CopERnIcus climate change Service Evolution

D2.3 Documentation on coupled skin temperature assimilation for coupled global reanalysis

Due date of deliverable	31/12/2025
Submission date	15/12/25
File Name	CERISE-D2-3-V1.2
Work Package /Task	WP2 Task 2.3
Organisation Responsible of Deliverable	ECMWF
Author name(s)	Christoph Herbert, Peter Weston, Patricia de Rosnay, Sébastien Massart and Massimo Bonavita
Revision number	V1.2
Status	ISSUED
Dissemination Level	PU

Funded by the European Union

The CERISE project (grant agreement No 101082139) is funded by the European Union.
Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.

1 Executive Summary

The purpose of this work is to develop the next generation of C3S (Copernicus Climate Change Service) global scale reanalysis codes infrastructure to support modular coupled assimilation and monitoring. In the ECMWF Integrated Forecasting System (IFS), which is used to produce the C3S global reanalysis, we developed and tested innovative coupled surface-atmosphere assimilation with coupled skin temperature data assimilation using an eXtended Control Variable approach (XCV). This document describes the methodology of the XCV framework in the atmospheric analysis. The impact of different experimental configurations regarding the enabling of the first-layer soil temperature within the XCV framework and the assimilation of the corresponding analysis fields as pseudo-observations into the Simplified Extended Kalman Filter (SEKF) using an outer loop coupling approach is investigated and results are presented.

Table of 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.

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, R&I 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 deliverable

This deliverable describes the methodology and initial assessment of assimilating land-surface-sensitive observations. It focuses on the technical implementation of the required infrastructure in the IFS and presents preliminary results from experiments on coupled skin-temperature assimilation for global reanalysis.

2.2.2 Work performed in this deliverable

The preliminary assessment of the methodology towards the optimal degree of coupling - as described in Deliverable D2.1. - was based on IFS developments in "research mode" outside the Object-Oriented Prediction System (OOPS) environment that is used in operations (English et al., 2017). In this deliverable, the current work in progress outlined in WP2 T2.3:

"*Development of coupled skin temperature assimilation over land and sea ice"* is presented together with baseline testing. Further developments and results will be documented in a forthcoming paper that will support the production of the ERA7 reanalysis Pv2.

Section 3 presents the outer loop coupling infrastructure developments in the IFS under OOPS, including a flexible coupling methodology. The flexibility allows the land data assimilation system (LDAS) to be activated in selected outer loop(s), determining whether updated land and atmospheric conditions are produced and passed on to subsequent loops. These developments improve both efficiency and operational suitability.

The SEKF observation vector is extended to allow the exploitation of surface-sensitive observations. Section 4 describes the methodology of the eXtended Control Variable (XCV) framework in the atmospheric analysis and examines the impact of different experimental configurations. These include enabling the first-layer soil temperature within the XCV framework and assimilating the resulting analysis fields as pseudo-observations into the SEKF using an outer-loop coupling approach.

2.2.3 Deviations and counter measures

No deviations have been encountered.

2.2.4 Reference Documents

[1] Project 101082139- CERISE-HORIZON-CL4-2021-SPACE-01 Grant Agreement
[2] Project 101082139- CERISE-HORIZON-CL4-2021-SPACE-01 D7.2 Albedo, vegetation and LST satellite datasets in the CERISE verification database
[3] Project 101082139- CERISE-HORIZON-CL4-2021-SPACE-01 D1.2 Unified, ensemble-based global land data assimilation system and documentation
[4] Project 101082139- CERISE-HORIZON-CL4-2021-SPACE-01 D2.1 Documentation of coupled assimilation infrastructure and methodology and preliminary assessment towards optimal degrees of coupling for coupled global reanalysis

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 Outer loop coupling for enhanced usage of surface-sensitive observations

This section presents the outer loop coupling data assimilation workflow developed in the ECMWF IFS within the OOPS. This implementation supports operational use – in line with the strategy outlined by de Rosnay et al. (2022) – while allowing flexible testing of various coupling configurations, as described in Section 3.1, and prepares the system for improved exploitation of land surface–sensitive observations. Section 3.2 details the extension of the SEKF observation vector in the IFS to enable assimilation of gridded observations such as those provided by the XCV framework discussed in Section 4.

