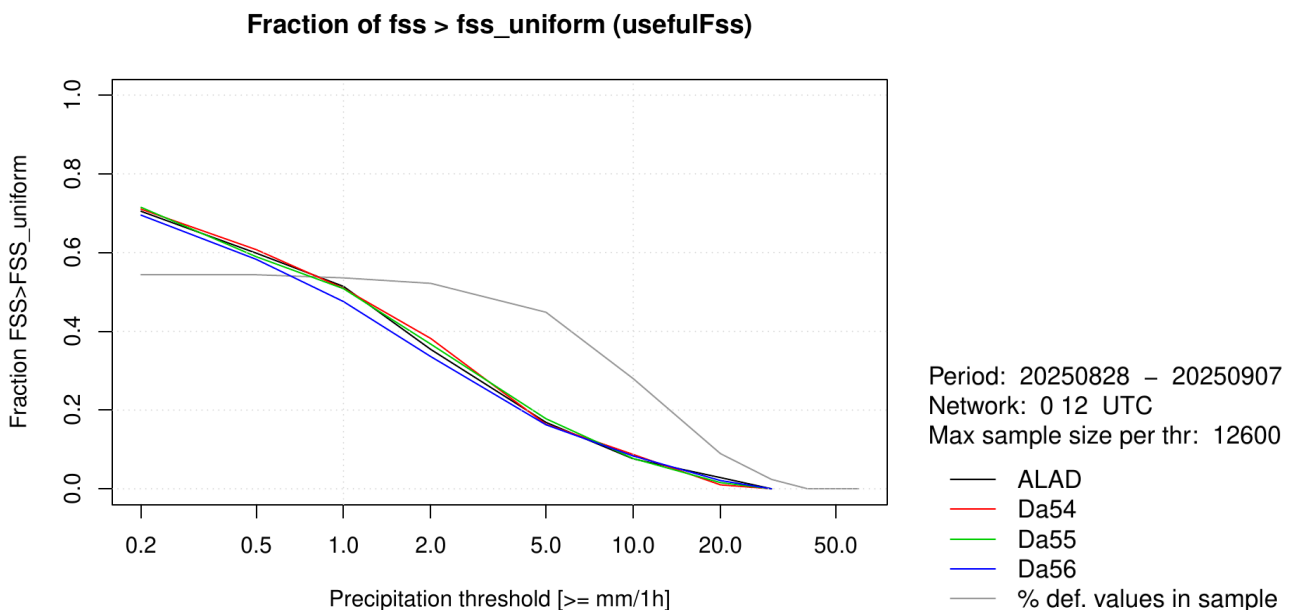


Figure 10: Fraction of useful forecasts ( $\text{FSS} > \text{FSS}_{\text{uniform}}$ ) for all experiments compared to the operational reference (ALAD). The grey line indicates the percentage of observations that meet the criterion for the FSS to be considered useful.

### Percentage of useful FSS over all FC ranges

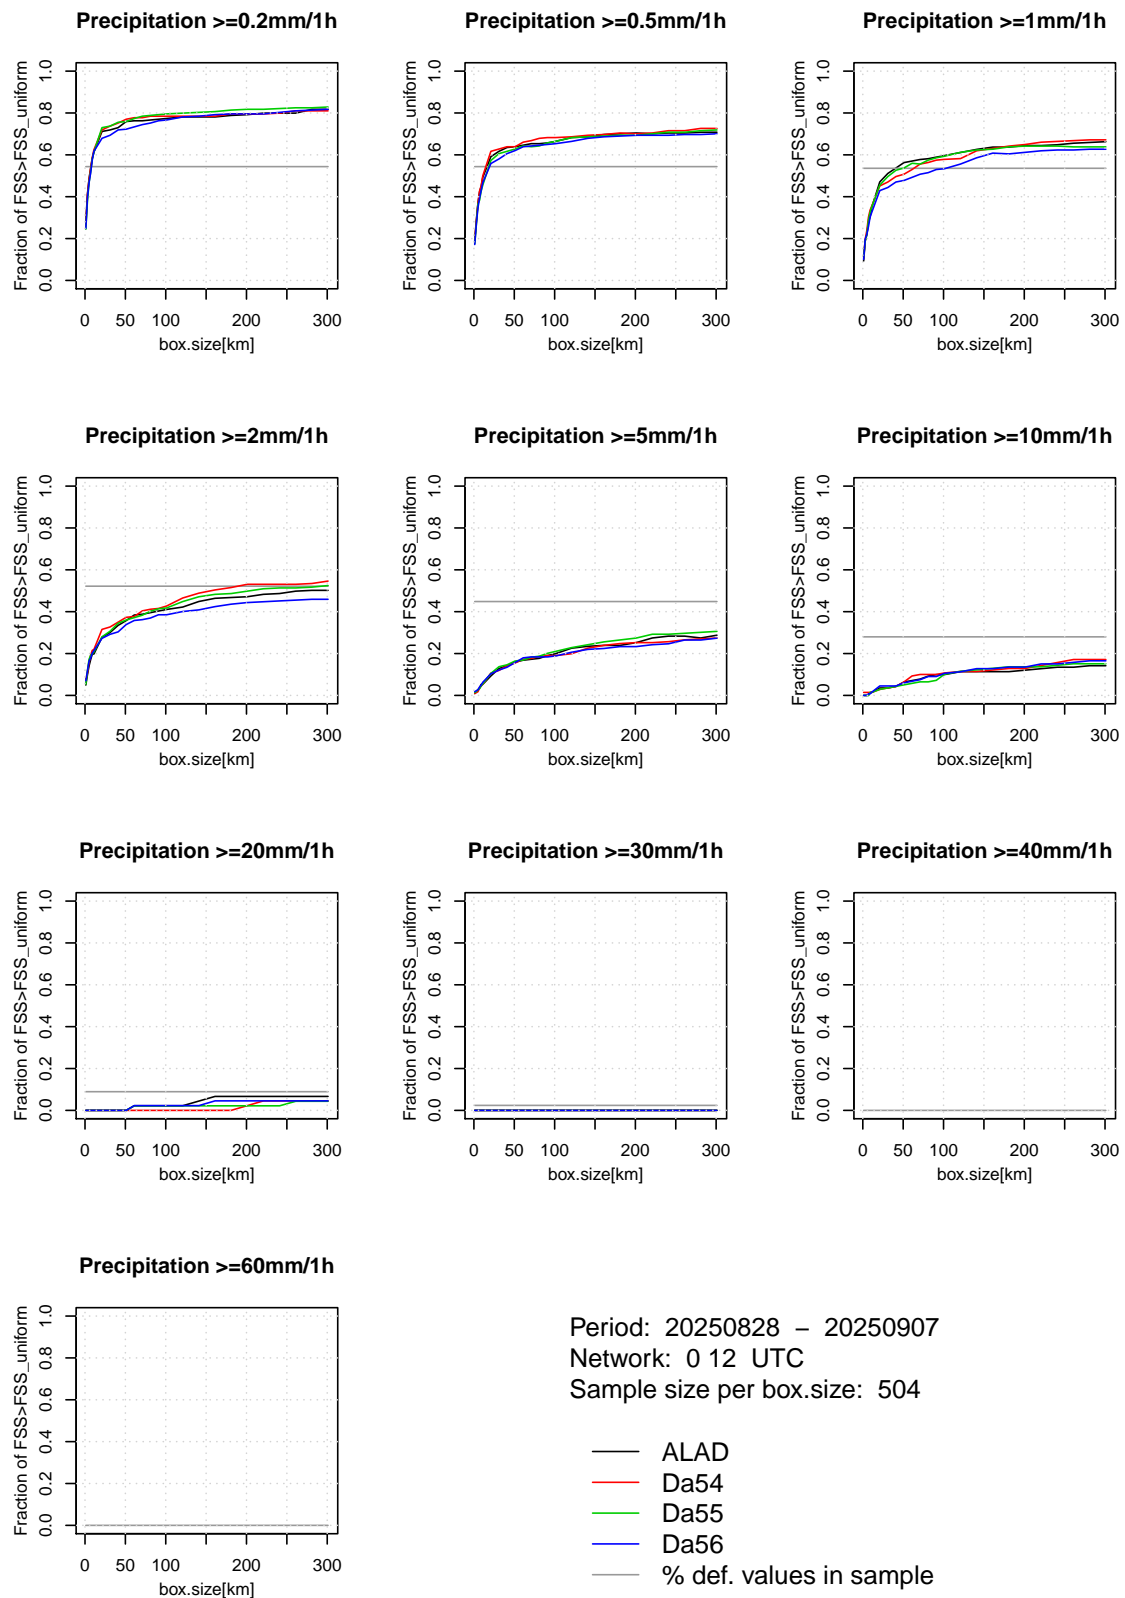


Figure 11: Percentage of useful forecasts as a function of box size across all forecast ranges for several precipitation thresholds (1-hour accumulation) for the three experiments relative to the operational run (ALAD). The grey line shows the percentage of observations that meet the criterion for the FSS to be considered useful.

## 4 Conclusion

---

This study examined the impact of different thinning-distance configurations on the assimilation of radar reflectivity in ALARO cy46, with the goal of identifying a setup that balances observational information content, error correlations, and computational efficiency. Using observation–error correlation diagnostics based on the Desroziers method [2], we estimated that a suitable thinning distance lies roughly between 7 and 15 km. At the same time, the diagnostics clearly showed that both the length of the evaluation period and the number of assimilated observations strongly influence the diagnosed correlations: longer periods and larger sample sizes produce smoother and generally lower correlation curves.

In light of these findings, three sensitivity experiments (a54, a55, a56) were carried out to assess how different thinning configurations affect the assimilation cycle and subsequent forecasts. The verification results indicated that thinning distance does have a measurable influence on assimilation behaviour and forecast quality, although the overall impact remained modest. Adjusting the pre-thinning in BATOR proved to be an effective way to address memory and computational-performance issues, as seen in experiment a56. However, increasing the sampling distance in BATOR—and thus filtering out many observations before later quality-control steps—may be compromising: across all verification metrics, a56 consistently performed somewhat worse than both the operational reference and the other thinning configurations. The other two experiments behaved similarly to the operational run, with a54 (10 km thinning in screening) showing the most consistently neutral or slightly positive impact.

Both the ObsTool diagnostics and the verification results suggest that a longer experimental period would be beneficial for obtaining more precise and reliable conclusions. Extending the evaluation period would help stabilise the diagnosed correlations and provide a more accurate assessment of the optimal thinning configuration.

In addition to refining the thinning distance itself, future work should also examine the potential of superobbing as an alternative means of reducing observation density, to assess whether it provides a more efficient or more physically consistent approach than the current thinning strategy applied to radar reflectivity data.

## 5 Acknowledgements

---

This work was supported by the ACCORD consortium, whose financial contribution is gratefully acknowledged. Special thanks are also extended to the Czech Hydrometeorological Institute for their warm hospitality and for providing an excellent working environment during the research stay.

## 6 References

---

- [1] Abel, S. J., and I. A. Boutle (2012), An improved representation of the raindrop size distribution for single-moment microphysics schemes, *Quarterly Journal of the Royal Meteorological Society*, 138(669), 2151–2162, doi:<https://doi.org/10.1002/qj.1949>.
- [2] Desroziers, G., L. Berre, B. Chapnik, and P. Poli (2005), Diagnosis of observation, background and analysis-error statistics in observation space, *Quarterly Journal of the Royal Meteorological Society*, 131, 3385–3396, doi:10.1256/qj.05.108.
- [3] Duncan, D., N. Bormann, A. Geer, and P. Weston (2023), Superobbing and Finer Thinning for All-sky Humidity Sounder Assimilation, doi:10.21957/5c3b9c8d9f.
- [4] EUMETNET OPERA Programme (n.d.), OPERA Radar Data Centre (ODC) Database.

- 
- [5] Liu, Z.-Q., and F. Rabier (2003), The potential of high-density observations for numerical weather prediction: A study with simulated observations, *Quarterly Journal of the Royal Meteorological Society*, 129, 3013–3035, doi:<https://doi.org/10.1256/qj.02.170>.
- [6] Marshall, J. S., and W. M. K. Palmer (1948), The distribution of raindrops with size, *Journal of Atmospheric Sciences*, 5(4), 165 – 166, doi:10.1175/1520-0469(1948)005<0165:TDORWS>2.0.CO;2.
- [7] Mittermaier, M., N. Roberts, and S. A. Thompson (2013), A long-term assessment of precipitation forecast skill using the fractions skill score, *Meteorological Applications*, 20(2), 176–186, doi:<https://doi.org/10.1002/met.296>.
- [8] Skamarock, W. C. (2004), Evaluating Mesoscale NWP Models Using Kinetic Energy Spectra, *Monthly Weather Review*, 132(12), 3019 – 3032, doi:10.1175/MWR2830.1.
- [9] Smerkol, P. (2022), *Documentation for the HOOF2 preprocessing tool*, ARSO, technical documentation for the HOOF2 radar preprocessing tool.
- [10] Wattrelot, E., O. Caumont, and J.-F. Mahfouf (2014), Operational Implementation of the 1D+3D-Var Assimilation Method of Radar Reflectivity Data in the AROME Model, *Monthly Weather Review*, 142, 1852–1873, doi:10.1175/MWR-D-13-00230.1.

# Back-phasing the Mode-S EHS thinning and adaptive weighting from AROME cy49 to cy46

Péter Elek<sup>1</sup> and Vivien Pourret<sup>2</sup>

<sup>1</sup>HungaroMet

<sup>2</sup>Météo-France

## 1 Introduction

The rapid update cycle data-assimilation benefits (e.g. 1 hour assimilation cycle) observation data that has a good spacial and time coverage on the model domain, especially in the upper atmosphere. The Mode-S EHS data from EMADDC perfectly satisfies these conditions. However the density of Mode-S EHS is too large to assimilate all available data, therefore it is important to find the right methodology (thinning, adaptive weighting) to effectively use the available data. A two-week long ACCORD/RC LACE stay took place on the topic of back-phasing the Mode-S EHS thinning and adaptive weighting from AROME cy49 to cy46 during 18 August - 30 August 2025 in Météo-France, Toulouse. The main motivations of this stay was to get familiar and adapt Vivien Pourret's work at Météo-France. Vivien Pourret achieved forecast improvements by using adaptive weighting on EMADDC Mode-S EHS data in AROME cy49. Since AROME cy46 is operational in several meteorological services in ACCORD, we decided to back-phase the adaptive weighting.

## 2 Methodology

In order to experiment and test the adaptive weighting we firstly back-phased the related code to the latest operational cy46 at Météo-France and then to the latest HungaroMet cy46 version.

### 2.1 Pre-thinning the Mode-S EHS data

At HungaroMet, we experimented with Siebren de Haan's box thinning (with box-heights of 300, 300, 600, and 1000 meters this setting produces the box-heights of 300 meter for the first 2 box, 600 meter for the 3rd and 1000 meter box height after that, and a box-width of 40 km this input translates to coordinate differences therefore it's smaller as we approach the poles), but its usage did not lead to any improvement in the forecasts. On the other hand, the colleagues from Météo-France achieved a forecast improvement with the usage of thinning and adaptive weighting in AROME cy49. This thinning is done in 2 steps:

step 1: arpifs/obs\_preproc/thiair.F90 performs a half resolution (e.g. 12.5km) thinning with a classical box thinning. This step is only used to reduce the number of the data to be handled by the next step.

step 2: arpifs/obs\_preproc/thiair2.F90 performs a true (e.g. 25km) thinning by computing inter-distances matrix by aircraft and by model level with the remaining output locations of step 1. Then we select data spaced by more than the thinning distance (25km) on the aircraft track by model level. After the thinning steps, the adaptive weighting is performed.

### 2.2 The method: Adaptive weighting method

The main idea behind adaptive weighting is the following: after the thinning which is only used aircraft by aircraft to avoid horizontal observation error correlations the Mode-S EHS data still remains dense with numer-

ous aircraft providing data, especially near airports ( $\sim 10000$  aircraft flying across AROME France Domain,  $\sim 6000$  aircraft flying across AROME Hungary Domain). We don't want to overfit the analysis for that data type therefore, to avoid overfitting, we select dense areas on a grid with a resolution of about the spatial representativeness of our model. By modifying the observation error of Mode-S data in dense boxes, the goal is to keep all observations and reduce their weights in the analysis process. If we choose the weights as the square root of the number of Mode-S EHS data in the dense box, then everything happens as if we had a homogeneous weighted grid of Mode-S data.

$$J(x_a) := \sum_{k,l \notin \text{modes}} (H(x_a)_k - y_k)R_{kl}^{-1}(H(x_a)_l - y_l) + \sum_{i \in \text{modes}} \frac{(H(x_a)_i - y_i)^2}{(w_i \sigma_i)^2} + \\ + \sum_{\substack{p,q \in \text{modes} \\ p \neq q}} (H(x_a)_p - y_p)R_{pq}^{-1}(H(x_a)_q - y_q) + \text{background terms}$$

Where  $R$  is the observation error covariance matrix,  $\sigma_i^2 := R_{ii}$ ,  $H$  is the corresponding observation operator,  $w$  is the adaptive weight,  $y$  is the observation value and  $x_a$  is the model state vector in the analysis time (during the minimization we calculate this:  $\arg \min J(x)$ ) and  $k, l, i, p, q$  is the indices of the data and  $\text{modes}$  is the index set of the Mode-S EHS data.

