CopERnIcus climate change Service Evolution

D2.2 Documentation of coupled assimilation infrastructure and methodology and preliminary assessment towards optimal degrees of coupling for coupled regional reanalysis

Due date of deliverable	31/12/2024
Submission date	Dec 2024
File Name	CERISE-D2.2-V1.0
Work Package /Task	WP2 Tasks 2.1 and 2.2
Organisation Responsible of Deliverable	MET-Norway
Author name(s)	Jostein Blyverket, Åsmund Bakketun, Benjamin Menetrier, Jelena Bojarova, Magnus Lindskog and Faúndez Alarcón José Manuel
Revision number	1.0
Status	Issued

Funded by the European Union

The CERISE project (grant agreement No 101082139) is funded by the European Union.

1 Executive Summary

The purpose of this work is to develop the next generation of C3S (Copernicus Climate Change Service) regional scale reanalysis codes infrastructure to support modular coupled assimilation and monitoring. In the HARMONIE-AROME system, which is used to produce the C3S regional reanalysis, we developed and tested innovative coupled surface-atmosphere assimilation. This document presents the infrastructure and coupled data assimilation methodology that was developed in the HARMONIE-AROME and it presents preliminary results.

Table of Contents

Contents

2 Introduction

2.1 Background

The scope of CERISE is to enhance the quality of the C3S reanalysis and seasonal forecast portfolio, with a focus on land-atmosphere coupling.

According to the project plan, it will support the evolution of C3S, over the project's 4 year timescale and beyond, by improving the C3S climate reanalysis and the seasonal prediction systems and products towards enhanced integrity and coherence of the C3S Earth system Essential Climate Variables.

CERISE will develop new and innovative ensemble-based coupled land-atmosphere data assimilation approaches and land-surface initialisation techniques to pave the way for the next generations of the C3S reanalysis and seasonal prediction systems.

These developments will be combined with innovative work on observation operator developments integrating Artificial Intelligence (AI) to ensure optimal data fusion fully integrated in coupled assimilation systems. They will drastically enhance the exploitation of past, current, and future Earth system observations over land-surfaces, including from the Copernicus Sentinels and from the European Space Agency (ESA) Earth Explorer missions, moving towards an all-sky and all-surface approach. For example, land observations can simultaneously improve the representation and prediction of land and atmosphere and provide additional benefits through the coupling feedback mechanisms. Using an ensemble-based approach will improve uncertainty estimates over land and lowest atmospheric levels.

By improving coupled land-atmosphere assimilation methods, land-surface evolution, and satellite data exploitation, Research and Innovation inputs from CERISE will improve the representation of long-term trends and regional extremes in the C3S reanalysis and seasonal prediction systems.

In addition, CERISE will provide the proof of concept to demonstrate the feasibility of the integration of the developed approaches in the core C3S (operational Service), with the delivery of reanalysis prototype datasets (demonstrated in pre-operational environment), and seasonal prediction demonstrator datasets (demonstrated in relevant environment).

CERISE will improve the quality and consistency of the C3S reanalysis systems and of the components of the seasonal prediction multi-system, directly addressing the evolving user needs for improved and more consistent C3S Earth system products.

2.2 Scope of this deliverable

2.2.1 Objectives of this deliverables

This deliverable describes the HARMONIE-AROME coupled assimilation infrastructure and methodology, and it presents preliminary assessment towards optimal degrees of coupling for coupled regional reanalysis.

2.2.2 Work performed in this deliverable

In this deliverable the work outlined in WP2 T2.1: *Infrastructure developments to support coupled DA developments, and code efficiency and modularity,* and in WP2 T2.2 *Development of outer loop land-atmosphere coupling* is described and evaluated.

The current land-atmosphere coupling approach used for the C3S regional reanalysis, consists of a weakly coupled DA system, where results from atmospheric and land surface analyses are fed back into the next 3-hour data assimilation window by the short connecting forecast. In the context of WP2 Task 2.1, we conducted infrastructure developments in the HARMONIE-AROME to improve the system efficiency and to prepare for coupling data assimilation workflow developments. More specifically we introduced the land data assimilation system from WP1 into the coupled system. In Sect. 3 we outline these infrastructure developments. As part of WP2 Task 2.2, we developed a weakly coupled land-atmosphere DA system based on re-running the land surface model in the assimilation window with analysed forcing, this methodology is described in Sect. 4.2.1. These developments give us the opportunity to correct the land surface state throughout the 3-hour assimilation window and utilize more observations and observation types than in the default setup. We also investigated how we use interface observations by assimilating screen level variables in the upper-air 3D-VAR (3-dimensional variational data assimilation). This is described in Sect. 4.2.2. In Sect. 5 we report on the outer loop coupling developments done in the HARMONIE-AROME 4D-VAR (4-dimensional variational data assimilation).

2.2.3 Deviations and counter measures

Due to recruitment delays at SMHI, we extended T2.2 in 2025 to develop the coupling between land DA and 4D-VAR. In 2024, we developed and implemented a simpler coupling in T2.2 based on atmospheric 3D-VAR instead of 4D-VAR. Met Norway developed the coupling between land and the 3D-VAR atmosphere in 2024, while SMHI will develop the coupling between land and 4D-VAR atmosphere in the extended T2.2 in 2025.

This deviation was discussed and agreed with the PO. It does not impact deliverables and other WPs.

D2.2 (M24) describes the coupling methodology developed in 2024 and D2.5 (M48) will describe coupling developed in 2025.

CARRA3-Pv1 (M36) will use the coupling developed in T2.2 in 2024. CERRA2-Pv2 (M45) will integrate the developments further conducted in WP2 from 2025 onwards as planned.

In terms of resources, 6PM from SMHI is being reallocated to Met Norway for T2.2 in 2024.

2.2.4 Reference Documents

[1] Project 101082139- CERISE-HORIZON-CL4-2021-SPACE-01 Grant Agreement

2.2.5 CERISE Project Partners:

ECMWF	European Centre for Medium-Range Weather Forecasts
Met Norway	Norwegian Meteorological Institute
SMHI	Swedish Meteorological and Hydrological Institute
MF	Météo-France
DWD	Deutscher Wetterdienst
CMCC	Euro-Mediterranean Center on Climate Change
BSC	Barcelona Supercomputing Centre
DMI	Danish Meteorological Institute
Estellus	Estellus
IPMA	Portuguese Institute for Sea and Atmosphere
NILU	Norwegian Institute for Air Research
MetO	Met Office

3 Infrastructure developments for code efficiency and modularity

3.1 Land data assimilation in the coupled system

Infrastructure developments were done to run the ensemble square root (EnSRKF) and local ensemble transform Kalman filter (LETKF) in HARMONIE-AROME. More details on these developments are reported in D1.3.