3.1 Implementation of outer loop coupled DA under OOPS

Outer loop coupling refers to a system where information produced in one outer loop – such as updated land or atmospheric analyses from independent DA approaches – are fed back into subsequent outer loops within the same data assimilation window. This iterative exchange allows the system to progressively refine the initial conditions and improve the overall analysis quality. The different outer loop-coupling configurations documented in Deliverable D2.1 were developed and assessed in “research mode” that is relying on both merged SEKF/4D-Var coupled trajectories and outer loop coupling based on IFS Cycle 49R1 outside OOPS, which is now used to run the operational assimilation system at ECMWF. This meant that all experiments documented in Deliverable D2.1 were run with OOPS switched off.

Further technical developments in the IFS were required for operational suitability to make the outer loop land-atmosphere coupling setup available under the OOPS environment. Building on the work in T2.1, the efficiency of the LDAS was improved by running the SEKF within the 4D-Var tasks using fields already stored in memory during the original trajectories. Now, the main 4D-Var trajectories also output the fields needed to provide the background, first guess, and quality control for the SEKF. This allows for more adaptive coupling regarding the initialisation of the LDAS in addition to addressing concerns about the resilience of the system whilst not affecting the time critical path. This development as part of T2.1 established the basis for a flexible setup regarding the use of atmospheric updated conditions that can be picked up by the LDAS.

The infrastructure has been developed in a flexible way to enable different fields to be output at different time frequencies as required by the SEKF. For example, the 2 metre temperature (T2m) and relative humidity (RH2m) fields are only output at analysis times, while the soil moisture fields are output at every time step to enable the most accurate first-guess departure calculations for ASCAT surface soil moisture observations. In addition, these fields can be output at each outer loop which is essential for the outer loop coupling developments. This development allows the SEKF to run without re-running the nonlinear trajectory and, together with other streamlining, makes it 80 times faster at Tco1279 (9 km) resolution than before. In this context, a “trajectory” refers to the model’s forward integration that simulates the evolution of the atmosphere over the assimilation window and provides the reference state for the data assimilation computation. This is a key aspect for operational applications and an important consideration for outer loop coupling where the SEKF will be required to run multiple times in a single data assimilation window.

The LDAS produces land-surface analysis fields as updated land initial conditions for all the SEKF control variables that are analysed in the unified system developed in WP1 (D1.2; Herbert et al., 2024), used for the atmospheric and land analyses in subsequent outer loops. Hereby, the analysis increments are computed and added to the background fields in the beginning of the DA window and interfaced with the OOPS environment to initialize the model trajectory.

Additional work involved IFS suites changes following the structure of the developments for ocean-atmosphere coupling (Laloyaux et al., 2016) where the LDAS is implemented in a separate family in parallel to the atmospheric minimization (similar to *nemovar* used for ocean DA). The LDAS family is configured for flexible activation up to the penultimate outer loop (see prepIFS settings in Figure 1), incorporating the two-dimensional Optimal Interpolation (2D-OI) and SEKF tasks in each outer loop, whereas the snow DA runs only in the first activated instance. The use of updated land conditions produced by the land analysis and subsequently used in the respective 4D-Var trajectory is provided under a switch (*LUPDATEFG*, default to true).

Figure 1: prepIFS settings to assimilate XCV fields into the LDAS using outer loop coupled DA including enabling outer loop coupling (LOUTLOOP_LDAS) with customizable activation of the land surface analysis tasks in the any up to the penultimate outer loop (ENABLE_OUTLOOP), and usage of most updated land conditions after each LDAS (LUPDATEFG).

Figure 2a illustrates the setup of the weakly coupled configurations as currently used in operations. Different outer loop coupled DA suites in the IFS are shown in Figures 2b and c. The configuration shown in Figure 2b is a baseline setup that – in case no updated land initial conditions are passed to the next outer loops (*LUPDATEFG* set to false) – is bit-identical with the weakly coupled system and is hence just an infrastructural change to the current system. Figure 2c shows the LDAS activated in *uptraj_1* and *uptraj_2*, triggered once the atmospheric trajectory (traj) in the respective outer loop is complete and the required fields are available.