Let's decompose the sum that includes the adaptive weighting. We could define  $n_a(i_g)$  which is the index set of Mode-S data in the AW grid with index  $i_g$ . In addition,  $n_g$  denotes the number of AW grid points. By using the fact that these weights only depend on the AW grid  $w(i_g) := w_i = w_j$  if,  $i, j \in n_a(i_g)$ .

$$\sum_{i \in \text{modes}} \frac{(H(x_a)_i - y_i)^2}{(w_i \sigma_i)^2} = \sum_{i_g=0}^{n_g} \sum_{i \in n_a(i_g)} \frac{(H(x_a)_i - y_i)^2}{(w_i \sigma_i)^2} = \sum_{i_g=0}^{n_g} \frac{1}{w(i_g)^2} \sum_{i \in n_a(i_g)} \frac{(H(x_a)_i - y_i)^2}{(\sigma_i)^2} = \\ = \sum_{i_g=0}^{n_g} \frac{1}{|n_a(i_g)|} \sum_{i \in n_a(i_g)} \frac{(H(x_a)_i - y_i)^2}{(\sigma_i)^2}$$

As we can see in the last term, if we choose  $w(i_g) = \sqrt{|n_a(i_g)|}$ , then we use the average in the AW grid during the minimization. In practice, in a simple example, we have 2 data types: one of them is Mode-S EHS. In that case, when we perform the minimization without adaptive weighting, we minimize this (in the simplest case  $R$  is diagonal with values  $\sigma_o^2$  and  $\sigma_i^2$  where the indices  $i \in \text{modes}$  are the values related to Mode-S EHS data. Also, the background terms are not detailed here because they do not change when we use adaptive weighting);

$$\frac{(H(x_a)_o - y_o)^2}{\sigma_o^2} + \sum_{i \in \text{modes}} \frac{(H(x_a)_i - y_i)^2}{\sigma_i^2} + \text{background terms}$$

Where  $x_a$  is the analysis,  $y_o$  is the non Mode-S EHS observation data, and  $y_i$  is the Mode-S EHS observation data when  $i \in \text{modes}$ . Let's assume that the total number of Mode-S EHS data is  $n = |\text{modes}|$ . Then, with adaptive weighting, we minimize this:

$$\frac{(H(x_a)_o - y_o)^2}{\sigma_o^2} + \frac{1}{n} \sum_{i \in \text{modes}} \frac{(H(x_a)_i - y_i)^2}{\sigma_i^2} + \text{background terms}$$

In the special case when we have identical Mode-S data (identical location and data fields), then:

$H(\cdot)_i = H(\cdot)_j \wedge y_i = y_j \wedge \sigma_i = \sigma_j \forall i, j \in \text{modes}$  let's define  $y = y_i$  and  $\sigma = \sigma_i$  where  $i$  is a fixed element of the set  $\text{modes}$ .

$$\frac{(H(x_a)_o - y_o)^2}{\sigma_o^2} + \frac{(H(x_a)_i - y)^2}{\sigma^2} + \text{background terms}$$

As we can see (in this very special case), this does not depend on the Mode-S EHS observation number anymore; therefore, we don't overfit on Mode-S EHS data due to their very high density compared to the other observation types.

At the code level, the `RFIND_CLSTR` namelist parameter sets the box size for adaptive weighting, which means that all observation errors in this box are multiplied by the square root of the number of aircraft in the box.

### 3 Experiments

During the stay, we ran several experiments at Météo-France to test the modifications. The test experiments span the time period of 2024.01.01-2024.01.31. In these tests we used the Mode-S EHS data available from EMADDC [1].

The EMADDC provides two different file format OBSOUL and BUFR. The OBSOUL files are available in every 15 minutes (`obsoul_2_ehs_eu*`). The BUFR files have two different versions, one of them is available in every 15 minutes just like the OBSOUL (`bufr_2_ehs_eu*`) and the other has a shorter update cycle in every 5 minutes (`bufr_2_ehs_fa*`). We used the `bufr_2_ehs_fa*` files because they are available in every 5 minutes.

With the test runs we verified that the LMDEHS (which allows or blocks the Mode-S EHS data use in the model) and LSPROBG (which turns on or off the adaptive weighting) variables works as expected.

The longer experiments are not finished during the stay, but they are finished now and with the internal tools at Météo-France, we calculated the Desroziers diagnostics to ensure that we obtained the results as expected (see Figure 1 and Figure 2). At the time of writing the report, we noticed that `cy49` (which produces correct behavior) and `cy46` produced different behavior with adaptive weighting.

As we can see in Figure 1 and Figure 2 the experiment with the back-phased `cy46` at Météo-France produced a bogusly high ( $\sim 10$  times larger) obs error compared to the correctly behaving `cy49` with adaptive weighting. The reason for this was a data selection implementation error in the back-phased adaptive weighting code.

After the stay we continued the back-phasing at HungaroMet. Investigating the `odbs` (ECMA,CCMA) produced by the version run on the supercomputer of HungaroMet, unexpected results are obtained (with suspiciously high adaptive weights), near to airports at high vertical coordinates ( $\sim 10000\text{m}$ ). This phenomenon occurs in the edge case when the adaptive weighting box partially covers multiple thinning boxes, which in rare cases (where the air traffic is dense, e.g., airports) causes a high number of aircraft in the adaptive weighting box. This could be resolved by double thinning with shifted grids or just simply changing slightly the thinning or adaptive weight box sizes.

For further details, there is an RC LACE report (Elek, 2025, [link](#)).

Table 1: Back-phased model settings

	HungaroMet	Météo-France
Model version	AROME cy46t1_bf07	AROME cy46t1_op1.11
Data assimilation methods	3D-Var + SEKF	3D-Var
Used aircraft data	AMDAR, Czech Mode-S MRAR	AMDAR
Update frequency	3 hourly	1 hourly
Horizontal and vertical resolution	2.5kmL60	1.3kmL90
Used Mode-S data	EMADDC, 5 minutes	EMADDC, 15 minutes
Mode-S data prethinning	Siebrén's thinning	No prethinning
RFIND_CLSTR	4500m	4500m
RFIND_AIREP	10000m	25000m

## Conclusions

The Mode-S EHS data has great potential to improve forecasts. Also the implementation of adaptive weighting in cy46 significantly benefited from the collaboration. However, there is a remaining task to optimize the parameters on our domain.

## Acknowledgements

I would like to express my gratitude to ACCORD for financing my stay. I would also like to thank the Météo-France colleagues for their exceptional hospitality and willingness to help. Lastly, I would like to thank Vivien Pourret for his continuous help and guidance during my stay.

## References

- [1] V. Pourret, J.-F. Mahfouf, V. Guidard, P. Moll, A. Doerenbecher, and B. Piguet, “Variational bias correction for mode-s aircraft derived winds,” *Tellus A: Dynamic Meteorology and Oceanography*, vol. 73, no. 1, pp. 1–27, 2021.

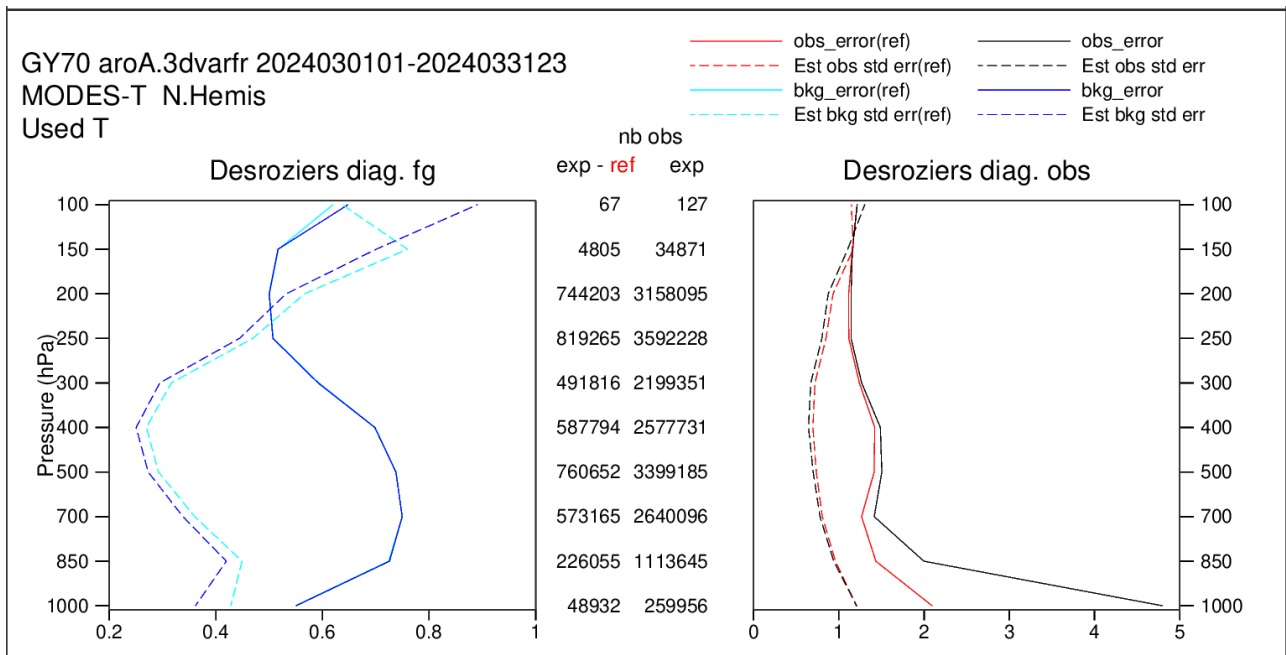


Figure 1: cy49 results, RFIND\_CLSTR = 4500 meter  
 reference: with Mode-S EHS without adaptive weighting  
 experiment: with adaptive weighting

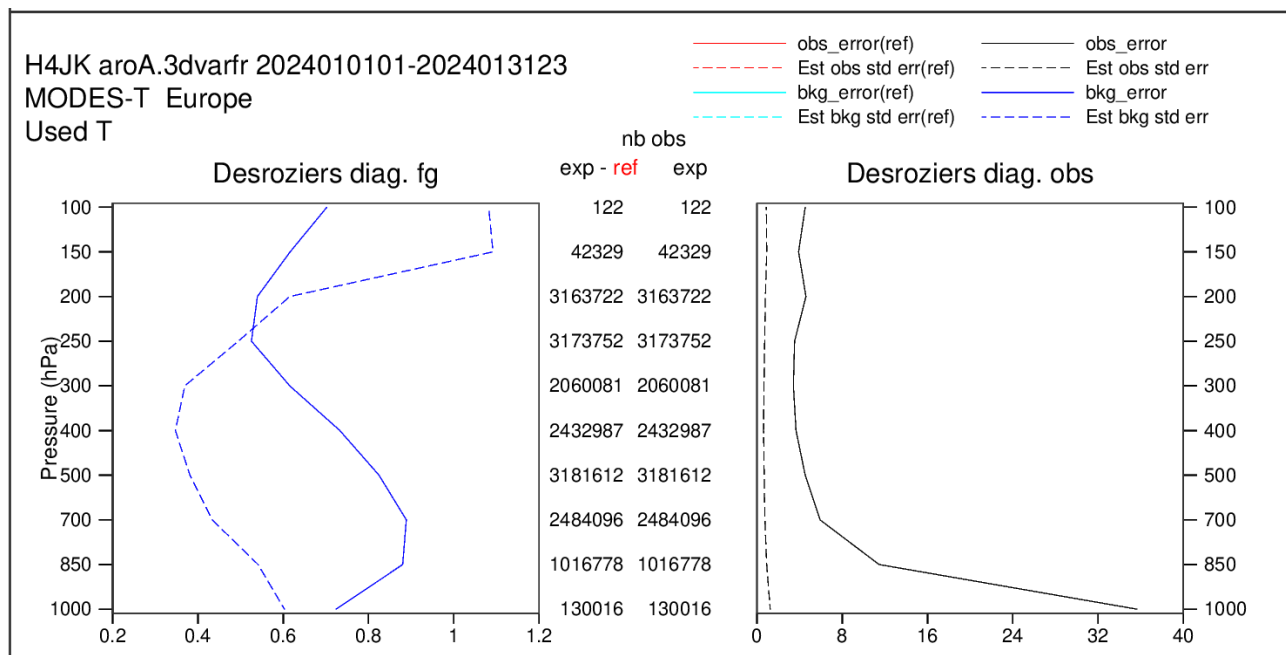


Figure 2: cy46 results, RFIND\_CLSTR = 4500m  
 no reference on the graph  
 experiment: with adaptive weighting

# All-sky IR assimilation of SEVIRI and FCI in AROME-Austria

Adhithiyan Neduncheran<sup>1</sup>, Florian Meier<sup>1</sup>

<sup>1</sup> GeoSphere Austria, Vienna, Austria

## 1 Introduction

---

In recent years, significant advancements were made in the assimilation of satellite-based infrared radiances for improving weather prediction models, particularly in conditions where clouds are present. Traditional data assimilation methods have typically filtered out cloud-affected observations due to the complexities involved in accurately modelling cloud properties and interactions. However, as the understanding and modelling capabilities of cloud microphysics improve, there is a growing interest in exploiting all-sky infrared radiance data, which includes both clear-sky and cloud-affected observations, for more comprehensive and accurate weather forecasting.

The assimilation of all-sky infrared radiances has emerged as a crucial approach for enhancing weather forecasts by incorporating cloud-affected observations alongside clear-sky data. A primary challenge in all-sky data assimilation lies in accurately characterizing observation errors, which vary due to the complex nature of cloud interactions and atmospheric conditions. Building on foundational work by *Okamoto et al. 2013* and *Geer and Bauer 2011*, this study implements an error model that adjusts observation error based on cloud amount. The error model considers the increased uncertainty in cloud-dense areas by dynamically assigning higher observation errors, while clearer regions are associated with lower error values, as *Geer and Bauer 2011* highlight the need for spatially adaptive error handling in all-sky conditions. The proposed model was tested within AROME-Austria CY48T1\_op1 (and CY49T1) in the OOPS framework with the usage of Water Vapour (WV) channels 6.2 and 7.3  $\mu\text{m}$  onboard MSG SEVIRI (and MTG FCI), evaluating its impact on obtaining a better representation of initial conditions and clouds. Results indicate that cloud-adaptive observation errors lead to improved assimilation of cloud-affected radiances, yielding more accurate and a refined depiction of atmospheric states. This work advances the application of all-sky radiance assimilation by optimizing error representation, aiming to capture cloud influence in numerical weather prediction effectively.

## 2 Data and model

---

MSG SEVIRI observations are used for the assimilation experiments. Only the WV channels are assimilated in all-sky conditions, while the other channels were assimilated in clear-sky conditions as done in operations.

The tests were carried out in AROME-Austria model having a grid spacing of 2.5 km with 90 vertical levels, centred over the Alpine region. IFS HRES boundary conditions is used in all experiments. RTTOV v12.2 is integrated within the model. Tests were also made in the CLAEF-1k, a 16 member ensemble prediction system at 1 km grid spacing.

## Comparison of simulated brightness temperatures

One case is evaluated in this study, 08-05-2023, init time 12 UTC.

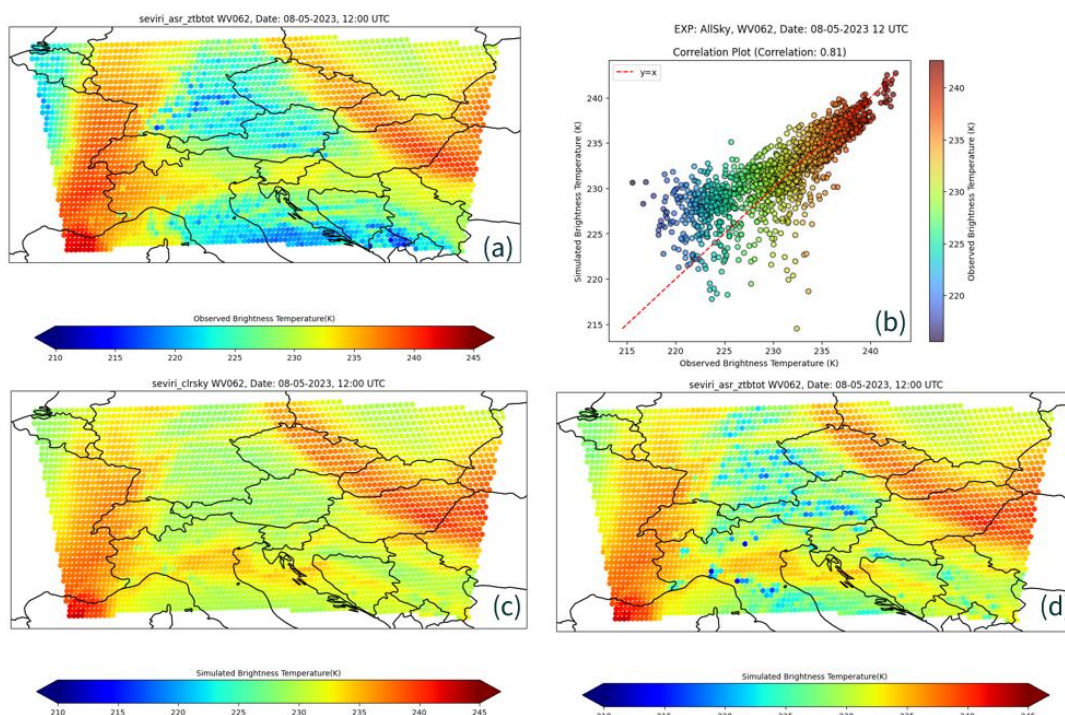


Fig. 1: All plots correspond to the case of 08-05-2023, WV channel  $6.2 \mu\text{m}$ : (a) Observed brightness temperature, (b) correlation between observed and all-sky simulated brightness temperature, (c) Clear sky simulated brightness temperature, (d) all-sky simulated brightness temperature.