3.2 OOPS

There is a fork of the OOPS C++ layer developed at the JCSDA (Boulder, Colorado) in the framework of the JEDI project, for which a model-agnostic background error covariance library called SABER has been developed. To enable the use of this SABER library, some classes and interfaces of the ECMWF OOPS have been modified. SABER includes a 2D correlation operator called “shadow-levels” that can take orography into account precisely. This correlation operator has been used to run a 2D-VAR for screen variables (2m temperature and humidity), that could replace the existing OI. It could also be used as a localization operator for a 2D-EnVAR method in future experiments.

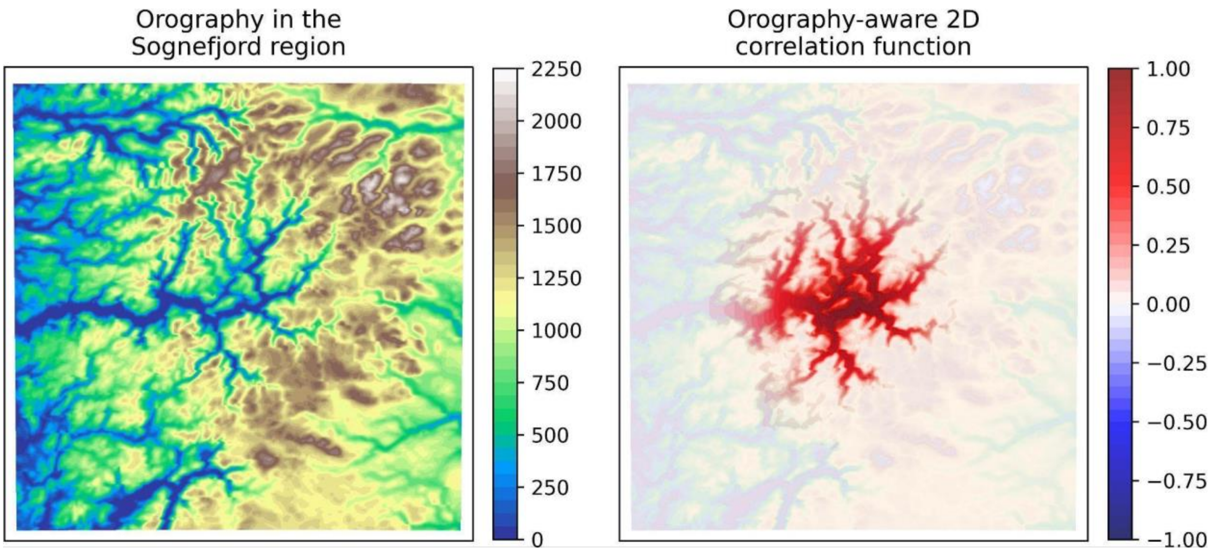

Figure 1: (Left) orography in the Sognefjord region in southern Norway, (right) orography-aware 2D correlation function.

4 Weakly coupled land-atmosphere data assimilation in 3D-VAR

4.1 Overview

Two approaches have been explored for the weakly coupled land-atmosphere data assimilation system in HARMONIE-AROME. First, we have implemented a methodology called analysed forcing, where we re-run the land surface model within the assimilation window using analysed forcing. Second, we have included the assimilation of 2m temperature and humidity observations in the upper-air 3D-VAR. This use of interface observations is an intermediate step towards more coupled data assimilation.

4.2 Methodology

In the current CARRA2 re-analysis we use the simplified extended Kalman filter (SEKF) for updating the soil temperature and moisture in levels 1-3 and 2-5, respectively based on SYNOP temperature and humidity observations. The SEKF is 1D which means that we first do a horizontal Optimal Interpolation (OI) to spread the observation information in space. The snow analysis in CARRA2 is based on horizontal OI of in-situ snow depth observations and satellite derived snow or no-snow observations from https://cryo.met.no/en/regional-snowcover-from-satellite. The upper-air DA system is based on the 3D-VAR and assimilates conventional and satellite observations.

4.2.1 Analysed forcing developments

The analysed forcing methodology was first described in (Bakketun et al., 2023). The motivation behind this methodology is that instead of sequentially updating the land surface at every 3-hour with observations valid only at that time, we hourly update the land surface throughout the 3-hour assimilation window using observations available throughout the window to analyse the forcing.
More specific:

- Run the coupled forecast model through the assimilation window.
- Use the coupled model run as first guess for analyzing the forcing using screen level and/or precipitation observations.
- Re-run the land surface model in an “outer-loop” sense for the assimilation window using the analysed forcing.
- Use new surface state as first-guess for surface assimilation of satellite observations and/or as initial conditions for the subsequent forecast cycle.

The methodology is illustrated in Fig. 2 from (Bakketun et al., 2023). In that paper the authors used the Nordic-Analysis product as input (https://github.com/metno/NWPdocs/wiki/MET-Nordic-dataset), while we in this work do our own gridded analysis.

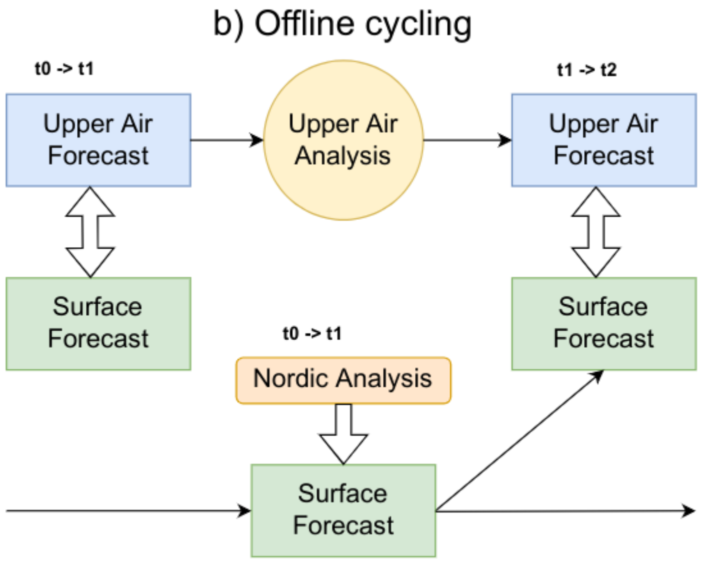

Figure 2: Schematics of the analysed forcing methodology. First we run the coupled forecast model through the assimilation window (blue and green box) from t0 to t1. The land surface model is then re-run in an "outer-loop" sense with hourly analysis of 2m temperature, humidity and precipitation. This is used as an improved first guess to the coupled run from t1 to t2.