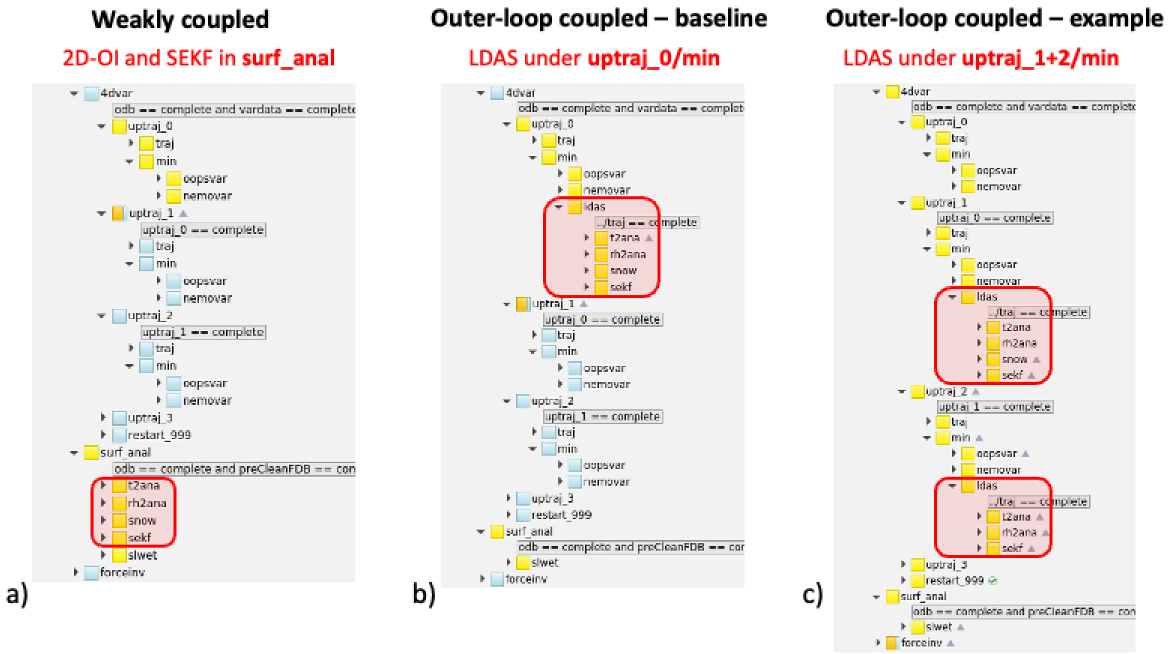

Figure 2: IFS suites setup in ecflow for different land-atmosphere coupled DA configurations; (a) currently operational weakly coupled system, (b) outer loop coupled DA with the LDAS (2D-OI, SEKF and snow DA) activated in the first outer loop (uptraj_0), and (c) outer loop coupled DA with the LDAS activated in uptraj_1 and 2).

3.2 Extending the SEKF observation vector in the IFS

The SEKF observation vector has been extended to enable enhanced use of surface-sensitive temperature observations in the LDAS. Like the T2m variables produced by the 2D-OI screen-level analysis that are used as pseudo-observations in the SEKF, the new implementation allows the surface temperature observations to be assimilated as gridded fields provided at the two synoptic times within each 12-hour data assimilation window (i.e., 00 and 06 UTC, and 12 and 18 UTC, respectively).

When using the XCV framework, which generates 4D-Var–based analysis fields – for e.g. skin or soil temperature – as surface-sensitive pseudo-observations (as described in Section 4.2), the layer to be analysed is required to be customisable depending on the information content of the observation. Thus, the development has the flexibility to use the new observations for analysing soil temperature at layers 1 to 3 (down to 1 metre depth), depending on the estimated depth sensitivity of the corresponding observation type. As the sensitivity between the new observation type and soil temperatures, the Jacobians are set to 1 for the first-layer soil temperature, while the EDA covariances between STL1 and STL2/3, respectively, are used to compute the Jacobians with respect to the deeper layers. The observation error can be estimated using Desrozier's method (see more detail in Section 4.2).