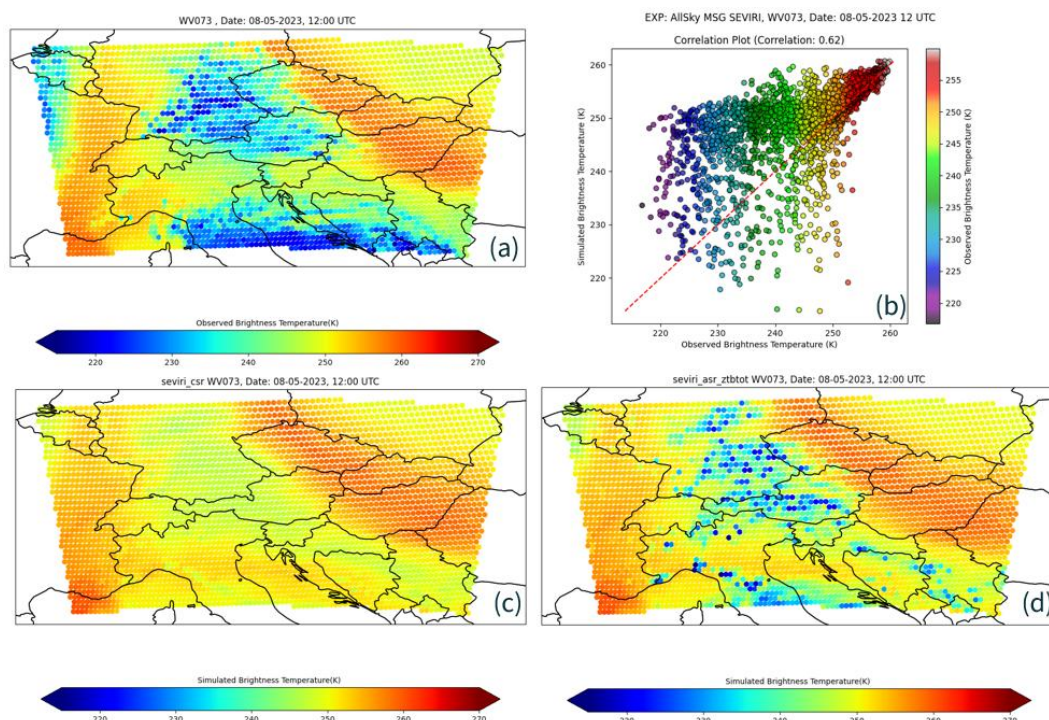


Fig. 2: All plots correspond to the case of 08-05-2023, WV channel  $7.3 \mu\text{m}$ : (a) Observed brightness temperature, (b) correlation between observed and all-sky simulated brightness temperature, (c) Clear sky simulated brightness temperature, (d) all-sky simulated brightness temperature.

### Cloud dependent observation error

The cloud amount(Ca) is calculated based on an approach by *Okamoto 2017* as below

$$Ca = (|Cm| + |Co|)/2,$$

where Cm and Co are cloud effect on model and observation, respectively, written by

$$Cm = B - Bclr,$$

$$Co = O - Bclr,$$

Where O and B are observed and (all-sky) simulated brightness temperature (BT), respectively. Bclr is background ClearSky BT when the cloud-scattering calculation is switched off in RTTOV.

The estimation of cloud amount is done using the window channel IR10.8  $\mu\text{m}$ .

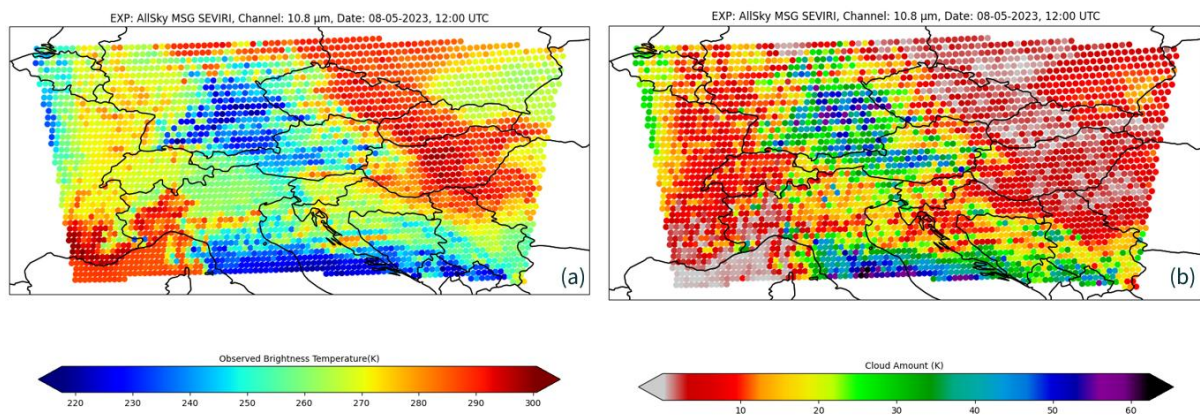


Fig. 3: (a) Observed brightness temperature for IR10.8  $\mu\text{m}$ , (b) Cloud amount calculated using the channel IR10.8  $\mu\text{m}$ .

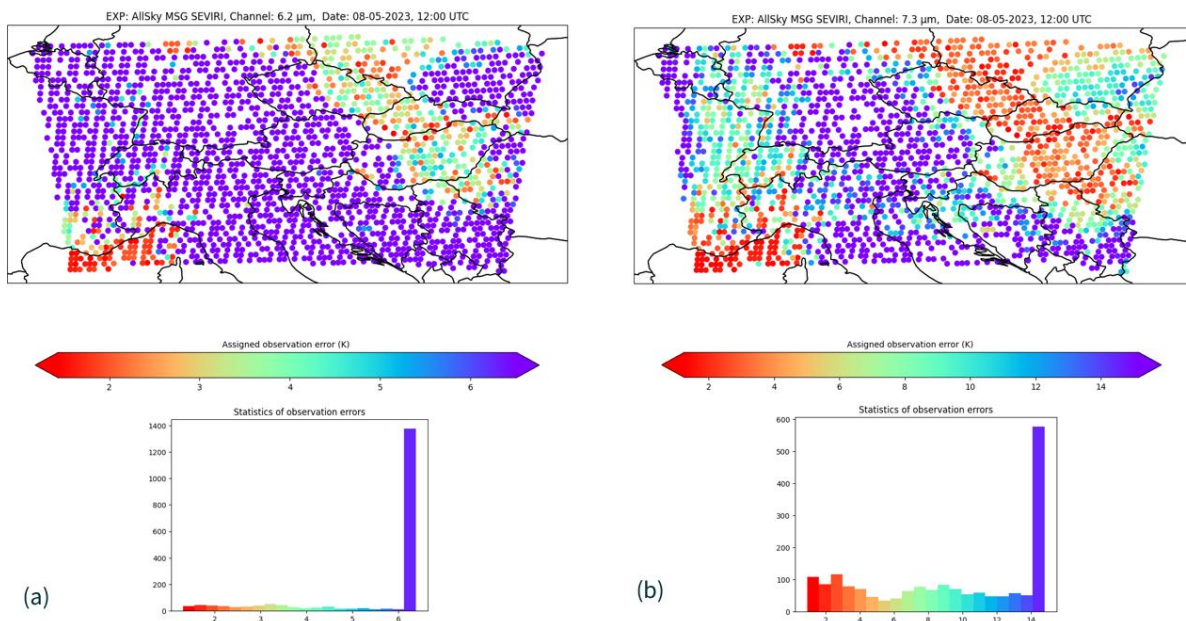


Fig. 4: Cloud amount dependent observation errors for (a) WV channel 6.2  $\mu\text{m}$ , (b) WV channel 7.3  $\mu\text{m}$

Figure 4 shows the distribution of the observation errors in all-sky assimilation which is now based on the calculation of cloud amount that is calculated dynamically using the window channel IR10.8  $\mu\text{m}$  and the error for a specific observation is assigned accordingly. The statistics of observation having different errors are shown in Figure 5 the respective channels, which are finally assimilated.

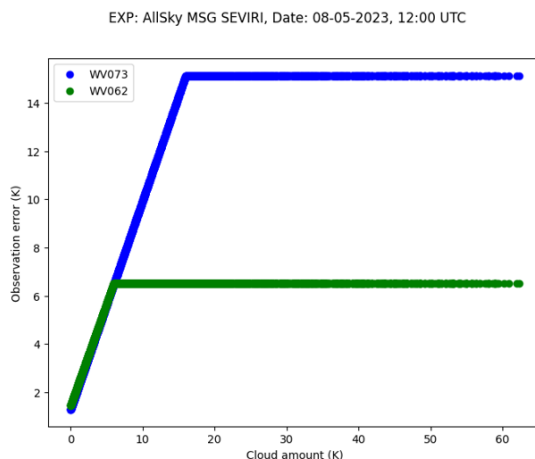


Fig. 5: Cloud amount vs observation error calculated based on the liner model proposed by Okamoto 2017

Earlier, all-sky assimilation experiments with SEVIRI water vapour channels established a robust framework for handling cloud-affected radiances through dynamically adaptive, cloud-dependent observation error modelling. Same principles are applied to FCI water vapour channels, enabling a consistent assessment of how these observations interact with existing assimilation systems. FCI’s enhanced spatial resolution, improved radiometric accuracy, and expanded spectral sampling promise greater information content but also introduce increased representativeness and cloud-related uncertainties, reinforcing the value of adaptive error characterisation.

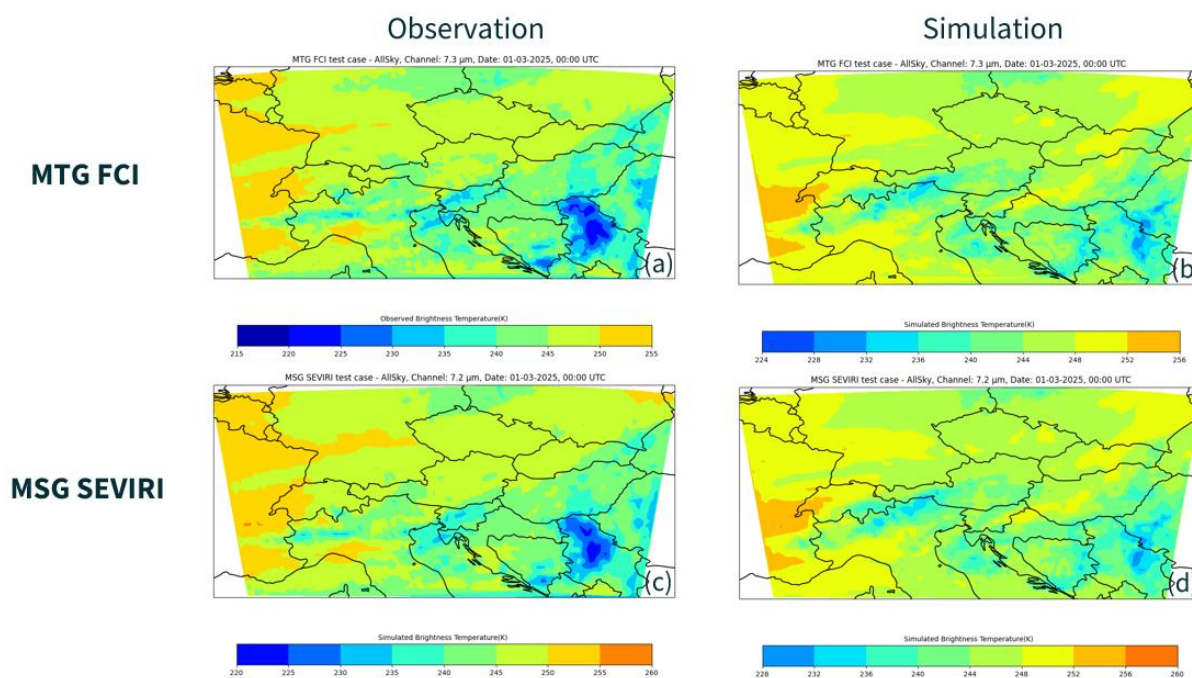


Fig 6: WV7.3  $\mu\text{m}$  channel from MTG FCI(top) and MSG SEVIRI (bottom), (a & c) Observed brightness temperature, (b & d) all-sky simulated brightness temperature

### 3 Discussion

---

**Impact of MTG FCI Characteristics:** The finer spatial resolution (1 km) and improved radiometric accuracy of MTG FCI is expected to enhance the simulation and assimilation of cloud-affected radiances. Compared to SEVIRI, MTG FCI provides more detailed cloud structures, which improves the realism of initial conditions in AROME Cy49T1.

**Role of cloud-dependent error modelling:** The adaptive error approach is critical for all-sky assimilation. Without dynamic error scaling, the assimilation system would either overfit uncertain cloudy observations or discard them entirely. By linking observation error to cloud amount, the system achieves a balance between utilizing valuable information and maintaining stability.

These findings confirm that all-sky assimilation of MTG FCI radiances is a promising step toward next-generation NWP systems. The approach leverages high-resolution satellite data and advanced error modelling to better capture cloud-radiation interactions, which are crucial for convective forecasting.

This work marks an important step toward operational use of MTG FCI radiances in high-resolution NWP systems. By implementing all-sky infrared assimilation in AROME Cy49T1 and applying a cloud-dependent observation error model, the study addresses the challenge of representing uncertainty in cloudy conditions. Currently, inter-channel error correlations between water vapour channels are handled through simplified approaches, such as block-diagonal structures and dynamic error scaling, rather than full covariance matrices. These innovations improve the assimilation of cloud-affected radiances, enhance humidity and cloud representation in the initial state, and demonstrate the practical pathway for exploiting next-generation geostationary satellite data in operational forecasting.

### 4 References

---

Okamoto, K., McNally, A.P. and Bell, W. (2014), Progress towards the assimilation of all-sky infrared radiances: an evaluation of cloud effects. *Q.J.R. Meteorol. Soc.*, 140: 1603-1614. <https://doi.org/10.1002/qj.2242>

Geer, A.J. and Bauer, P. (2011), Observation errors in all-sky data assimilation. *Q.J.R. Meteorol. Soc.*, 137: 2024-2037. <https://doi.org/10.1002/qj.830>

Okamoto, K. (2017), Evaluation of IR radiance simulation for all-sky assimilation of Himawari-8/AHI in a mesoscale NWP system. *Q.J.R. Meteorol. Soc.*, 143: 1517-1527. <https://doi.org/10.1002/qj.3022>

Wastl, Clemens, et al. "C-LAEF: convection-permitting limited-area ensemble forecasting system." *Quarterly Journal of the Royal Meteorological Society* 147.735 (2021): 1431-1451.

Geer, A. J., et al. "The growing impact of satellite observations sensitive to humidity, cloud and precipitation." *Quarterly Journal of the Royal Meteorological Society* 143.709 (2017): 3189-3206.

Geer, Alan J., Stefano Migliorini, and Marco Matricardi. "All-sky assimilation of infrared radiances sensitive to mid-and upper-tropospheric moisture and cloud." *Atmospheric Measurement Techniques* 12.9 (2019): 4903-4929.

## **5 Acknowledgement**

---

This work was carried out under the project “CloudyRadiances” at GeoSphere Austria funded by Austrian Federal Ministry of Climate Action, Environment, Energy, Mobility, Innovation and Technology, via the ASAP 18 call from FFG, Austrian research promotion agency grant No. 892666 in collaboration with the Institute for Meteorology and Geophysics at University of Vienna, Austria. We’d also like to thank Meteo France for hosting two scientific research stays at CNRM/GMAP/OBS funded by ACCORD.

# 3DEnVar with flow-dependent background error covariances

Chahrazed BOUZERMA<sup>(1)</sup>, Idir DEHMOUS<sup>(2)</sup>

<sup>(1)</sup> Numerical Weather Prediction Department, Office Nationale de la Météorologie, Algiers, ALGERIA

<sup>(2)</sup> Meteorological and climatological research Department, Royal Meteorological Institute, Brussels, BELGIUM.