Analysis of 2m temperature and relative humidity

Hourly 2m temperature and relative humidity observations undergo quality control (QC) using the titanlib package (https://github.com/metno/titanlib). After the QC of the observations we do horizontal optimal interpolation (OI) using gridPP (https://github.com/metno/gridpp) and hourly first guess values of 2m temperature and relative humidity, see Fig 3 for observation usage. After the OI analysis, we create surface forcing based on these analysed variables and the model first guess for non-analysed variables. We then re-run the land surface model throughout the 3-hourly assimilation window to produce a new land surface state for the subsequent forecast (and or additional surface analysis, e.g., assimilation of snow and/or satellite observations). Figure 3, shows the number of active T2m observations in the analysed forcing (orange) experiment vs the default CARRA2 setup. We note that for the analysed forcing we have observations for every hour, while for the default setup we are not utilizing observations between the 3-hourly analysis times. Another point to note is that the number of observations in the analysed forcing drops at e.g. 01, 04.. UTC etc. this is most likely due to the range in the observation files is not large enough for offset 3, i.e., t-3.

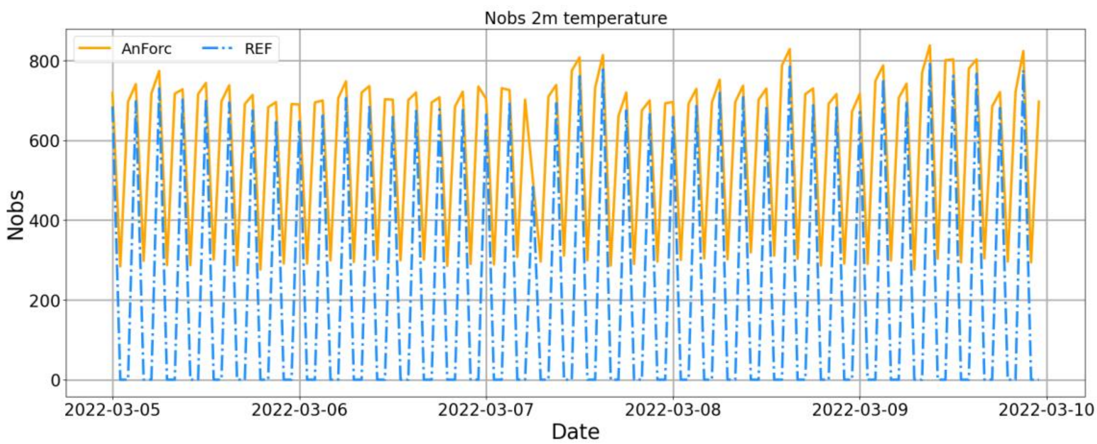

Figure 3: Number of active observations in the Analysed forcing (AnForc, orange) methodology vs reference run (REF, blue).

Precipitation analysis

The precipitation analysis was not technically ready to run at the time when we started the verification experiments, see Sect. 4.3. We will set up an additional experiment also including analysed precipitation when the technical readiness level is higher.

The precipitation analysis works in the same way as the 2m variables, however we had to add fetching of 1-hourly accumulated precipitation data in the Bufr files. As first guess values we use the 1-hour accumulated precipitation from the upper-air files for every hour. In the pysurfex framework (https://github.com/metno/pysurfex) we added i) 1h accumulated precipitation in Bufr2Json creating json files from the Bufr observations, ii) 1h accumulated precipitation from FA (internal HARMONIE-AROME format) to netCDF, iii) quality control of 1h accumulated precipitation and iv) horizontal OI. The analysed precipitation can then be used when we create the SURFEX offline forcing, in the same way as 2m temperature and humidity.

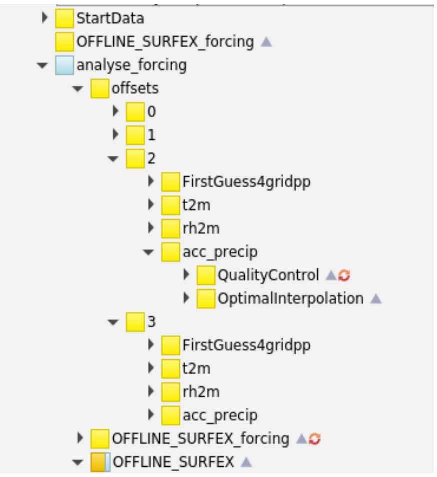

Figure 4: Ecflow setup of the analysed forcing methodology. For the different offsets we perform QC and OI on 2m temperature, humidity and 1h accumulated precipitation. The analysed variables enter the SURFEX offline forcing in the OFFLINE_SURFEX_forcing task and are used in the surface loop (OFFLINE_SURFEX).

4.2.2 Assimilation of 2m temperature and humidity in upper-air 3D-VAR

After moving to a more physically based surface scheme in CARRA2, our experience is that the surface analysis (when assimilating screen level variables) has less impact than in the default schemes. The sensitivity calculated through the SEKF indicates that updating the soil based on these observations has very little impact.

Recently, ECMWF introduced the assimilation of SYNOP temperature and humidity data into their 4D-VAR (Ingleby et al., 2024). This gave large improvements in the short-range T2m forecast. Here we explore this methodology of assimilating 2m temperature and humidity in the upper-air 3D-VAR in HARMONIE-AROME.