4 Investigating surface temperature in the outer loop coupled setup

4.1 Overview

Developments focused on the outer loop coupled assimilation methodology for enhanced use of interface observations for increased consistency in land-atmosphere assimilation for global reanalysis purposes.

The outer loop land-atmosphere coupling methodology implemented in T2.2 was advanced in T2.3 as described in Section 3 to investigate the usage of surface temperature in a coupled setup. Several configurations were tested to identify the impact of enabling the first-layer soil temperature (STL1) as part of the atmospheric 4D-Var, and the assimilation of the provided STL1 fields into the SEKF to analyse first-layer soil temperature in the land. The experiments were carried out based on outer loop land-atmosphere coupled DA by running short term numerical experiments with the global IFS.

4.2 Methodology

The physical variables currently analyzed in the 4D-Var are only atmospheric variables. The first-layer soil temperature over land was added to these control variables thanks to the XCV framework (Massart, 2023). The IFS includes a tangent linear and adjoint version of the physical processes related to STL1. All observations sensitive to the surface within the assimilation window are therefore contributing to the optimization of the STL1 field. This is to allow an optimal usage of the observation at the interface between land and atmosphere. Within the XCV framework, the analyzed STL1 (as referred to XCV-STL1 in the following) obtained at the beginning of the atmospheric assimilation window is propagated in time by the land-atmosphere coupled model and the forecasted value is used for the next DA window. This forecasted value can also be used in the surface analysis.

The gridded XCV-STL1 analysis fields are written out in the 4D-Var trajectory at synoptic times (00, 06, 12, and 18UTC) in each outer loop. These 2D fields can be directly assimilated as pseudo-observations into the SEKF with the expanded observation vector developments described in Section 3.2. The Desroziers (2009) method has been used to estimate the observation error of the XCV-STL1 field when added as a new observation type to the LDAS, in addition to assimilating T2m used to analyse multi-layer soil temperature. An observation error of ~1.08K has been computed with Desroziers based on the computation of the covariance matrix between the unbiased first-guess and analysis departures – using the departures obtained from experiments based on active and passive observations data.

Preliminary experimentation showed that inflating the observation error to 2K – relatively to using 1K – has no benefit, revealing slightly degraded forecast scores, e.g. for temperature at 1000hPa (Figure 3a). Because the XCV-STL1 is representing the model STL1 to be analysed, the Jacobian element between the XCV and the STL1 control variable is selected to 1 (Section 3.2). The updating of STL2 was tested to propagate the information content deeper into the soil beyond the first model layer of 7 cm depth – using the EDA-derived Jacobians computed from the covariances between model STL1 and STL2. However, initial verification over a short testing period, in which both STL1 and STL2 were analysed, in comparison to analysing STL1 only, has shown no benefit (see comparison against surface-sensitive T2m and Q2m observations in Figure 3b). In all experimental configurations described in the following, XCV-STL1 is assimilated into the SEKF using an observation error of 1 K, with updates applied solely to STL1.