## 1. Introduction

---

Data assimilation is a fundamental component of modern numerical weather prediction. Although 3DVar has been extensively used in operational systems, its reliance on static background error covariances limits its ability to represent flow-dependent atmospheric uncertainties. Three-dimensional ensemble variational methods (3DEnVar) overcome this limitation by incorporating ensemble-based, flow-dependent background error covariances, enabling a more realistic and situation-dependent spread of observational information and an improved representation of anisotropic structures. The implementation of such advanced methods is facilitated by the Object Oriented Prediction System (OOPS), which provides a flexible and modular software framework.

In this context, the present work aims to develop and test a prototype 3DEnVar system within OOPS, using flow-dependent covariances derived from an ensemble of perturbed ARPEGE forecasts downscaled to ALARO Belgium model, without considering observational perturbations at this stage.

## 2. Practical implementation

---

### 2.1. Methodology

In order to progress towards 3DEnVar, we first successfully implemented and tested a simple 3D-Var experiment within the *OOPS* framework. The main goal of this step was to become familiar with the implementation of data assimilation within *OOPS*. The detailed outcomes of these preliminary experiments will be presented in the following sections.


After the implementation of a simple 3D-Var, the next step was to move on to the 3D-Ensemble-Var (3DEnVar). This approach differs from the classical 3D-Var mainly at the minimization stage, as it uses an ensemble of forecasts from multiple members as input. Consequently, implementing 3DEnVar also requires modifications of the *oops.json* configuration file, including the description of the ensemble forecasts and options specific to ensemble processing. These aspects, along with the 3DEnVar scheme tuning, the perturbation method, and the experimental setup will be detailed later on.

## 2.2. Experimental Setup

Our experiments with 3DVar and 3DEnVar have been carried out with the ALARO model. The model setting is shown in the *Table 1*, while *Figure 1* illustrates the spatial domain covered by the model :

*Table 1: Summary of the ALARO model configuration and experimental environment*

<b>Model version</b>	Belgium ALARO-1.3km
<b>Resolution</b>	1.3 km
<b>Levels</b>	87
<b>Number of gridpoint</b>	576 x 576
<b>Cycle</b>	Cycle 46
<b>Domain center</b>	(4.55°E ; 50.80°N)
<b>Coupling</b>	ARPEGE (7 km)
<b>Surface initialisation</b>	None
<b>Expriment environment</b>	Belenos (MF machine)
<b>Framework</b>	OOPS



*Figure 1: Belgium-ALARO-1.3km Domain*

## 3. Comparison of 3D-Var (OOPS) Implementation vs MASTERODB:

Using the CY46 3D-Var scripting system on Belenos ( provided by A. Trojakova and B. Strajnar ) under the directory /home/gmap/mrpa/strajnarb/sample\_3dvar\_cy46t1, we successfully performed the OOPS-based 3D-Var minimization.

This setup allowed us to follow their workflow and reproduce their results by comparing the two binaries, OOVAR and MASTERODB.

The objective was to check that the OOPS and non-OOPS executables produced nearly identical results. This verification was successfully achieved, as demonstrated in the following Single AMDAR Observation experiment, which compares the 3D-Var OOVAR and 3D-Var MASTERODB runs.

### 3.1. Single AMDAR Observation Assimilation Experiment:

To evaluate the implementation of the OOPS-based 3D-Var minimization and compare it with the 3D-Var system using MASTERODB, we conducted an assimilation experiment using a single AMDAR observation in BUFR format located at 48.72° N, 7.18° E, at 10.500 m altitude. The observation, corresponding to 02 February 2025 at 00 UTC was used to assess the impact of assimilating temperature (T) and horizontal wind (U) parameters.

The increments (analysis - first guess) of temperature (T) and horizontal wind component (U) at the 19th model level are shown in *Figures 2 and 3* respectively, with OVAR results on the left and MASTERODB on the right.

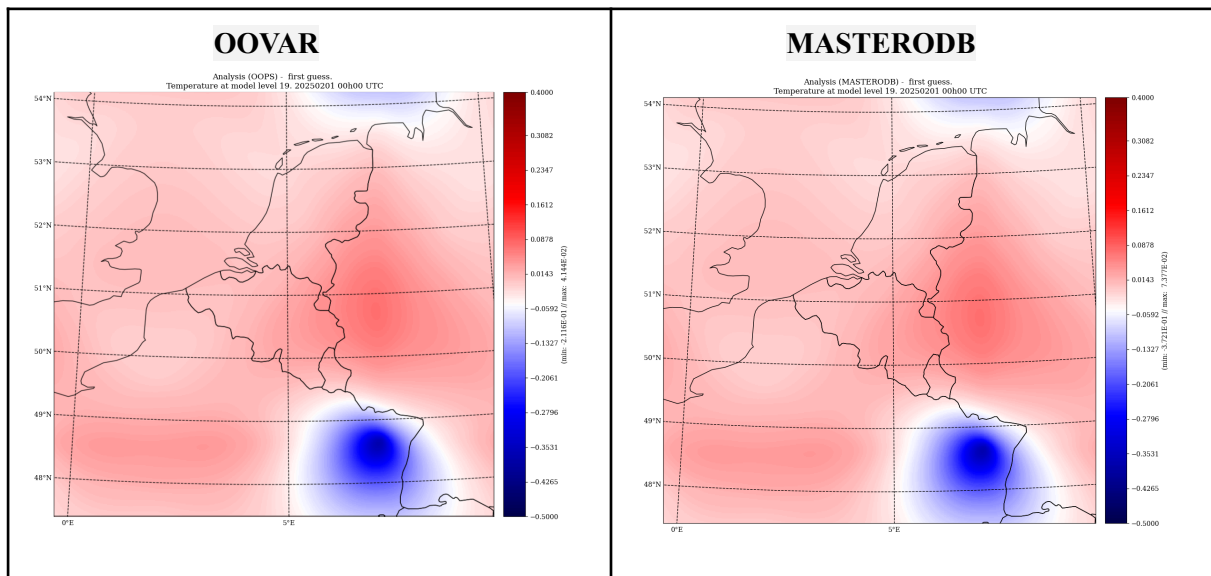


Figure 2: Temperature Increment at level 19

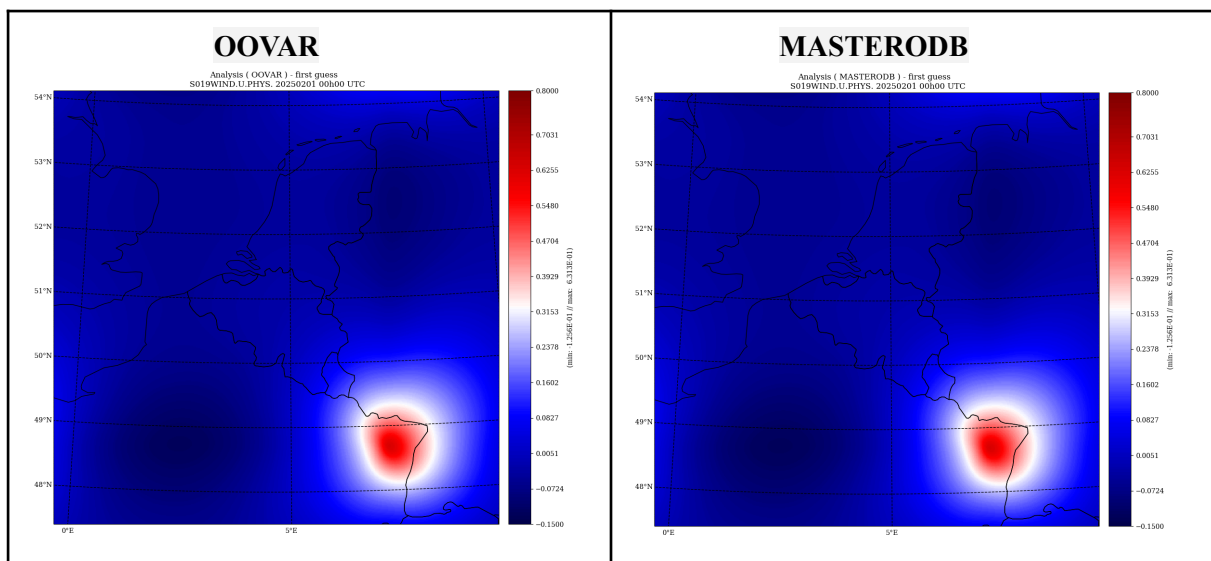


Figure 3: Horizontal wind Increment at level 19

It can be observed that:

- For both parameters, the results from OOVAR and MASTERODB are nearly identical.
- The assimilation of this AMDAR observation has a localized impact around the observation point, leading to a *negative temperature increment*, This indicates that the assimilation produces a colder analysis, effectively cooling the model state within ~150 km of the observation location. In addition, the positive increment in the zonal wind component (U) shows that the assimilation strengthens the zonal wind by increasing the U-component in the analysis in a radius of about 100 km around the observation point.
- However, as identified by RC-LACE colleagues in their initial comparison between the MASTERODB and OOVAR binaries using AROME 3DVAR, there is a difference between the two binaries, particularly for specific humidity (q), where noticeably different results are observed. The impact of assimilating the AMDAR observation with MASTERODB is characterized by a smooth, well-defined structure centered on the observation point, with a positive increment in its vicinity. In contrast, the impact structure with OOVAR is not discernible, appearing noisy and spurious across all model levels. The magnitude of the increments in MASTERODB is roughly 100 times larger than in OOVAR. This discrepancy is illustrated in *Figure 4*, which shows the q increment at the 67th model level.

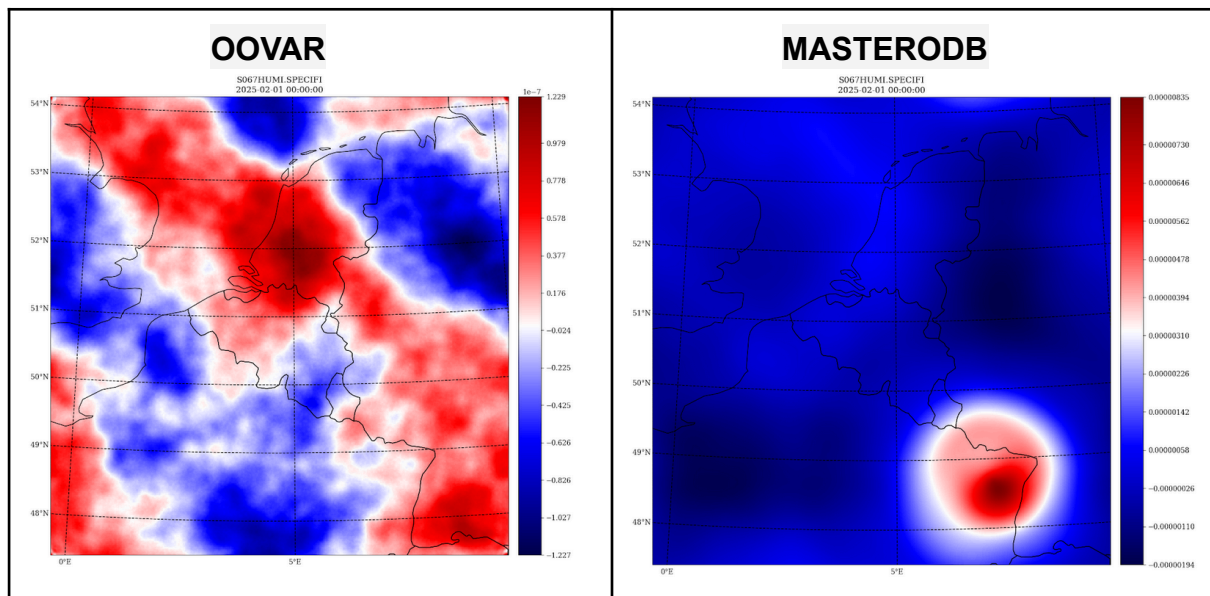


Figure 4: Increment of Spec.Hum at level 67

## 4. Implementation of 3DEnVar within the OOPS framework

### 4.1. 3DEnVar Scheme Tuning

Building on the successful evaluation of the OOPS-based 3D-Var minimization, the next step was to explore the 3DEnVar scheme.

In 3DEnVar, several parameters need to be carefully tuned. We are aware that these parameters are strongly interdependent, and that modifying one of them can influence the optimal values of the others. Considering these interactions, and given the limited time available within the framework of our internship, our main objective is to set up and take in hand a functional prototype of the 3DEnVar system. The tuning therefore relies on parameter values taken from previous studies, notably (Brousseau et al, 2025), where:

- The most efficient horizontal localization length scale was identified as a configuration varying from 25 km at the lowest vertical level to 150 km in the highest one.
- The optimal vertical localization length scale was set to  $0.2 \ln(hPa)$ .

Additionally, one of the key parameters to be configured is the inflation factor of the background errors. Based on previous work on the tuning of the ALARO-Belgium 3DVar system at RMI (Dehmous, 2025), using the Desroziers diagnostic method (Desroziers et al., 2005) :

- The optimal value for inflation factor was found to be 1.05, in order to moderately enhance the ensemble spread and mitigate potential under-dispersion issues.

Whereas, we are also aware that further adjustments and sensitivity tests should be considered in future work, once the prototype has been fully implemented, regarding the identification of the most optimal values for our configuration.

The chosen settings are illustrated in *Figure 5* below, which presents the corresponding configuration structure in the *oops.json* file used for the 3DEnVar experiments.

```

"Jb" {
  "Background" {
    "state" [
      {
        "date" "2024-08-01T00:00:00Z"
        "term" "PTH"
        "ifile" "0"
        "variables" "0"
        "expver" "OOPS"
      }
    ]
  }
  "ModelBias" {}
  "ObsBias" {
    "filename" "no-varbc"
  }
  "Covariance" {
    "covariance" "ensemble"
    "localization" {
      "loc_type" "1"
      "levl_001" "1"
      "hlen_001" "150000"
      "levl_002" "88"
      "hlen_002" "25000"
      "v_length_scale" "0.2"
    }
  }
  "ensemble" [
    {
      "members" "32"
      "inflation" "1.05"
      "variables" "1"
      "sub_ens" "1"
      "state" [
        {
          "ifile" "1"
          "expver" "unused"
          "filename" "/scratch/work/bouzermac/bxl_cy46_alaro13_3denvar/data/gp_field/2024080100/ICMSHA013_0000_mb001_gp"
          "dont_check_last_file" "yes"
        }
        {
          "ifile" "1"
          "expver" "unused"
          "filename" "/scratch/work/bouzermac/bxl_cy46_alaro13_3denvar/data/gp_field/2024080100/ICMSHA013_0000_mb002_gp"
          "dont_check_last_file" "yes"
        }
        {
          "ifile" "1"
          "expver" "unused"
          "filename" "/scratch/work/bouzermac/bxl_cy46_alaro13_3denvar/data/gp_field/2024080100/ICMSHA013_0000_mb003_gp"
          "dont_check_last_file" "yes"
        }
      ]
    }
  ]
}

```

Figure 5: Fragment of the *oops.json* Configuration Detailing the 3DEnVar Settings

### 4.2. Perturbation method

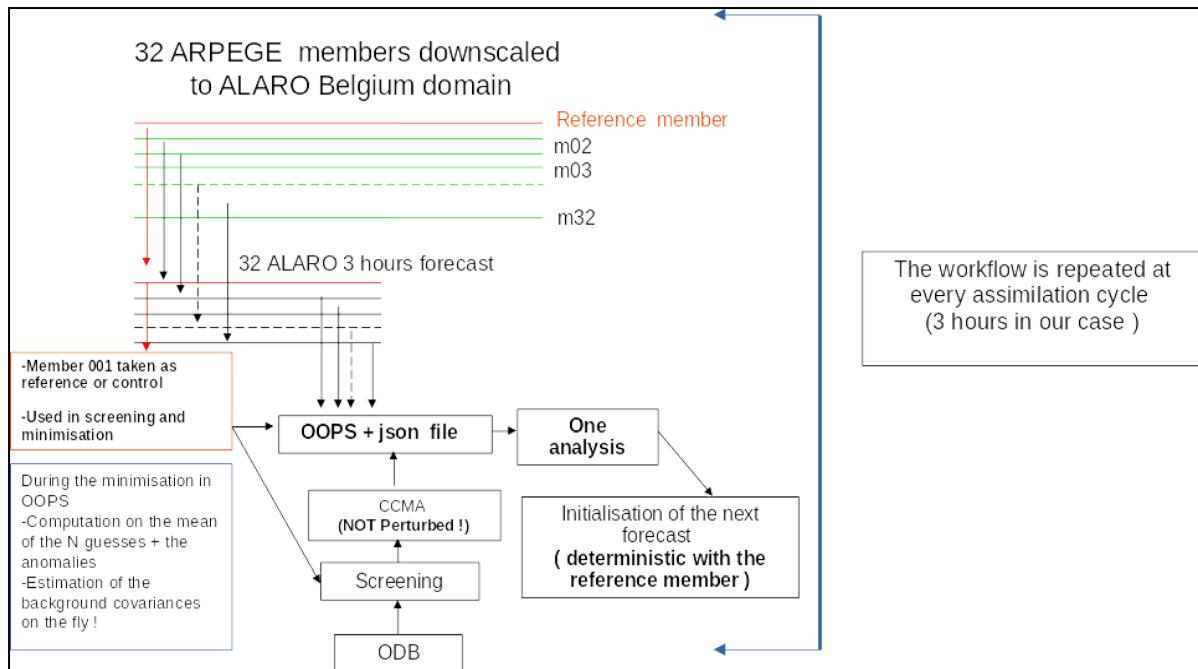


Figure 6: Overview of the Perturbation Approach

Our experiments used an ensemble of forecasts from ALARO at 1.3 km resolution with a 3-hour-cycling, whose initial and lateral boundary conditions are provided by a set of perturbed ARPEGE coupling files from the Météo-France 4DEnVar-EDA system.