In the HARMONIE-AROME code we add 2m temperature and relative humidity observations to the upper-air 3D-VAR. In the quality control we have excluded 2m temperature and humidity observations taken from ships, we only use automatic and manual reports over land. We also filter out observations where the land-sea mask is less than 0.75. Most likely this value could be set to 1 to reduce representation errors in coastal regions. We also check that the difference in station vs model height is less than 100m. The 2m model first guess is lapsed to the observation height using a lapse rate of 6.5 K/km. The observation errors of 2m temperature and relative humidity are set to 0.1 and 1.4K, respectively.

4.3 Experiments

We run three sets of experiments, two evaluating the analysed forcing methodology for March to May and June to August 2022 (named CARRA3_Pv1_AnForc_Spring and Summer, respectively). The assimilation of 2m temperature and relative humidity (named CARRA3_Pv1_screenDA) is run for the beginning of June 2022 and compared to the default setup (named CARRA3_Pv1_REF_spring and Summer, respectively) and analysed forcing.

Table 1: List of experiments and experiment settings.

Experiment	Assimilation settings	Analysed forcing	Time-period
CARRA3_Pv1_REF_Spring	3D-VAR + SEKF and snow DA	No	1st March - 1st May
CARRA3_Pv1_AnForc_Spring	3D-VAR and snow DA	Yes	1st March - 1st May
CARRA3_Pv1_REF_Summer	3D-VAR + SEKF and snow DA	No	1st June - 1st August
CARRA3_Pv1_AnForc_Summer	3D-VAR and snow DA	Yes	1st June - 1st August
CARRA3_Pv1_screenDA	3D-VAR with 2m assim and snow DA	Yes	1st June - 20th June

4.4 Results

Here we present preliminary results for the analysed forcing methodology and assimilation of 2m temperature and humidity observations in the upper-air 3D-VAR.

4.4.1 Analysed forcing

In Fig. 5, we have plotted the normalized mean root-mean-squared-error temperature difference between the Def (CARRA2 settings) and the analysed forcing settings (anFOR), for March, April, June and July. Positive values indicate that the analysed forcing improves over the CARRA2 and vice versa. From the figure it is clear that at analysis time the analysed forcing improves by ~5-8%. However, this improvement at t0 (we use values from time-step zero, not analysed in the verification), is not preserved and for longer lead times as the differences are mostly zero. Except for July, where we see a degradation of the 2m temperature at longer lead times. This is something that we need to investigate further. One hypothesis could be non-optimal QC and OI settings. In these first experiments we used the default CARRA2 settings, which might not be optimal for the analysed forcing methodology. For example, in the Nordic-Analysis product which uses the similar approach for postprocessing NWP output, the OI correlation lengths are much shorter than what we have. When having too large correlation lengths we could have increments from stations that are far away from each other and not representative of the station being analysed.

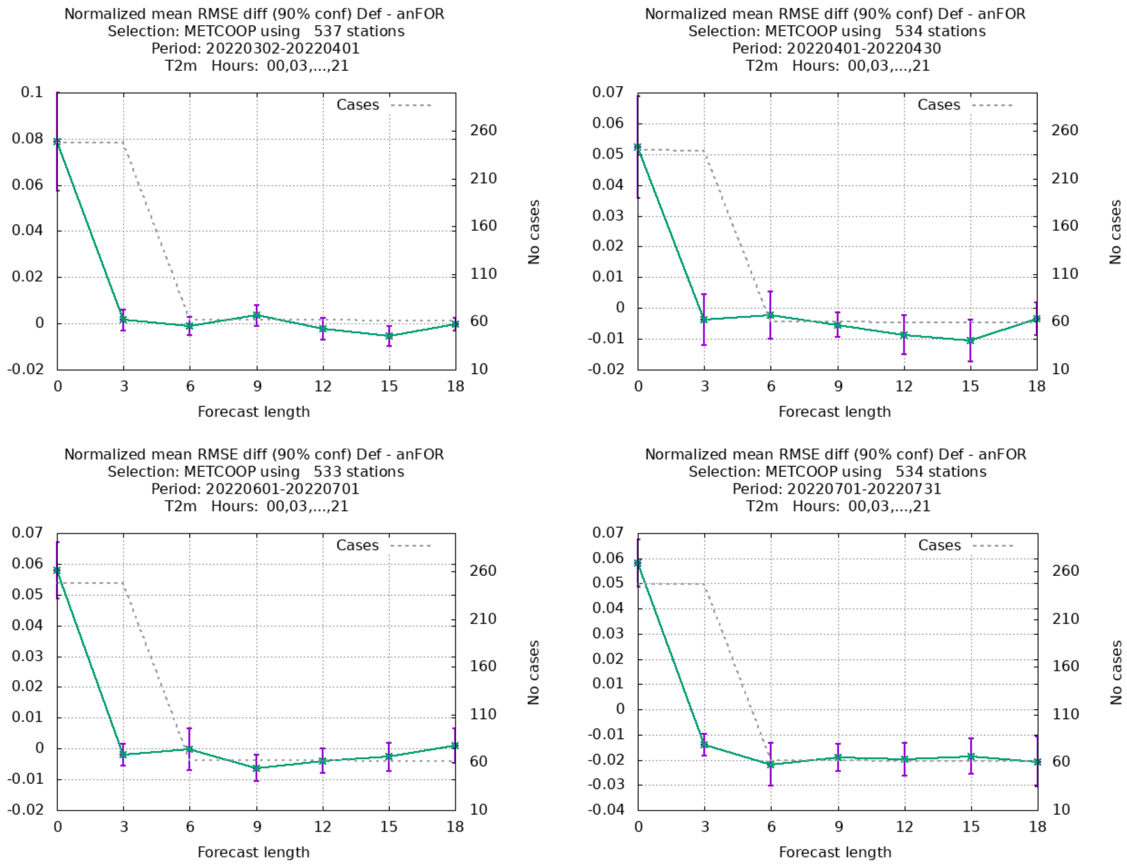

Figure 5: Normalized mean root-mean-squared-error temperature difference (green line) between the Def (CARRA2 settings) and the analysed forcing settings (anFOR), for March, April, June and July. Box and whiskers (purple) are 90 % confidence level and dashed black line is the number of cases.

A second way to improve is by analysing the precipitation, which will have a longer memory in the soil through soil moisture corrections, than when we only analyse 2m temperature and humidity.