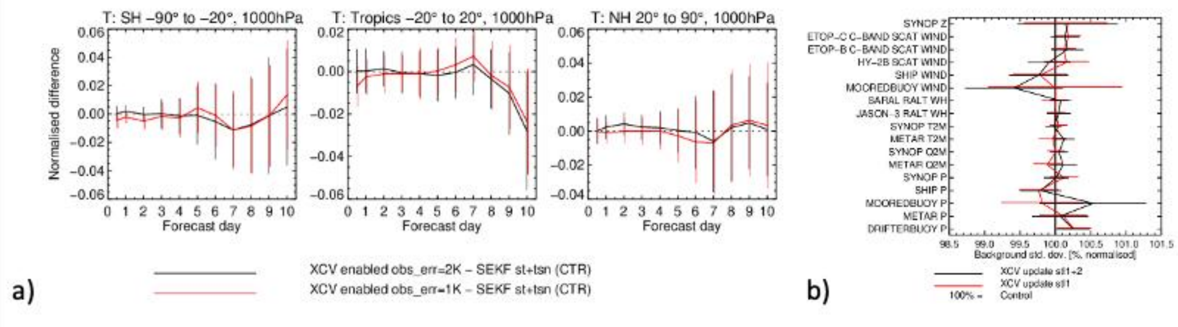

Figure 3: Verification of preliminary experiments to test the observation error and the feasibility of analyzing soil temperature at deeper layers using XCV: (a) RMS forecast error for air temperature at 1000hPa using an observation error of 1K (red curve) compared to 2K (black curve) ; (b) First-guess departures of analyzing STL1/2 (black curve) compared to STL1 only (red curve) .

4.3 Experiments

This section describes the setup of the NWP coupling experiments for a preliminary assessment of the impact of enabling first-layer soil temperature (STL1) in the control vector of the atmospheric 4D-Var, and assimilating the analyzed fields into the land – in the context of outer loop coupled DA. Enhanced land-atmosphere coupling in a "quasi-strong" manner can be achieved either through balancing updated land and/or atmospheric initial conditions between the two Earth-system components, or via assimilating the XCV-STL1 from the atmosphere into the land by exploiting their information content as surface-sensitive pseudo-observations. Enabling the XCV-STL1 in the 4D-Var control vector – without utilizing the STL1 analysis produced during the optimization – reveals largely neutral atmospheric scores. Therefore, the evaluated experiments and results that are solely based on enabling the XCV-STL1 are not discussed further.

The scientific experiments are conducted based on the implementation described in Section 4.2. Hereby, the impact of the XCV-STL1 is verified against the current weakly coupled DA system using IFS Cycle 50R1 developmental version 14. The following configurations (summarized in Table 1) are tested in NWP experiments over a three-month period during the boreal summer (JJA) of 2022.

(1) The weakly coupled system serves as the reference control setup.

The following setups are tested in which the XCV-STL1 is enabled in 4D-Var with the analysed STL1 fields produced in each outer loop:

(2) Enhanced coupling through pseudo-observations: the LDAS is activated in uptraj_2 and the most accurate XCV-STL1 analysis fields – produced in the final tangent-linear model trajectory – are assimilated into the SEKF. Hereby, the LDAS is initialised from the background (BG) fields without exchanging land and atmospheric initial conditions. This is to capture only the isolated effect of using the surface-sensitive pseudo-observations, without mixing in the influence of updated atmospheric conditions.

(3) Enhanced coupling through balanced conditions: enabling the XCV-STL1 (as sink variable) without assimilating the analysed STL1 fields, using the strongest coupling configuration investigated in D2.1, where the LDAS is activated in three outer loops (uptraj_0-2). In each outer loop, updated land and atmospheric conditions are exchanged between components by using the most updated first-guess (FG) fields available in the respective trajectory.

(4) Enhanced coupling through pseudo-observations and balanced conditions: combining (2) and (3) where the LDAS is activated in three outer loops (uptraj_0-2) and the propagated XCV-STL1 analyses by each model trajectory is assimilated into the SEKF in each outer loop.

Table 1: List of 3-month NWP experiments along boreal summer (JJA) 2022. In all experiments, LDAS is activated in uptraj_2 (or more) outer loop(s).

	Description	LDAS in outer loop(s)	XCV enabled in 4D-Var	XCV used in LDAS	LDAS First Guess
1	Weakly coupled (CONTROL)	1 (no land feedback)	-	-	BG*
2	XCV-STL1 enabled and assimilated	1	yes	yes	BG*
3	LDAS in uptraj0-2, XCV-STL1 enabled	1, 2, 3	yes	-	Updated 4D-Var outer loop
4	LDAS in uptraj0-2, XCV-STL1 enabled and assimilated	1, 2, 3	yes	yes	Updated 4D-Var outer loop

* The LDAS is based on the background (BG) without using updated initial conditions

4.4 Results

This section shows the impact in atmospheric forecast skill of all experiments based on configurations (2) - (4), compared to (1), as described in Table 1 in Section 4.3, verified against the operational IFS Cycle 47R3 analysis.