The performance of ensemble data assimilation is known to be highly sensitive to ensemble size, particularly for sizes below approximately 100 members (Houtekamer and Zhang, 2016). However, in practice, the ensemble size is constrained by computational resources and by the availability of ensemble boundary conditions. Since the ARPEGE ensemble currently provides a maximum of 50 members, the ensemble size in this study was set to 32. This represents a compromise between computational feasibility and the expected benefit of increasing ensemble size.

The 32 ARPEGE ensemble coupling files are downscaled into 32 ALARO 1.3 km 3-hour forecasts. The forecast from the first member (mb001) is retained as the reference or control member, constituting the first guess during both the screening and minimization steps. It is specified in the *oops.json* configuration file as **ICMSHOOPSINIT**, and is also listed in the “ensemble” section together with the remaining 31 members.

The minimizer used during the minimization was DRPCG (Derber-Rosati Preconditioned Conjugate Gradients) since it is the one which offers more verbosity compared to Plancoz algorithm for example.

The estimation of the background error covariance is flow-dependent and computed on the fly using the ensemble members. During both the screening and minimization steps, the observations used are not perturbed. The resulting analysis is then used to initialize the subsequent forecast.

Ideally, the full system would then run in a 3-hour assimilation cycle, but due to time constraints during the internship, it was not possible to complete the implementation of the cycling.

## 5. Comparison of 3DVar-OOVAR vs 3DEnVar

### 5.1. Assimilation experiment with Two Single SYNOP Observations

With the 3DEnVar setting shown in figure 5, the next step was to evaluate the implementation of the OOPS-based 3DEnVar minimization and compare its behavior with the classical 3D-Var system implemented through OOPS.

In order to achieve this, an assimilation experiment using two SYNOP observations in BUFR format from stations at distinctly different elevations has been conducted ( Koksijde :10 m and Maastricht : 150 m), both valid on 01-08-2024 at 00 UTC. This elevation contrast tests the impact of orography-induced background errors: while 3D-Var's static covariances spread corrections uniformly, 3DEnVar's flow-dependent B dynamically represents altitude-sensitive error structures and their vertical propagation.

The increments for the temperature at model level 87 are shown in *Figure 7* below, representing 3DEnVar increments obtained with three ensemble sizes (4, 16 and 32 members) compared with corresponding 3D-Var increments from the same experiment. For visual comparison, the 3D-Var increments were scaled down by a factor of 10 because their amplitude is considerably larger than those from 3DEnVar.

In addition, *Figure 8* illustrates the increments of the horizontal wind component U at model level 87, comparing the 3D-Var and 3DEnVar (using 32-member) experiment.

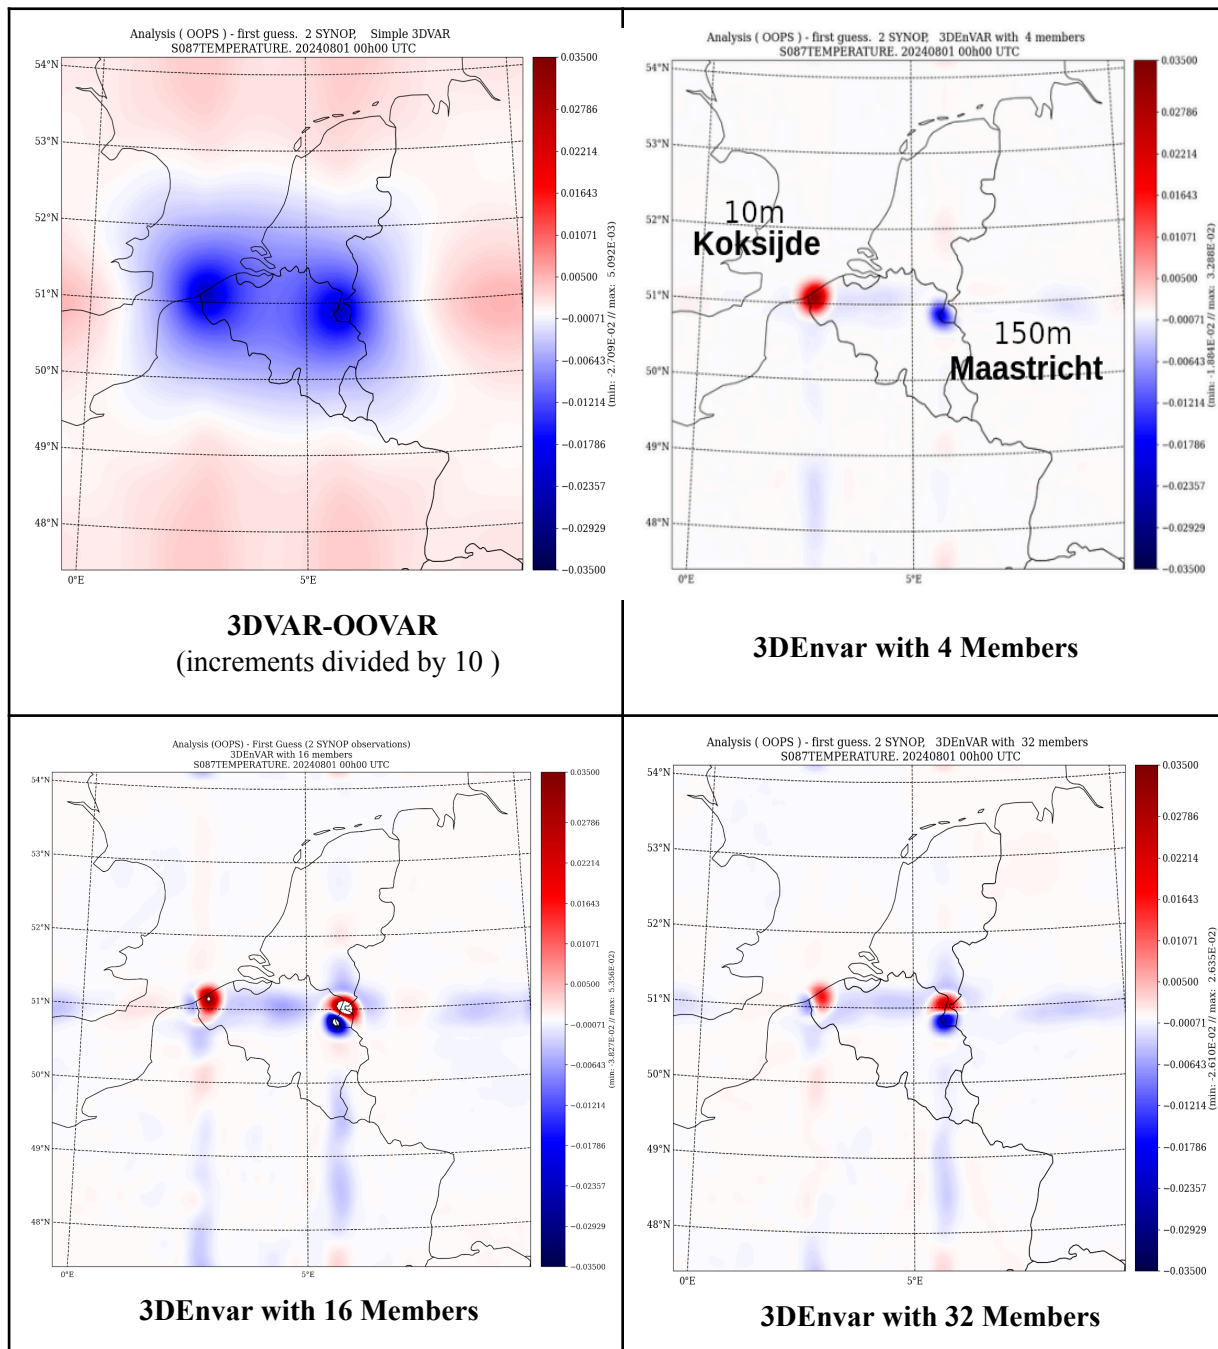


Figure 7: Temperature increments at level 87 for 3DVAR-OOVAR experiment compared with 3DVAR's using ensembles of 4, 16 and 32 members.

The 3D-Var temperature increment displays wide, smoothly damped negative values around both observation sites, reflecting a cooling of the analysis extending roughly ~200 km from the observation locations. Beyond this radius of influence, the increment reverses sign and becomes positive, indicating a tendency of the assimilation of the two synop observations to produce a warmer analysis farther away. The overall homogeneous and overly smooth structure highlights 3D-Var's tendency to diffuse the observational impact uniformly across the domain.

By contrast, 3DENVAR produces spatially confined, punctiform increments with a much smaller radius of influence (~60 km at the surface). 3DENVAR also yields distinct increment structures and signs at the two stations, whereas 3D-Var produced similar negative increments at both sites. With the 4-member ensemble, 3DENVAR gives a positive increment at Koksijde (10 m), indicating a local warming, and a localized negative increment at Maastricht (150 m), indicating local cooling.

Increasing the ensemble to 32 members changes the local structure further, producing opposite-signed lobes at the same observation point, while the rest of the domain shows minimal increments.

The generally smaller amplitude of 3D $\text{EnVAR}$  increments compared with 3D-Var may partly result from the chosen horizontal localisation (surface length scale = 25 km), which strongly limits the spatial extent and magnitude of ensemble-derived covariances.

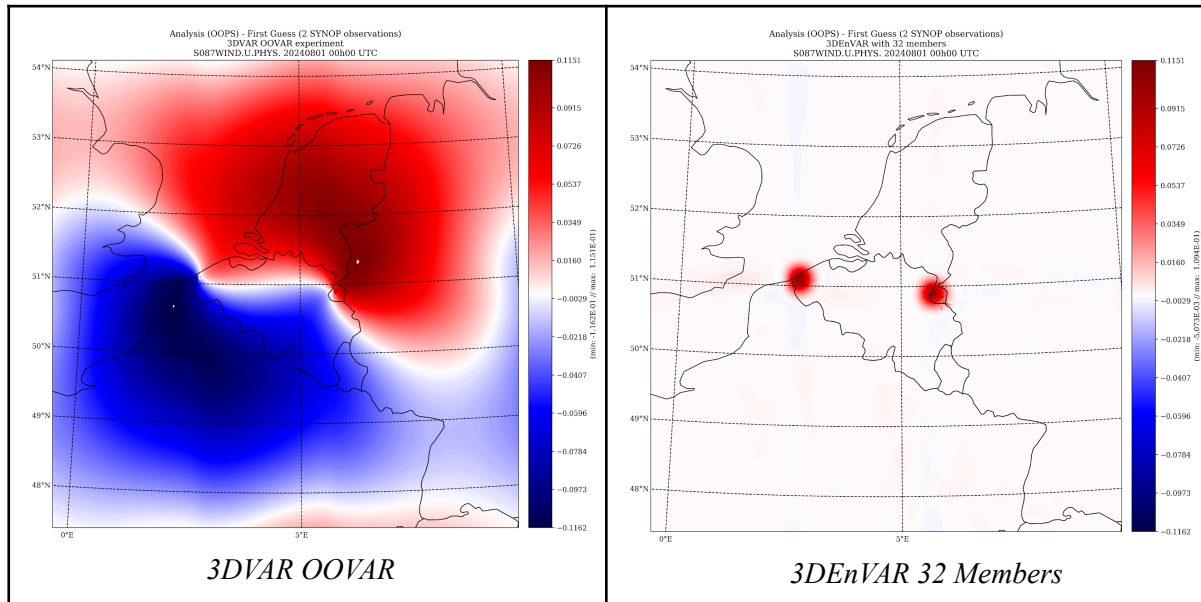


Figure 8: Horizontal wind component increments at level 87 for 3DVAR-OOVAR experiment compared with 3D $\text{EnVAR}$ 's 32 members.

The 3D-Var U-wind increment at model level 87 exhibits a wide, smooth, and nearly uniform dipole structure dividing the domain into two large regions of opposite sign. Around Koksijde (10 m), the increment is negative, indicating a reduction of the zonal wind component, with the influence extending over the southern part of the domain and gradually weakening with distance. Conversely, the assimilation around Maastricht (150 m) SYNOP observation induces a positive increments with the strongest amplitude near the observation point and a smooth decay outward. This large-scale pattern reflects the nature of the 3D-Var background-error covariances, which spread corrections uniformly across the domain.

In contrast, the 3D $\text{EnVAR}$  experiment with a 32-member ensemble produces localized positive increments around both observation sites, confined within a small radius of influence (~50 km). Beyond this distance, the increments vanish rapidly, indicating a more flow-consistent and spatially constrained adjustment. At Maastricht (150 m), the increment field exhibits a slightly granular structure characteristic of 3D $\text{EnVAR}$ .

### 5.2. Assimilation experiment using all available observation types across the domain

To further evaluate the behavior of our 3DVar setup using the largest ensemble size (32 members), and to compare it with 3DVar-OOVAR, an additional experiment has been carried out. The analysis date and time was 2024-08-01 at 00 UTC, assimilating all available observation types within the ALARO Belgium domain (SYNOP, TEMP, GNSS, and AMDAR, all provided in BUFR format).

The *Figures 9 and 10* present, respectively, the temperature and zonal wind (U) increments at model level 76 for both the 3DVar and 3DVar configurations.

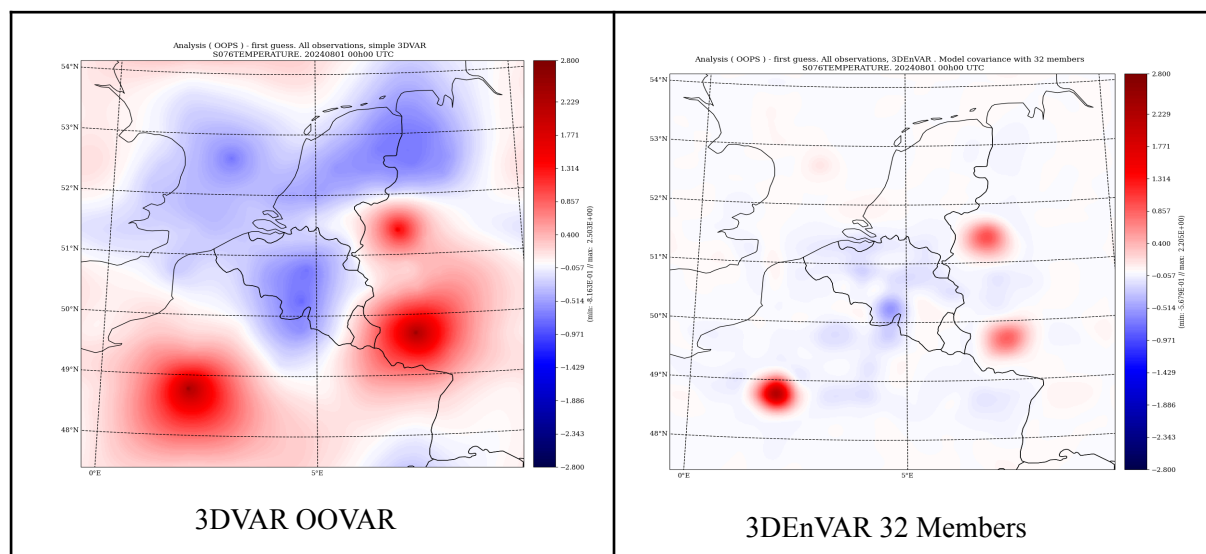


Figure 9: Temperature Increment at level 76 - Comparison of 3DVAR vs 3DVAR using 32members

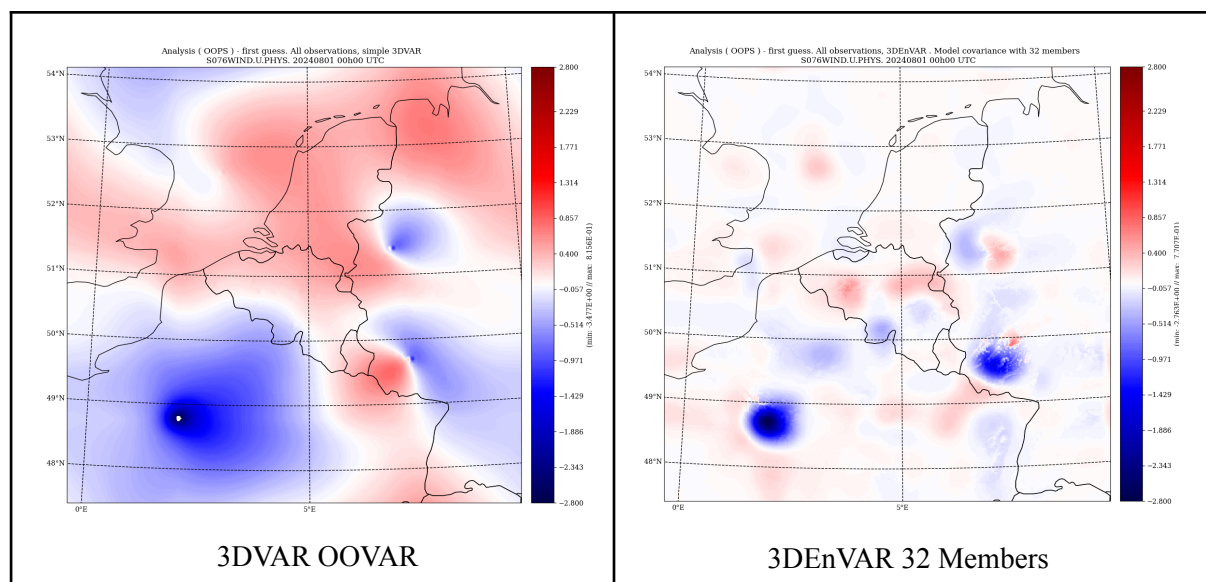


Figure 10: Horizontal wind U Increment at level 76 - Comparison of 3DVAR vs 3DVAR using 32members

Overall, the results are consistent with the previous experiment with the two isolated SYNOP observations. Compared with the smooth, domain-wide structures produced by 3D-Var, the 3DVar increments display a much more fragmented and irregular texture. The effect is particularly striking for the U-wind increments, where the 3DVar increments show strong local variations in stark contrast with the broad, continuous patterns of 3D-Var. This hachured appearance reflects the direct imprint of the flow-dependent sampling of the background-error covariances.