In Fig 6, we have plotted the same metrics but for 2m specific humidity. Here we see that there are statistically significant improvements at lead time t0 and longer, for March, April and June. While we in July see a degradation of 2m specific humidity scores after t0.

Further analysis will look into how the analysed forcing methodology affects the snow during spring and fit to other observations. For independent verification of the horizontal OI we also blacklisted 33% of the WMO SYNOP stations in the domain, at random locations.

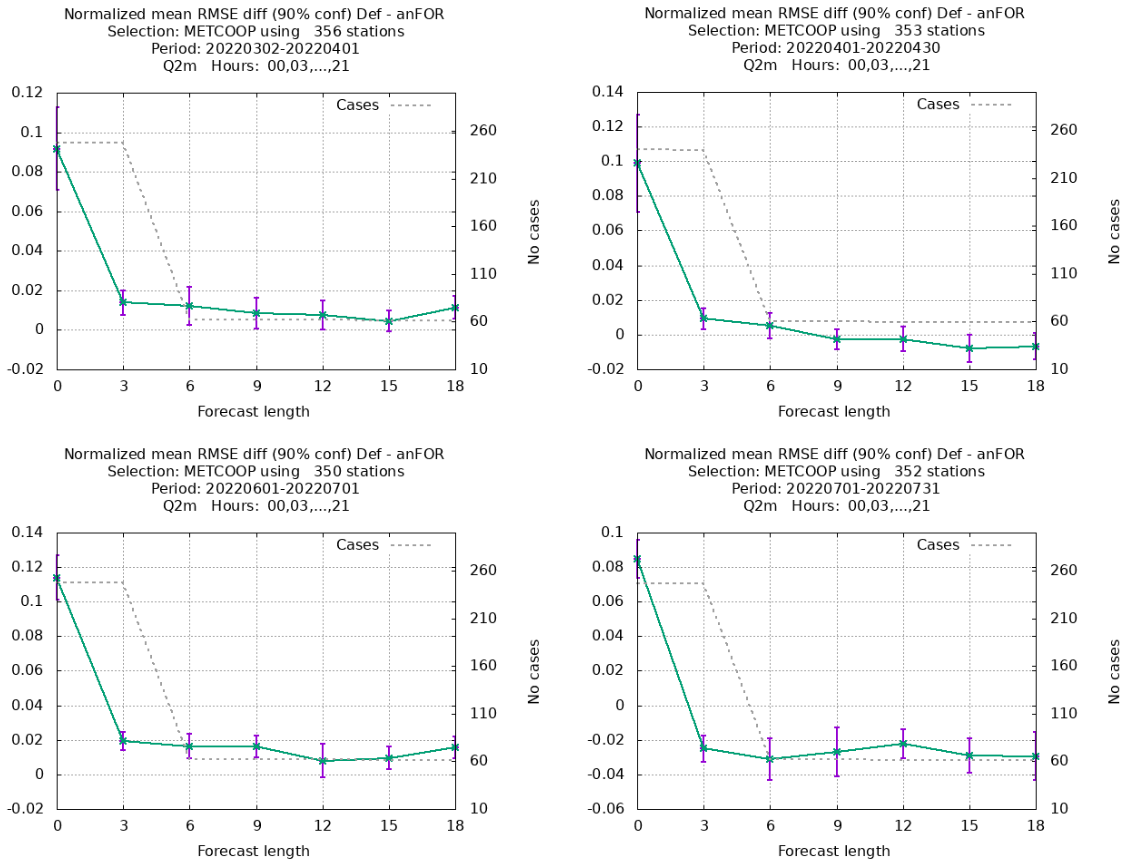

Figure 6: Normalized mean root-mean-squared-error specific humidity difference (green line) between the Def (CARRA2 settings) and the analysed forcing settings (anFOR), for March, April, June and July. Box and whiskers (purple) are 90 % confidence level and dashed black line is the number of cases.

4.4.2 Assimilation of screen level variables in 3D-VAR

We now look into the results from assimilating 2m temperature and relative humidity in the upper-air 3D-VAR. At the time of writing this report this experiment CARRA3_Pv1_screenDA Tab. 1. had been running for 20 days. In Fig. 7, we have plotted the bias and std of the errors for a) mean sea level pressure, b) 2m temperature, c) 2m relative humidity and d) 2m specific humidity. The CARRA3_Pv1_REF_Summer is shown in purple, CARRA3_Pv1_anFOR_Summer in green and the CARRA3_Pv1_screenDA in blue. It is clear from all variables that the assimilation of 2m temperature and humidity observations has a positive impact on the verification scores, lasting for most variables several hours into the forecast (18-hour here to replicate what is done in CARRA2).

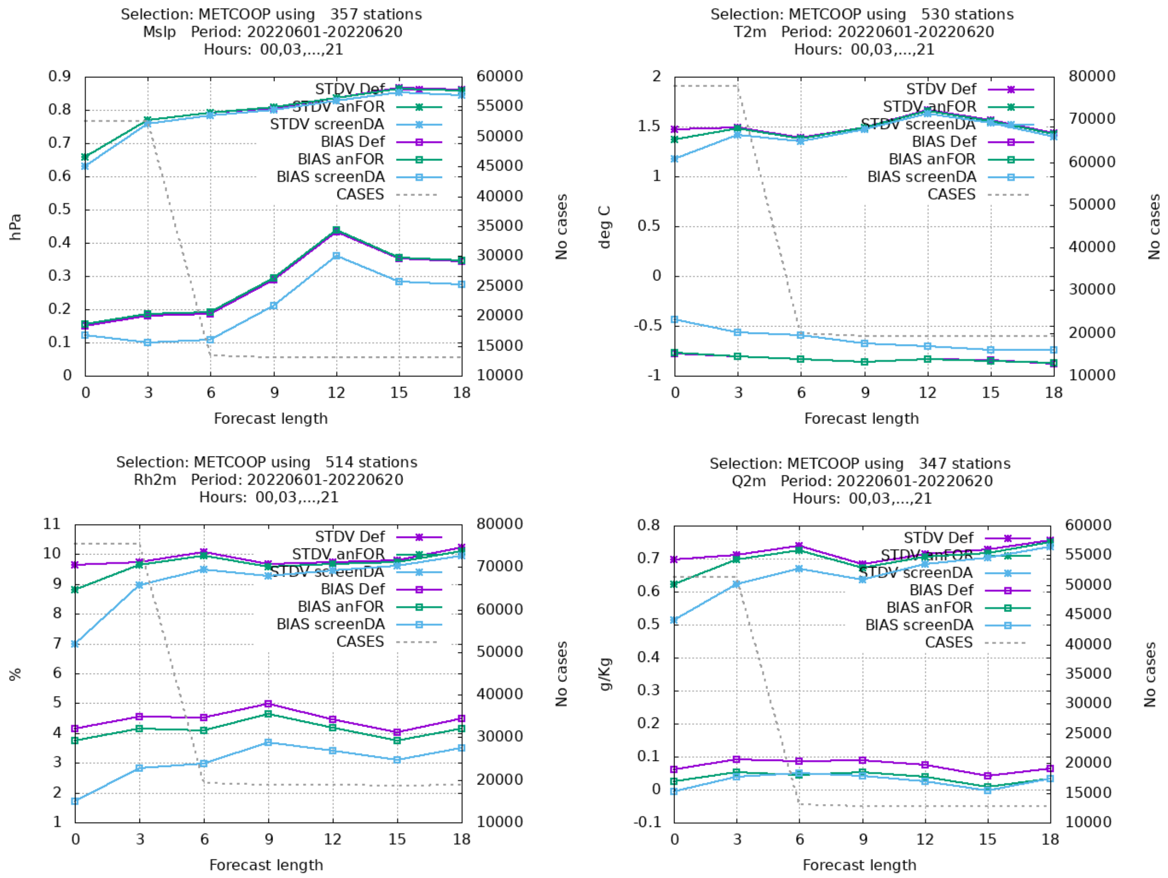

Figure 7: Bias and standard deviation of the errors for mean sea level pressure, 2m temperature, relative humidity and specific humidity. CARRA3_Pv1_REF_Summer (purple), CARRA3_Pv1_AnForc_Summer (green) and CARRA3_Pv1_screenDA (blue).

We also note the positive impact of the analysed forcing during this time-period. The effect of assimilating 2m variables on the upper air verification is shown in Fig. 8. Here we see that close to the surface the temperature bias and standard deviation of the errors are improved for the CARRA3_Pv1_screenDA experiment. The effect is less for specific humidity but the errors close to the surface are improved over the experiments not assimilating 2m variables.

In Fig. 9, we plot the normalized mean root-mean-squared-error temperature and specific humidity difference between the Def (CARRA3_Pv1_REF_Summer) and the analysed forcing settings (CARRA3_Pv1_screenDA). Positive values indicate that the assimilation of 2m temperature and specific humidity has improved the forecast verification scores. We see that the impact is positive throughout the 18-hour forecast and statistically significant (magenta bars).

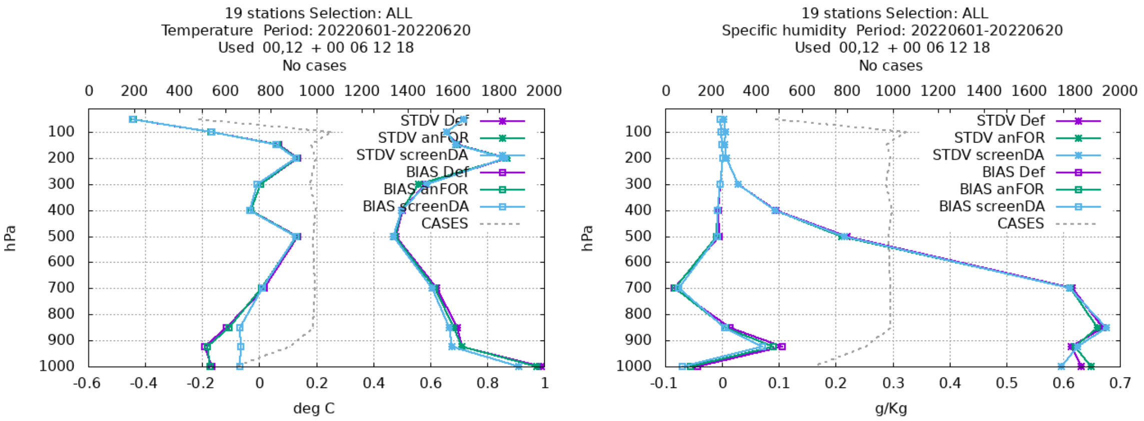

Figure 8: Bias and standard deviation of the errors for temperature and specific humidity. CARRA3_Pv1_REF_Summer (magenta), CARRA3_Pv1_AnForc_Summer (green) and CARRA3_Pv1_screenDA (blue).

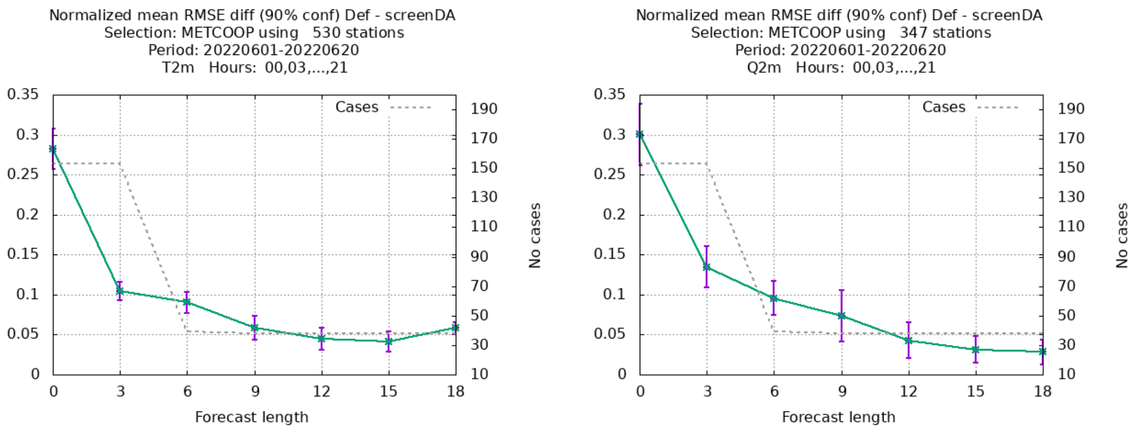

Figure 9: Normalized mean root-mean-squared-error specific difference between the Def (CARRA3_Pv1_REF_Summer) and the assimilation of 2m variables (CARRA3_Pv1_screenDA).