## 6. Conclusion and Perspectives

---

This work focused on the implementation and initial assessment of a 3DEnVar prototype within the OOPS framework for the ALARO 1.3 km Belgium configuration. After successfully reproducing a reference 3D-Var implementation using OOVAR and ensuring consistency with the classical MASTERODB-based system, the configuration was extended to incorporate flow-dependent background error covariances derived from a 32-member ensemble of downscaled ARPEGE forecasts. The setup relied on parameter values inspired by previous papers, particularly regarding horizontal, vertical localization and background error inflation.

The comparative assimilation experiments revealed clear structural differences between the two schemes. The 3D-Var system produced smooth, large-scale increments with uniform spatial spread, reflecting the inherent limitations of static climatological covariances. In contrast, 3DEnVar generated spatially confined, anisotropic and granular increments sensitive to local flow structures and orographic variability. These behaviors were consistent both in isolated-observation experiments and in full-domain assimilation involving SYNOP, TEMP, GNSS, and AMDAR datasets.

While these results align with theoretical expectations, several numerical aspects require further refinement. A systematic tuning of localization scales and inflation factors is required to optimize the configuration for ALARO Belgium. Moreover, discrepancies observed in humidity-related increments indicate the need to incorporate recent bug fixes developed at Météo-France to ensure a more consistent and physically realistic treatment of moist variables in the ensemble space. Finally, a comprehensive evaluation of forecast impact under a fully-cycling system remains necessary to assess the operational relevance of the prototype.

In summary, the implementation presented here provides a functional baseline for further developments of ensemble-based assimilation within the ALARO model. Although the current configuration reproduces the expected structural behavior of 3DEnVar, its practical impact on forecast skill has not been assessed due to time constraints during the internship and remains a key priority. The system should therefore be considered as a foundation for ongoing research rather than a finalized alternative to the operational 3D-Var framework.

## References

---

- Brousseau, P., Vogt, V., Arbogast, E., Martet, M., Thomas, G., & Berre, L. (2025). The operational 3DEnVar data assimilation scheme for the Météo-France convective scale model AROME-France.
- Desroziers, G., Berre, L., Chapnik, B., & Poli, P. (2005). Diagnosis of observation, background and analysis-error statistics in observation space. *Quarterly Journal of the Royal Meteorological Society*, 131, 3385-3396.
- Dehmous, I. (2025). *ALARO Belgium 3DVar – Tuning using Desroziers method and diagnostics*. [https://opensource.umr-cnrm.fr/attachments/download/6545/DAWW\\_alaro\\_3DVAR\\_belgium\\_tuning.pdf](https://opensource.umr-cnrm.fr/attachments/download/6545/DAWW_alaro_3DVAR_belgium_tuning.pdf)
- Houtekamer, P.L., & Fuqing Zhang. (2016). Review of the Ensemble Kalman Filter for Atmospheric Data Assimilation. *Monthly Weather Review*, 144, 4489–4532.

# Large eddy simulations with ALADIN input

Coupling MicroHH LES model to output from AROME CSC

P. Smerkol, N. Kastelec, J. Cedilnik  
Slovenian Environmental Agency  
(ARSO)

## 1 Introduction

---

When developing new physical schemes of turbulence for a meteorological model, we usually lack reference data against which we can compare the results, as measurements of, for example, third-order moments or 3D turbulence are rare. Results of large eddy simulations (LES) are in such cases often used as a reference, as the simulation is run on a much finer grid than the meteorological model, so it's dynamics can resolve finer turbulent processes which we want to parameterize in the meteorological model.

In most cases, large eddy simulations are run on idealized cases that have well known meteorological situations and often also dedicated measurements. Some well known such cases are for example BOMEX [2], GABLS [3] and DYCOMS-II [4].

At Slovenian Environmental Agency (ARSO), as a part of a masters degree [5], we tried a different approach. We ran large eddy simulations on a chosen subdomain of the domain of the meteorological model, and set the orography, initial and boundary conditions for the large-eddy model to the values interpolated from the meteorological model. In this way, we hoped that results of large eddy simulations would reflect actual meteorological situations even more.

Below, the meteorological model and large eddy simulation model are described, and then their results are compared for two different cases.

## 2 Cases

---

Based on meteorological situations and availability of data, we chose two cases:

### 2.1 Barje case

First interesting case has a domain centered on the flat moor south of Ljubljana called Ljubljansko Barje (see Fig.1 left). In a location on the moor, ARSO measures profiles of vertical wind components and virtual temperature with a Sodar-RASS system for heights up to 600 m above ground.

In the morning on the 4th of September 2025, there was a typical late summer situation with a weak ground inversion (300-400 m above ground) and a subsidence inversion (500-700 m above ground) and a very weak wind throughout the vertical cross-section of the atmosphere as can be seen from Sodar-RASS measurements (Fig.1 right). Because of this, fog occurred in the lower part, which then dissolved, leaving an adiabatic temperature profile in the higher levels and a minimal residual inversion in the lower levels for times up to 10:00 AM local time.

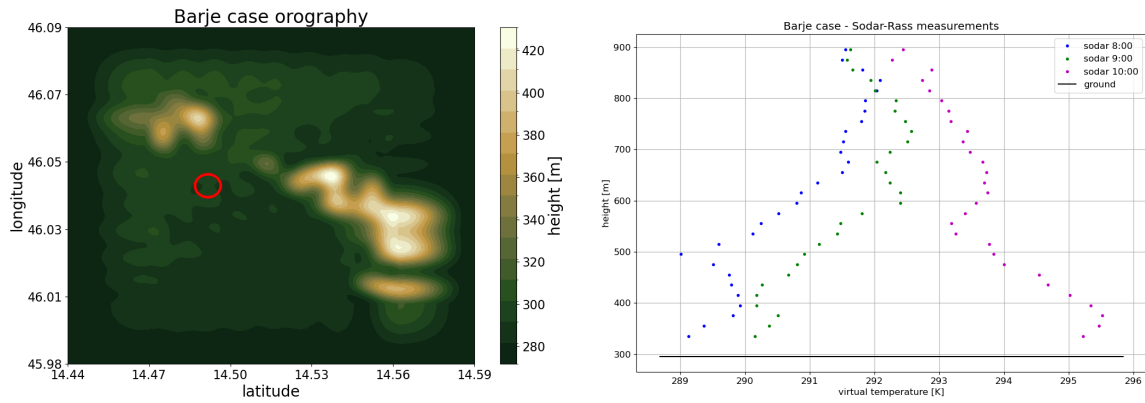


Figure 1: Orography and sodar measurements for Barje case. Left - orography, where red circle denotes the location of Sodar-RASS measurements, right - Sodar-RASS measurements at different times.

## 2.2 Planica case

The second chosen case has a domain centered on the site of the Planica ski jumping hill, situated in the Tamar alpine valley (see Fig.2). At the time of the ski jumping competition on 27th of March 2025 at 9:00 AM local time, there was an early spring situation, where there was a strong northeastern wind upward of 1000 m above ground and below a strong mixing occurred due to sunrise and rugged alpine orography. This resulted in wind gusts and a reversal of the general wind to a southwest direction in the lower part at the ski jumping hill site (from oral communication with the official ARSO meteorologist present at the competition).

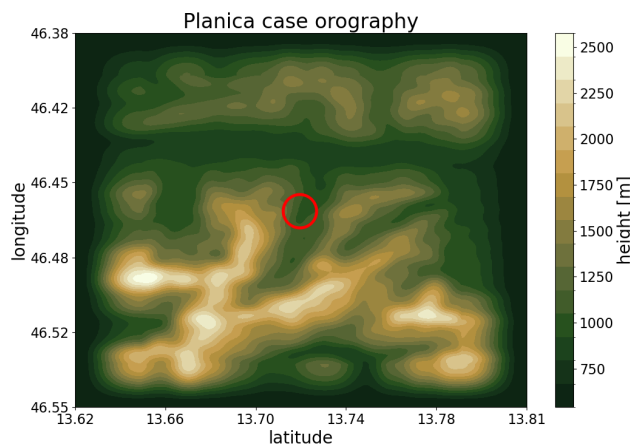


Figure 2: Orography for Planica case. The red circle denotes the location of the Planica ski jumping hill.

## 3 Meteorological model

Meteorological models used in both cases are AROME cy48t3 deode models [1] with 200 m resolution and 90 vertical levels, where initial and boundary conditions were taken from an AROME 500 m model, which was in turn coupled to the Digital Twin model. For the surface, we used SURFEX with Ecoclimap SG. Domains of

200 m and 500 m models for both cases are shown in figure 3.

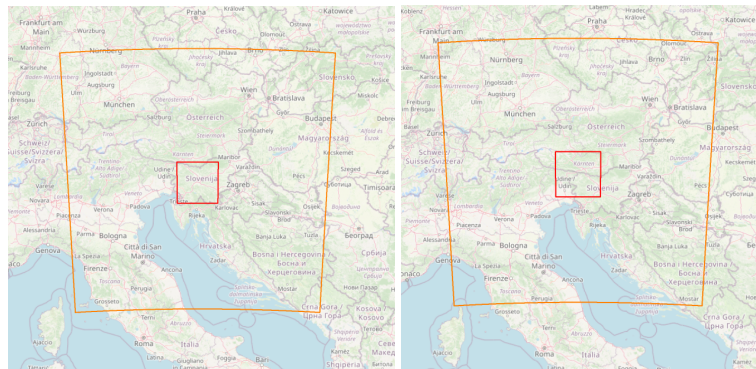


Figure 3: Domains of used meteorological models. Left - models used in Barje case, right - models used in Planica case. Red - domain of 200 m AROME cy48t3 deode 200m model used in analysis, orange - AROME 500 m model used for initial and boundary conditions of the 200 m model.

## 4 Large eddy simulation

For large eddy simulations, we chose the MicroHH model [6], a model specifically developed for simulating the atmospheric boundary layer. We chose it because its version 2 (in development) has an option to include an immersed boundary, which can be used to simulate orography. It also has an adaptive time step, which reduces when the CFL condition is above some value to ensure convergence.

It numerically solves equations for conservation laws of mass, momentum, energy and humidity in the Boussinesq approximation, using a 4-th order Runge-Kutta method on a grid that is doubly periodic in lateral boundary conditions. For parameterization of sub-grid processes it uses the Smagorinsky-Lilly model [7].

MicroHH also has various options for physical parameterizations, from which we chose a Monin-Obukhov similarity theory [8] based model for surface fluxes and scalar fields, an one-moment ice scheme for microphysics [9] capable of producing snow and cloud ice, and the GCSS parameterization for radiation [10], which is a simplified parameterization, but fast and numerically stable.

At the model top, a buffer layer is used, which absorbs the energy of gravitational waves, thus preventing them to reflect back into the domain.

MicroHH can be run either in parallel using MPI, or on a single GPU using CUDA.

## 5 Experiment setup

Both experiments were setup by first running the forecast of the meteorological model, starting with the 500 m models at 0:00 the day before case times, and the 200 m model a few hours before case times, to allow for spin up. In Table 1, model setups of the 200 m model for both cases are shown.

In the MicroHH models, we first interpolate the orography of the meteorological models to the MicroHH grid and set it as an immersed boundary in the model. Because of doubly cyclic lateral borders in MicroHH, we have to modify the orography at domain edges, as there orography is not continuous, which can result in artificial sharp gradients that lead to numerical instabilities.

*Table 1: Setup of 200 m meteorological models used in cases*

	Barje case	Planica case
Run start (UTC)	2025-09-03 12:00	2025-03-26 23:00
Vertical levels	90, hybrid	90, hybrid
Resolution [m]	200	200
Grid points (X,Y)	600, Lambert	648, Lambert
Domain center (lat, lon)	46.02, 14.48	46.48, 13.72

*Table 2: Setup of MicroHH models used in cases*

	Barje case	Planica case
Run start (UTC)	2025-09-04 03:00	2025-03-27 05:00
Domain size (x,y,z) [km]	10.80, 11.52, 5.12	14.80, 17.92, 7.04
Resolution (x,y,z) [m]	40, 40, 20	40, 40, 20
Number of levels (x,y,z)	272, 288, 256	352, 448, 352
Domain center (lat,lon)	46.04, 14.51	46.47, 13.72
Buffer zone width [km]	1.5	1.5
Initial time step size [s]	2	2

This is done by setting values at the edges to the minimum value of orography in the whole domain and then do a cosine smoothing in 10-gridpoint strips at all four lateral borders from the minimum value at the edges to the real orography value at the other end of the strip. The effect of this procedure can be clearly seen in Figs. 1 - left and 2. Minimum value at the edges was chosen to minimize the effects of the artificial orography in the center of the domains, where we compare results.

Next, 2D fields for surface boundary conditions, namely the liquid water potential temperature and total specific humidity, are interpolated to the MicroHH grid and smoothed at the domain edges to again avoid potential sharp gradients at the lateral boundaries. For wind components, a no-slip condition is set at the bottom boundary, so wind surface fields are not needed.

We next run an initial MicroHH model step, to obtain the result files with 3D fields, which we then switch with interpolated smoothed 3D fields obtained from the meteorological model. Fields that we switch are both horizontal wind components, liquid water potential temperature, total specific humidity and specific contents snow, liquid water and graupel.

With this setup, MicroHH is then run effectively from the meteorological model's state. We start the run at least an hour before case times, to allow the MicroHH model some relaxation from the meteorological state to its own dynamics. In Table 2, some data about the model setups are shown.

MicroHH was run on a Nvidia L40S GPU, as unfortunately, version 2 with included immersed boundary does not yet work with MPI parallelization. A 5 hour forecast was on average calculated in about 12 hours.

## 6 Results

### 6.1 Barje case

Here, we focus on values of virtual temperature and relative humidity at the site of Sodar-RASS measurements, to see how the models reproduce inversions and fog. We look at results at 08:00 UTC (10:00 local time), 5 hours after the MicroHH model start.

In the virtual temperature (Fig. 4 - left), we see that the subsidence inversion layer at around 600 m is well

simulated by both models, but below the meteorological model mixes to the ground, while in MicroHH, the ground inversion layer is still present. We can also see that the bottom boundary condition for MicroHH, which is static during one model run is not suitable anymore after 5 hours of forecast. This is likely also the reason for the systematic difference between Sodar-RASS measurements and MicroHH results.

In relative humidity (Fig. 4 - left), we notice the overexaggerated values in MicroHH in the lowest layer, which is again a consequence of the wrong bottom boundary condition, but nevertheless we see a small inversion layer at height 350 and 80% relative humidity. This indicates that here in MicroHH we get fog which is slowly dissolving. above we have a humidity profile typical for adiabatic mixing (almost constant between 400 and 700 m). In the meteorological model, there is no fog present and the adiabatic mixing is weakly expressed.




Figure 4: Results for Barje case at 08:00 UTC at the Sodar-RASS measurement site. Left - profile of virtual temperature, right - profile of relative humidity.

## 6.2 Planica case

In this case, we focus on values of the horizontal wind components and virtual temperature at the site of Planica ski jumping hill, to see how the models reproduce the wind reversal and possible retention of cold air in the valley. We look at results at 10:00 UTC, 5 hours after the MicroHH model start.

In wind profiles (Fig. 5), we see that both models capture the general strong northeastern wind above 1000 m, but in MicroHH the wind reversal in the lower layers is also present, where in the meteorological model, the northeastern wind direction is present to the valley floor, which does not correspond to the real situation. This effect is even more visible if we look at the meridional wind component cross-section in the west-east direction (Fig. 7), where the effect of the buffer layer (from 5500 to 7000 m) is also visible.

In the virtual temperature profiles (Fig. 6), we see that above 1300 m there is an adiabatic temperature profile in both models, which is a consequence of strong mixing due to strong wind in these layers. In the meteorological model the mixing is present down to the valley floor, where in MicroHH model the mixing in the valleys is smaller, due to retention of cold air, which is more realistic.

## 7 Conclusion and future work

Above results present our first try of using a high resolution LES, where the model is not run on an idealized experiment, but instead initialized from fields of a meteorological model. We analyzed results on two cases with relatively well known meteorological situations and compared them to the results of a meteorological model.

Despite the simplicity of our implementation typical of first tries, the results are qualitatively quite good. The model captures effects such as orography induced wind rotor in a very narrow valley, fog creation and ground




Figure 5: Results for Planica case at 10:00 UTC at the Planica ski jumping hill site. Left - profile of the zonal wind component, right - profile of the meridional wind component.

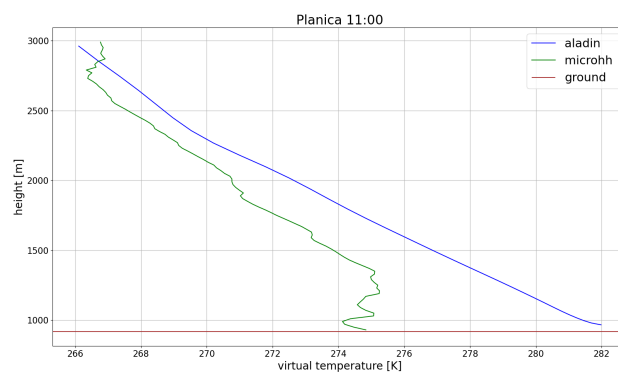


Figure 6: Results for Planica case at 10:00 UTC at the Planica ski jumping hill site for virtual temperature.




Figure 7: MicroHH meridional wind component cross-section in the west-east direction at 10:00 UTC. Red circle denotes the site of Planica ski jumping hill

inversion.

Biggest problems for joining meteorological fields with a LES model occur mainly on boundary conditions - the bottom conditions are kept constant for the whole time of integration, and because the lateral conditions are cyclic, we are forced to implement artificial smoothing and orography to prevent sharp gradients on domain edges.

The model also ran for 5 hours of forecast without any external forcing, which means that during the run, MicroHH does not receive information about general changes in meteorological situations (for example, a change in the general wind direction after 3 hours, or a change of shortwave flux due to night-day transition). This issues are being addressed.

Scripts for preparing initial and boundary conditions for MicroHH from the meteorological model are written in Python and already enable us to routinely run different cases on different domains, but some work will also be put in to write the documentation and make the scripts more user friendly, to enable users to easily run the MicroHH models with their own meteorological data.

For comments, critiques and questions, please write to [peter.smerkol@gov.si](mailto:peter.smerkol@gov.si).

## References

- [1] Termonia, P., Fischer, C., Bazile, E., Bouysse, F., Brožková, R., Bénard, P., Bochenek, B., Degrauwe, D., Derková, M., El Khatib, R., Hamdi, R., Mašek, J., Pottier, P., Pristov, N., Seity, Y., Smolíková, P., Španiel, O., Tudor, M., Wang, Y., Wittmann, C., and Joly, A.: The ALADIN System and its canonical model configurations AROME CY41T1 and ALARO CY40T1, *Geosci. Model Dev.*, 11, 257–281, <https://doi.org/10.5194/gmd-11-257-2018>, 2018.
- [2] (BOMEX) : Holland, J. Z., and E. M. Rasmusson, 1973: Measurements of the Atmospheric Mass, Energy, and Momentum Budgets Over a 500-Kilometer Square of Tropical Ocean. *Mon. Wea. Rev.*, 101, 44-55
- [3] (GABLS): Holtslag, B. J. W. (2006). Preface: GEWEX Atmospheric Boundary-layer Study (GABLS) on Stable Boundary Layers. *Boundary-Layer Meteorology*, 118(2), 243–246.
- [4] Stevens, B., and Coauthors, 2003: Dynamics and Chemistry of Marine Stratocumulus—DYCOMS-II. *Bull. Amer. Meteor. Soc.*, 84, 579–594, <https://doi.org/10.1175/BAMS-84-5-579>.
- [5] Kosmač, U. 2025: Turbulence in simulaciji velikih turbulentnih vrtincev, Masters degree, University of Mathematics and Physics in Ljubljana.
- [6] van Heerwaarden, C. C., van Stratum, B. J. H., Heus, T., Gibbs, J. A., Fedorovich, E., and Melado, J. P.: MicroHH 1.0: a computational fluid dynamics code for direct numerical simulation and large-eddy simulation of atmospheric boundary layer flows, *Geosci. Model Dev.*, 10, 3145–3165, <https://doi.org/10.5194/gmd-10-3145-2017>, 2017.
- [7] Lilly, D. K.: Models of cloud-topped mixed layers under a strong inversion, *Q. J. Roy. Meteor. Soc.*, 94, 292–309, 1968.
- [8] A. S. Monin in A. M. Obukhov, Basic laws of turbulent mixing in the surface layer of the atmosphere, *Trudy Akademii Nauk SSSR, Geofizicheski Institut* 24(151) (1954) 163–187.
- [9] H. Tomita, New microphysical schemes with five and six categories by diagnostic generation of cloud ice, *Journal of the Meteorological Society of Japan. Ser. II.* 86A (2008) 121–142.

- [10] S. K. Krueger, H. Morrison in A. M. Fridlind, *Cloud-resolving modeling: arm and the gcss story, v: The Atmospheric Radiation Measurement Program: The First 20 Years* (ur. D. D. Turner in R. G. Ellingson), *AMS Meteorological Monograph 57*, American Meteorological Society, 2016, str. 25.1–25.16, doi: 10.1175/AMSMONOGRAPHS-D-15-0047.1.

# The implementation of the sweep interpolation in the semi-Lagrangian scheme of the ACCORD system

Natalia Szopa, Petra Smolíková

## 1 Introduction

---

The sweep interpolation method as described and demonstrated by Mortezaadeh and Wang is motivated by its ability to reduce both dissipation and dispersion errors, which is particularly critical near sharp gradients. Even though it is a 4th-order semi-Lagrangian scheme it achieves a high-order accuracy at the computational cost of a 3rd-order method by alternating backward and forward interpolations. Theoretical explanation of the method can be found in [1] and [2].

Understanding how the sweep interpolation works and implementing it in the LAM code, as well testing the quality of the new method was the focus of my stay at CHMI, Prague in November 2025.

## 2 Implementation

---

The code modification was done based on the operational version of CY48t3mas available at CHMI in Prague on lada: /home2/mma010/pack/cy48t3mas\_op1.02.NEC400DPMPI226.x and may be found on lada: /local/mma277/pack/CY48t3mas\_sweep.01.NEC400DPMPI226.x. We implemented the sweep interpolation method in the LAM code by phasing Filip Váňa's original implementation done in the global IFS to cy48t3mas available in Prague and by adapting it to the local model geometry.

Since interpolation weights in the limited area model use plane geometry while the global model employs spherical coordinates, we had to implement corresponding modifications to the LAM code. We introduced a new logical key *LSWEEP* as a global switch activating sweep weights computation. We added new conditional cases in the routine *elascaw.F90* that compute quadratic horizontal and vertical weights required for sweep interpolation when *LSWEEP* = *.T.*. The following routines have been modified:

aladin/interpol/elascaw.F90  
arpifs/adiab/larcina.F90  
arpifs/adiab/larcinha.F90  
arpifs/oops/interpolator\_ad\_mod.F90  
arpifs/oops/interpolator\_mod.F90  
arpifs/op\_obs/slntad.F90  
arpifs/op\_obs/slnt.F90

## 3 Experiment

---

To evaluate the results we conducted a test on the classical Robert's rising thermal bubble. We define it to be a sharp edged 2-dimensional bubble with perturbed temperature: +1K in the centre of bubble and wind speed of 0m/s (no advection) in the field of constant potential temperature of 300K. The experiment was carried out on a domain of 256 gridpoints, 200 vertical levels with meridional grid spacing and vertical spacing of half levels both set to 10m.

We compared the solutions of the sweep interpolation with reference solutions using quadratic and cubic interpolators on Figures 1 - 4.

## 4 Conclusion

The sweep interpolation method produces results visually very similar to both quadratic and cubic reference solutions, making it difficult to draw definitive conclusions about its relative performance. However, the close agreement with both established interpolators suggests the method successfully captures the essential dynamics of the rising thermal bubble. Thus the method is promising and additional testing across varying parameters and a 3-dimensional test cases is necessary to further validate the sweep interpolation method for atmospheric modelling applications.




Figure 1: Bubble test: difference between sweep and quadratic Lagrange interpolation of perturbation of potential temperature after 7 minutes.

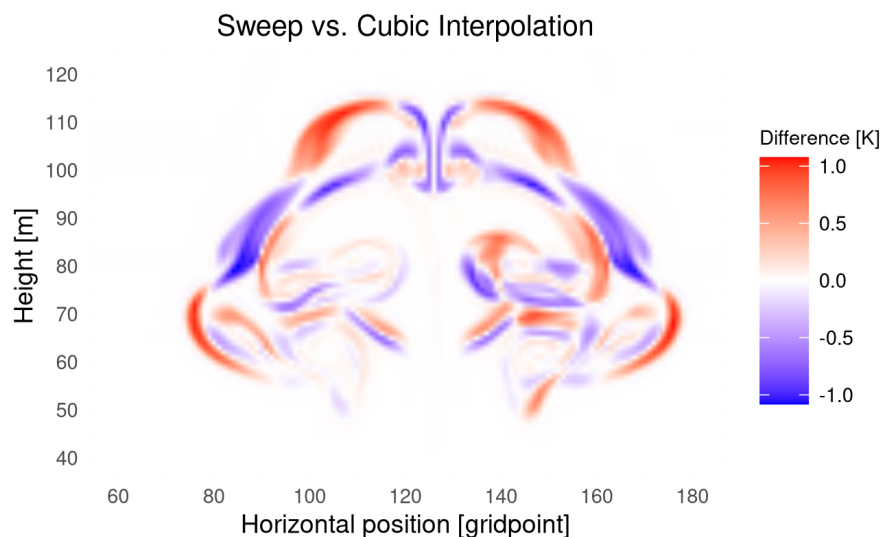


Figure 2: Bubble test: difference between sweep and cubic Lagrange interpolation of perturbation of potential temperature after 7 minutes.

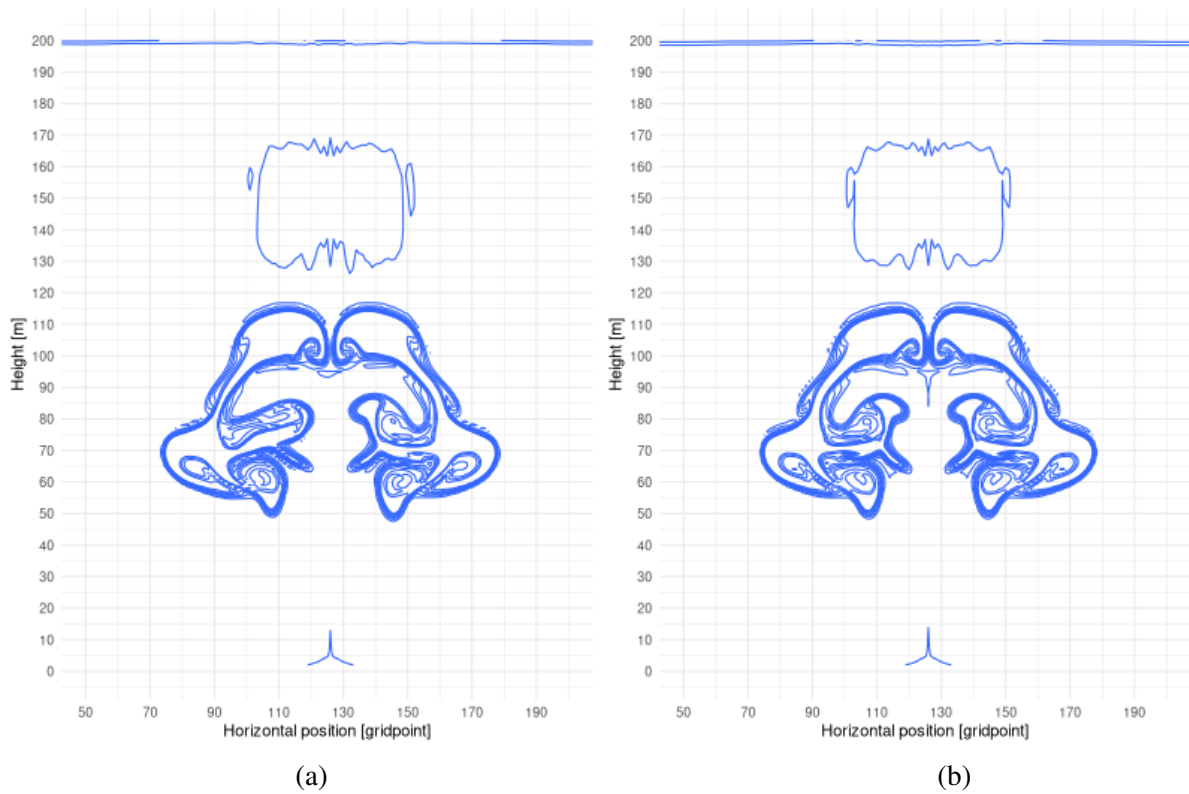


Figure 3: Bubble test: perturbation of potential temperature after 7 minutes - (a) sweep interpolation, (b) quadratic interpolation.

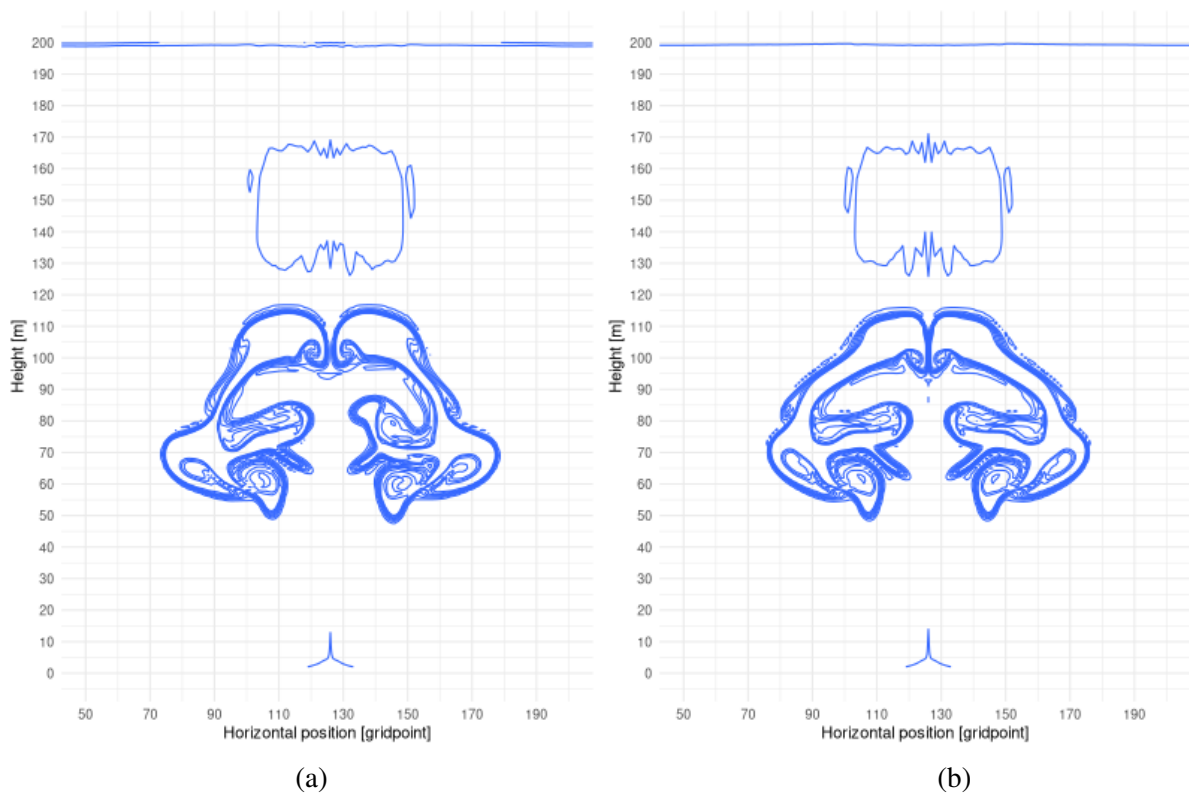


Figure 4: Bubble test: perturbation of potential temperature after 7 minutes - (a) sweep interpolation, (b) cubic interpolation.