5 Outer loop land-atmosphere coupling

5.1 Overview

The 4D-VAR and its multi-incremental formulation with multi-step character provides a good framework to evaluate different levels and paths of coupling land and atmosphere data assimilation. As a first step, we study weak coupling by performing surface analysis at different stages of the atmospheric 4D-VAR, utilising non-linear trajectory runs.

5.2 Methodology

Two configurations for weakly coupling surface and upper air assimilation have been under consideration. Both of them exploit the influence that an updated initial surface field can have on the trajectory runs performed at the end of each outer loop in the atmospheric 4D-VAR. The surface analysis has been done through CANARI/OI (based on optimal interpolation), but the setup can be generalized to other methods of surface data assimilation.

The first setup consists of surface analysis performed before and after the upper air 4D-VAR. The analysis before is done at the beginning of the 4D-VAR assimilation window, impacting all the trajectory runs within the atmospheric 4D-VAR, and therefore the final upper air analysis field. On the other hand, the upper air analysis influences the surface field by also performing a surface analysis after 4D-VAR. In this case, the surface analysis is done at the time corresponding to the model initial state for the next forecast. The left diagram below shows the script construction for this setup.

For the second setup, surface analysis is performed within each outer loop in the atmospheric 4D-VAR. These analyses are done at the beginning of the 4D-VAR assimilation window and before the inner loops to influence the trajectory runs. The sequential nature of this setup has the benefit of continuously updating both surface and upper air fields for a more consistent model initial state. The right diagram below shows the script construction for this second setup.

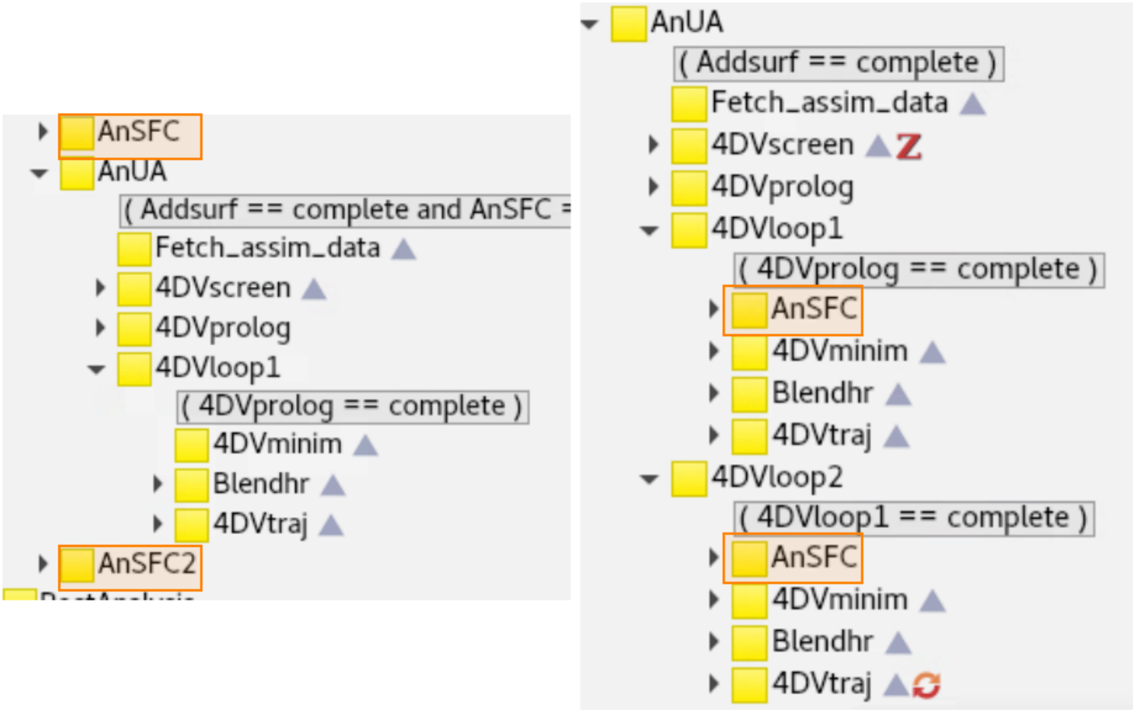

Figure 10: Block construction for the two set-ups under consideration, highlighting the surface analyses.

5.3 Results

The first setup has been run with one outer loop to explore the feasibility and technicalities of the approach. A reference case was also run, where surface analysis was only performed after the atmospheric 4D-VAR. In this way, one can compare directly the influence on the increments by performing a surface analysis before. This is shown in Fig. 11, where the increments of the soil temperature for the reference case, the weakly coupled case, and their differences are shown from left to right. From these preliminary results, one can see the surface analysis performed before the atmospheric 4D-VAR influences the one done after. A similar behavior has been seen for the upper air fields.

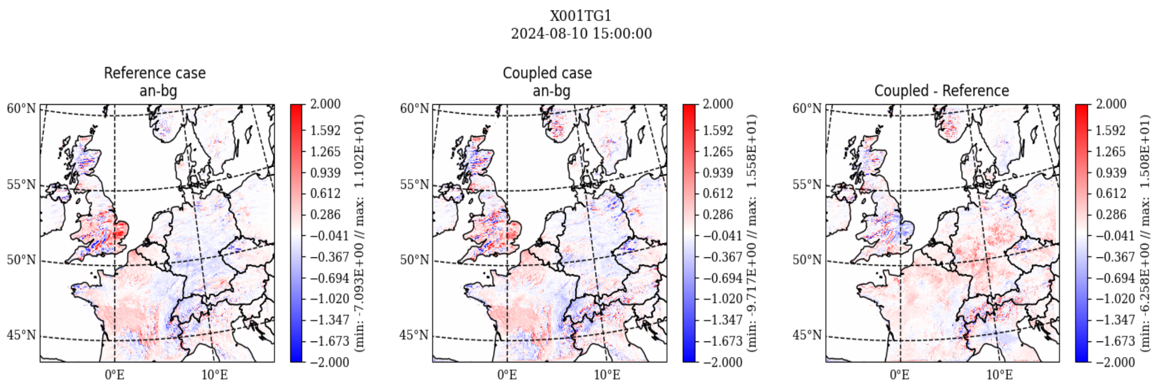

Figure 11: Soil temperature increments at model initial state time for reference case (left), first weakly coupling setup (middle), and their difference (right).