## References

- [1] Mortezaazadeh, M., & Wang, L. L., 2017: *A high-order backward forward sweep interpolating algorithm for semi-Lagrangian method*, International Journal for Numerical Methods in Fluids, 84(10), 584–597.
- [2] Mortezaazadeh, M., Cossette, J.-F., Dastoor, A., de Grandpré, J., Ivanova, I., and Qaddouri, A., 2024: *Sweep interpolation: a cost-effective semi-Lagrangian scheme in the Global Environmental Multiscale model*, Geosci. Model Dev., 17, 335–346.

# Visualization of norms

Harold Petithomme, Météo France CNRM

## 1 Introduction

---

In the model, norms represent the average of a field, may it be 2D or 3D.

They can be useful for scientific or technical activities:

- assert that values are realistic: no NaN or Infinity, no negative values where relevant (like contents or concentration, for example)
- assert that values are unchanged after some technical (or scientific) modification
- learn about basic range, mean or statistics of fields: diurnal or seasonal cycle, trends,...

Actually, playing with norms is a very good way for learning about the model. When one knows how to activate their printing in NODE files, then it is possible to use them in a smart way.

This page presents 2 small tools that can parse norms printed in NODE files and produce nice output like light web pages with graphics.

## 2 Layout

---

This article first presents what kind of norms can be printed in NODE files and how to control the printing. Two tools are then presented: one for plotting the different kinds of norms and the other one for comparing norms from 2 NODE files. Finally, some further details and perspective about this topic are mentioned.

### Norms in NODE files

In the code, there are 2 great kinds of norms that can be printed:

- norms for spectral fields (or SP norms)
- norms for grid-point fields (or GP norms)

A field is either a 2D field (surface pressure, 10-metre wind speed) or a 3D field (temperature, U and V wind components, liquid water content,...)

Printing norms is controlled by a few keys in namelist NAMCT0. Among them:

- NSPPR: an integer value controlling whether to print, for 3D fields only, only an average of levels (0), each level (1) or each wave of spectral fields (2)
- NFRSDI/NFRGDI: an integer value giving the frequency (in steps) of norms printing for spectral / grid-point fields respectively
- NSDITS/NGDITS: in case of NFRSDI=0 (resp. NFRGDI=0), values of time-steps (hours if negative) for which to print spectral norms (resp. grid-point norms)

These keys control norms printing for model variables. When FullPOS is in use, norms for post-processed fields (FP norms) are also printed. However, control parameters are different for these norms: NFRPOS and NPOSTS (same characteristics as previous keys).

### Specific norms

Down to cycle 49, supplementary output of norms are available. These are called specific norms in the following.

In cycle 49, a few GP norms have been added. New prints in NODE files are available for GMV and GFL variables as well as variables from SL buffer 2. For GMV and GFL variables, there are prints for model states t0 and t1, that is current and future states. There are even several printings for state t1. The same configuration is applied to SL buffer 2. In the end, all these new output of GP norms are placed in the GP part of the model. In short:

- in gp\_model.F90: before entering cpg\_drv, output of GP norms of GMV/GFL at t0
- in cpg\_drv.F90: after call to cpg, output of GP norms of GMV/GFL at t1 and SL buffer 2
- in call\_sl\_drv.F90:
  - after call to call\_sl, output of GP norms of GMV/GFL at t1 and SL buffer 2
  - after call to trmfixers, output of GP norms of GMV/GFL at t1
- in cpplag\_drv.F90: after call to cpplag, output of GP norms of GMV at t1

In cycle 51 to come, some other norms have been added :

- in stepo.F90: after calls to transdirh/etransdirh, output of SP norms

These new output of norms are conditional. They are activated by new keys from NAMCT1, LGPDIAG (for GP norms) and LSPDIAG (for SP norms). They also follow their respective frequency of printing: NFRGDI and NFRSDI (but not NGDITS, nor NSDITS).

### Plotting norms

The tool norms.sh provides a very nice way for plotting norms printed in any NODE file.

This shell script reads in any input NODE file and produces a small set of graphics (PNG files), which are available locally. For a smart display, these PNG files are linked in a simple HTML file, letting the user load this HTML page in any web browser.

The script norms.sh has several options to be passed on the prompt, offering a fairly wide range of use cases. Here is its syntax:

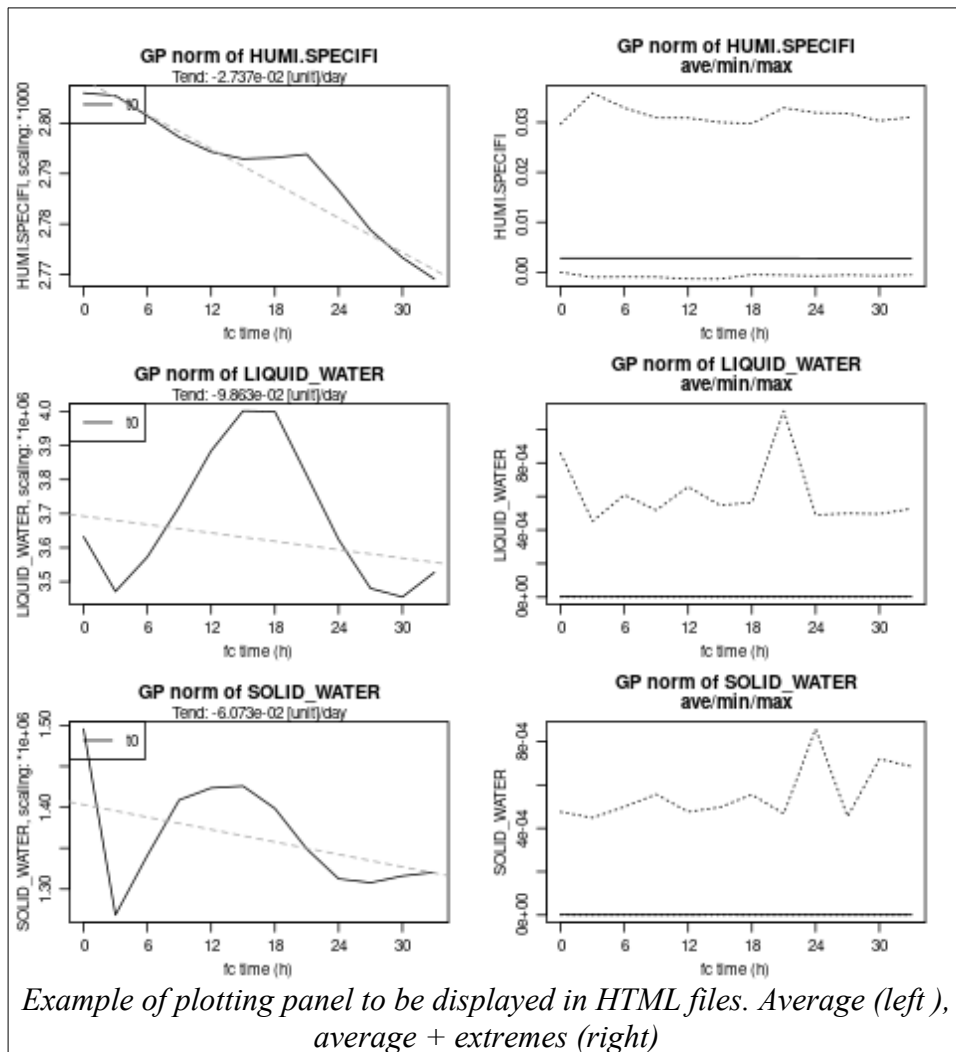
```
norms.sh NODE1 NODE2... -o HTML [-lev (LEV|I1:I2:...)] [-detail] [-nogfl] [-nogmv] [-noadiab] [-noslb2] [-nospt1] [-nogpt1] [-not1] [-opt OPTIONS] [-h]
```

The description of the options and details about the usage is given by the doc/help on the command line (option '-h'). One can notice the 3 main use cases of the script:

- simple mode: 1 NODE file as input, 1 level. Plots show norms of variables along time-steps. On each plot, curves are for the same variable, but potentially for different "aspects" (time-level, mainly).
- advanced mode: several NODE files as input. Plots still contain norms along time-steps but now with one curve per input file (ie comparing experiments)
- vertical mode: in simple or advanced mode, plots are now vertical profiles of variables at 3 fixed time-steps instead of one level along the time-steps

Remarks on levels :

1. the vertical mode is managed by option '-lev' where a value of -1 should be passed. Only 3D fields can be plotted.
2. in the "horizontal" mode, passing 0 to option '-lev' means plotting norms for the average of 3D fields. Norms for 2D fields can only be plotted for this value.



Remark for PC scheme:

One other point to mention is about PC scheme ("Predictor / Corrector"), which is part of the dynamics. For each P/C iteration, SP and GP norms are output when relevant. In this situation, the tool handles P/C this way:

- at predictor iteration, time of integration is considered as the time of the current time-step times the time-step duration
- at corrector iteration, time of integration is considered as the same time as for predictor iteration + half the time-step duration (more precisely '+TSTEP/(NSITER+1)', cf code for details)

This way, curves on plots are kept continuous with or without PC scheme. This may be seen as strange but is a fairly simple interpretation of partial completion (at predictor) / full completion (at corrector) of steps.

### Comparing norms

The 2nd tool is another shell script : normdiff.sh. Its aim is not to plot norms from NODE files but rather to precisely compare these norms.

This simple script is well designed for technical work. It lets the user know the difference between norms from 2 NODE files, for instance given from 2 versions of code (different cycles or branches or simply after some local modification of code). The difference in norms is simply printed on the screen and is expressed as the number of differing digits. Hence, an output of 0s only indicates a strict

equality of the compared norms. On the contrary, when things go wrong, differences will show up 17s, ie the maximum number of digits in a norm value.

The script is simple to use, with a few set of options available. Here is the syntax:

```
normdiff.sh NODE1 NODE2 [-spre RE|-nosp] [-gpre RE|-nogp] [-nosurf] [-nofp] [-mnx]
[-ftz] [-h]
```

By default, the norms compared and output are:

- those for spectral fields (SPEC)
- those for grid-point fields (GFL)
- those for post-processed fields, spectral or grid-point (FP)
- those for surface fields (SURF), initial or final norms (eg in CANARI)

Disabling any kind of norms printing is possible: just use one or several intuitive options among '-nosp', '-nogp', '-nofp' or '-nosurf'.

step	VORT	DIVE	TEMP	HUMI	KINE	SP
0	0	0	0	0	0	0
1	15	16	9	14	15	0
2	15	16	10	15	15	10
3	15	16	10	15	15	10
4	16	16	11	15	15	10
5	16	16	11	15	15	10
6	16	16	11	15	15	11

*Example of output for SP fields (average only)*

Let's now present the main use cases:

- standard comparison: comparison of standard norms for SPEC and GFL fields
- specific comparison: for SP norms (resp. GP norms), using option '-gpre RE' (resp. '-spre RE') let's the user pass a regular expression (RE) which is the header of specific norms to be found in the NODE files

Two other features are worth to mention.

First, for GP norms, it is possible to output comparison for min and max values in addition to the average value. This is the goal of option '-mnx' ("Mean, miN and maX"). Output is kept in a fairly compact form:

step	PROFTEMPERAT	PROFRESERV.E	PROFRESERV.G	SURFRESERV.N	SURFALBEDONE	SURFDENSIT.N	SURFALBEDOHI	SURFTEMPERAT
pp1	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
pp2	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
pp3	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
pp4	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
pp5	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
... 3 more 0 lines								
step	SURFRESERV.E	SURFRESERV.I	SURFRESERV.G	SURFZ0.FOIS.	SURFALBEDO	SURFEMISSIVI	SURFET.GEOP0	SURFIND.TERR
pp1	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
pp2	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0

*Example of output for FullPOS fields with min/max ('-mnx')*

The other point is a bit more tricky. Indeed, some norms printed in NODE files may not be fully consistent. Actually, for grid-point norms, the average can be non significant compared to min and max values. This is the case for example for wind components (available in GMV structure, so in specific GP norms) : when the average is more or less 10<sup>15</sup> times smaller than the extremes (min and max), it may not be significant. This is why option '-ftz' ("flush to zero") has been introduced. When used, the average is rounded up to a significant value (cf code for details).

Finally, the tool normdiff.sh has been made the most generic possible. It can indeed handle a wide variety of situations:

- in integration jobs: standard norms as well as specific norms
- in PC scheme: norms for predictor as well as for corrector
- in TL/AD testing jobs : norms for TL, for AD and for trajectory storing
- if post-processing is on : norms for FP events

Last remark : if one NODE file is uncompleted or contains fewer norm events, the tool handles it fine and limits comparison to events in common.

### **Perspective for norms**

Tools about norms from NODE files can be very useful for both technical and scientific work. Now these tools are available, it is interesting to think of adding ever more norms printing in the code. One can imagine adding some output for variables and arrays of great interest, in any scientific domain : dynamics (SI, SL or HD arrays), physics (internal arrays of parametrization, tendencies), DA, observation, ensembles, etc.

Once norms are available, it would be easy to adapt the tools to these new norms. They are indeed very generic and easy to extend.

In addition, a new similar tool is under development for plotting internal variables without printing norms. It is rather based on printing local statistics (ie mean, min, max) of each MPI task (in a logging manner). This would make it possible to plot and compare a proxy-distribution of variables, which is a really cheap way to enter deep down the model.

## **3 References**

---

Technically, the 2 tools presented here are made of 2 shell scripts that call 1 or 2 other scripts (in R language) for doing all the treatment (parsing, plotting or printing).

These other R scripts also use an R package, `mfnode`.

The code of the tools is available on Github : <https://github.com/HaroldMTO/node>

The R package `mfnode` must be previously installed. Get it from Github as well: <https://github.com/HaroldMTO/mfnode>

Installation procedure is the classical one for R packages. Look at README file in node.

For more information and support, please contact the author: [harold.petithomme@meteo.fr](mailto:harold.petithomme@meteo.fr)

# Patricia around ACCORD 26 beautiful countries (4)

Patricia Pottier

## 1 An autumn escape to Croatia and Slovenia: Water in all its states

Late October 2025, while countless beautiful places remained to be discovered, the time had come to head west and return home, taking with us so many wonderful memories.

Throughout this journey, water was our constant companion, from the salty spray of the Adriatic to the emerald lakes and thundering waterfalls featured here. But as we know, water is also a formidable force; it manifested itself most spectacularly through torrential rains and flooding during two extreme weather events. These were very accurately predicted by ARSO and DHMZ, allowing us to safely adjust our itinerary between the interior and the coast.

Beyond the scenery, this escape offered the joy of reconnecting with LACE friends at the LCS dinner in Ljubljana.



Explore More: <http://pottier.pat.meteo.free.fr>



- [Book 4](#): published in December, featuring an extensive collection of photographs.



- [the Padlet](#): our digital travelogue, updated throughout the journey.

## 2 Club of early ALADINists

For those who shared the first lines of code and the earliest workshops in the 90s, the ALADIN adventure was never just about numerical weather prediction. Today, this group of "early pioneers" continues to keep that heritage alive through an informal club.

During our first two meetings in 2025, we shared updates on our current lives and projects. To stay connected, we have launched a shared GoogleDoc where members can post a short "bio" to help us get to know each other better without repeating introductions at every meeting. We are also collaborating on a shared GoogleSlide document to collect and enjoy photos from our decades of cooperation. Beyond reminiscing over shared photos, we decided to create a dedicated digital space to archive ALADIN documents, ensuring our human history is preserved. Work in progress, all material welcome !

Save the Date: our next meeting is scheduled for Wednesday, 25 March 2026, at 18:00 CEST. See <http://pottier.pat.meteo.free.fr/aladin.html> .




# Previous editions of the ACCORD Newsletter

Anne-Lise Dhomps, ACCORD Consortium Scientific Secretary



[1st ACCORD Newsletter](#),  
published on 5 October 2021



[2nd ACCORD Newsletter](#),  
published on 28 February 2022



[3rd ACCORD Newsletter](#),  
published on 24 October 2022



[4th ACCORD Newsletter](#),  
published on 19 June 2023



[5th ACCORD Newsletter](#),  
published on 27 March 2024



[6th ACCORD Newsletter](#),  
published on 23 December 2024



[7th ACCORD Newsletter](#),  
published on 29 July 2025

The joint ALADIN-HIRLAM Newsletter (2013-2021, 16 editions) can be found on the [ALADIN dedicated webpage](#).

### Elaboration of the ACCORD Newsletters

For the concrete writing of an article, please refer to the editorial guidelines, accessible at: <http://www.accord-nwp.org/?Recommendations-templates>

The newsletter content is based on voluntary contributions by the scientists and the teams in the consortium. We want it to be a useful tool for sharing both “practical” information and experience (code engineering, quality assurance, system aspects) and “more fundamental” results (advances in research work, outcome of specific meetings or working days etc.).

NL9 is expected for end-2026, however contributors can upload their material at any time during the year [on this shared folder](#). Do also not hesitate to encourage scientific contributions by young scientists (PhDs, post-docs etc.) or technical contributions (codes, porting, optimization etc.).