The second setup under study has been run with two outer loops, and preliminary results are shown in Fig. 12. Here, the 2m temperature increments are shown for each outer loop. It can be noticed how smaller and smoother increments are achieved in the second loop.

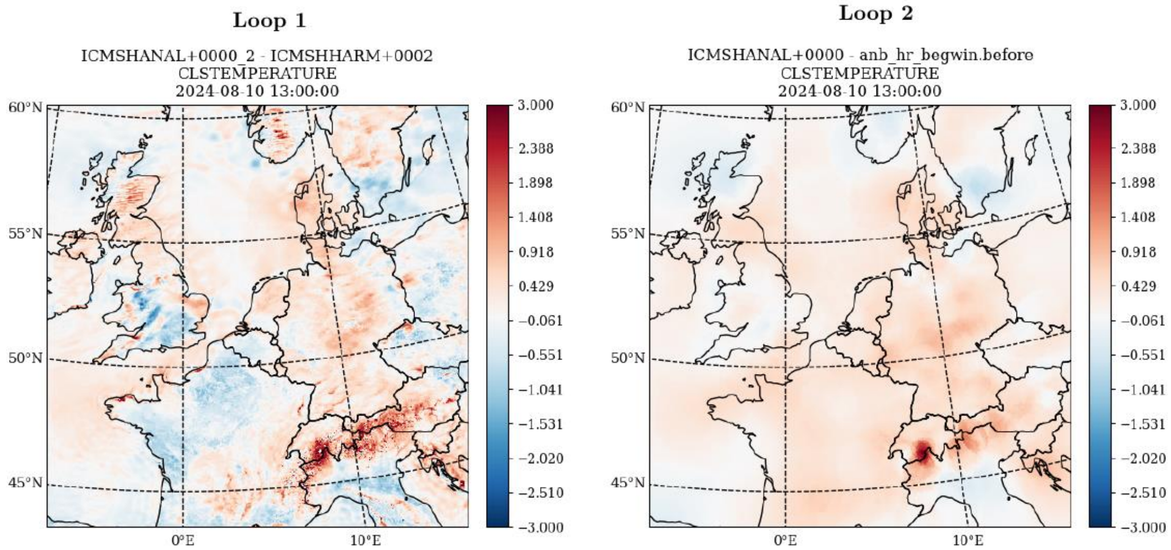

Figure 12: Analysis increments of the 2m temperature fields from each outer loop for the second set-up.

6 Summary and Conclusions

This report summarizes the infrastructure developments that were done in the HARMONIE-AROME system in preparation for the next generation of the C3S regional reanalysis. They include the capability to do ensemble surface data assimilation developed in WP1 in the coupled framework. Development of OOPS includes a 2D correlation operator that can take orography into account precisely. This correlation operator has been used to run a 2DVAR for screen variables (2m temperature and humidity), that could replace the existing OI.

Weakly coupled land-atmosphere data assimilation was developed for the 3D-VAR using analysed forcing and assimilation of interface observations (2m temperature and humidity). The analysed forcing methodology allows us to increase the observation usage within the assimilation window and use observations that are currently not used in our system (precipitation). The preliminary results show that the analysed forcing methodology mainly improves the time-step zero scores. This is by design, as we present results where we verify against the same observations as we assimilate (but not the analysis directly). In the analysed forcing experiments we blocklisted 33% of the observations with a valid station identifier from the analysis. Later evaluation will include these independent observations in the verification. Further work will investigate the results in more detail and perform tests with different settings for the horizontal OI and adding precipitation to the analyzed variables.

The assimilation of 2m temperature and humidity has a large positive impact on the surface (throughout the forecast) and upper-air verification scores. This is promising given the static background error covariance applied here and more tests will be done to tune the assimilation system. It is also possible to extend this to the 3D-EnVAR and the orography aware correlation function.

Two approaches for introducing a weak coupling between surface and upper-air have recently been introduced into the HARMONIE-AROME system. First evaluation of the functionality of these two approaches has been done. Next steps are to extend the first verification to include a longer time-period, to include analysis of precipitation (in the analysed forcing methodology) and to combine with the unified land data assimilation system developed in WP1. Work has also started on the outer loop coupling in the HARMONIE-AROME 4D-VAR in the reference system, and technical tests have been carried out performing the surface analysis within the outer loops. As it is outlined in the deviation plan the work on the implementation of the outer loop 4D-VAR coupling in the CERISE HARMONIE-AROME code base will be done in the extended T2.2.

7 References

Bakketun, Å., Blyverket, J., & Müller, M. (2023). Using a Reanalysis-Driven Land Surface Model for Initialization of a Numerical Weather Prediction System. *Weather and Forecasting*, *38*(11), 2155–2168. https://doi.org/10.1175/WAF-D-22-0184.1

Ingleby, B., Arduini, G., Balsamo, G., Boussetta, S., Ochi, K., Pinnington, E., & de Rosnay, P. (2024). Improved two-metre temperature forecasts in the 2024 upgrade. https://doi.org/10.21957/BI49S20QA8

Document History

Version	Author(s)	Date	Changes
0.0	Jostein Blyverket	27.10.2024	Initial version
0.1	Jostein Blyverket, Åsmund Bakketun, Benjamin Menetrier, Jelena Bojarova, Magnus Lindskog and Faúndez Alarcón José Manuel	28.11.2024	First full draft
0.2	Jostein Blyverket, Åsmund Bakketun, Benjamin Menetrier, Jelena Bojarova, Magnus Lindskog and Faúndez Alarcón José Manuel	02.12.2024	Version submitted for internal review
0.3	Jostein Blyverket, Åsmund Bakketun, Benjamin Menetrier, Jelena Bojarova, Magnus Lindskog and Faúndez Alarcón José Manuel	16.12.2024	Version addressing reviewer comments from NILU and MF.
1.0	As above	Dec 2024	Updated after internal reviews

Internal Review History

Internal Reviewers	Date	Comments
Jean-Christophe Calvet MF and Yvan Orsolini NILU	Dec 2024	Initial version