Figure 4 shows the relative differences in RMS forecast error of T2m for different lead times and latitude ranges – Southern (SH) and Northern hemispheres (NH) and tropics. A negative (positive) relative difference indicates improvement (degradation), and results are considered significant when the error bars do not cross the zero line, which represents the weakly coupled setup as the reference. Enhanced coupling through simply assimilating STL1 into the land (green curve) leads to small but non-significant improvements in the tropics with neutral scores at higher latitudes. The strongest outer loop coupling configuration in combination with the

XCV-STL1 assimilated into the SEKF (black curve) in each outer loop reinforces the positive impact in comparison to the purely outer loop coupled setup with solely enabled XCV (red curve). Overall, improvements in T2m are small (~1-2 %) but significant for more than five forecast days, spanning the NH and the tropics where most of the land masses are located.

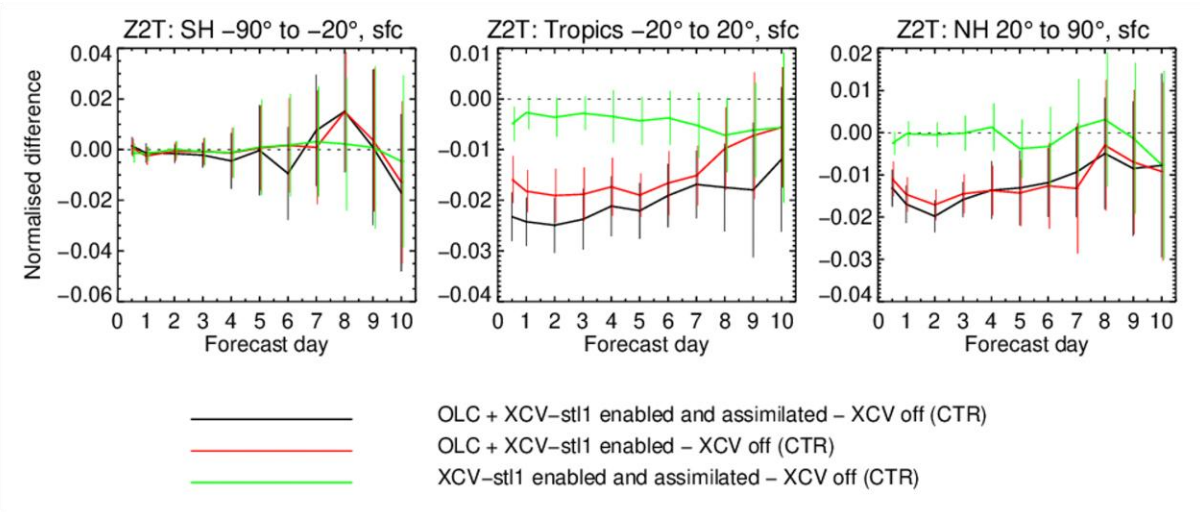

Figure 4: Normalised differences in RMS forecast error of T2m along boreal summer 2022, comparing configurations regarding STL1 enabled and assimilated in uptraj_2 without exchanging land and atmospheric initial conditions (green line), the LDAS activated in uptraj_0-2 with STL1 enabled (red line), and LDAS activated in uptraj_0-2 with STL1 enabled and assimilated into the SEKF in each outer loop (black line), compared to the weakly coupled system used as control (zero horizontal dotted line). Scores are verified against the operational IFS analysis.

Figure 5 shows the spatial distribution of relative differences in RMS forecast error of T2m and RH2m for the first three forecast days, comparing the strongest coupling configuration with XCV-STL1 data assimilation (4) to the weakly coupled control (1). Regarding T2m forecast (see Figure 5a), except for Australia, most regions within the tropical belt exhibit small improvements that consistent throughout the forecast period. This might show the potential of increased coupling in tropical areas, where forecasts are generally difficult to improve through changes in the LDAS alone, especially over convective zones. However, it has to be further investigated and validated to confirm its effectiveness. Similar as for T2m, the corresponding results for RH2m (Figure 5b) indicate that enhanced coupling causes consistent degradation over Australia at all lead times. Elsewhere, the impact is mostly neutral, with some additional degradation observed over southern Africa.

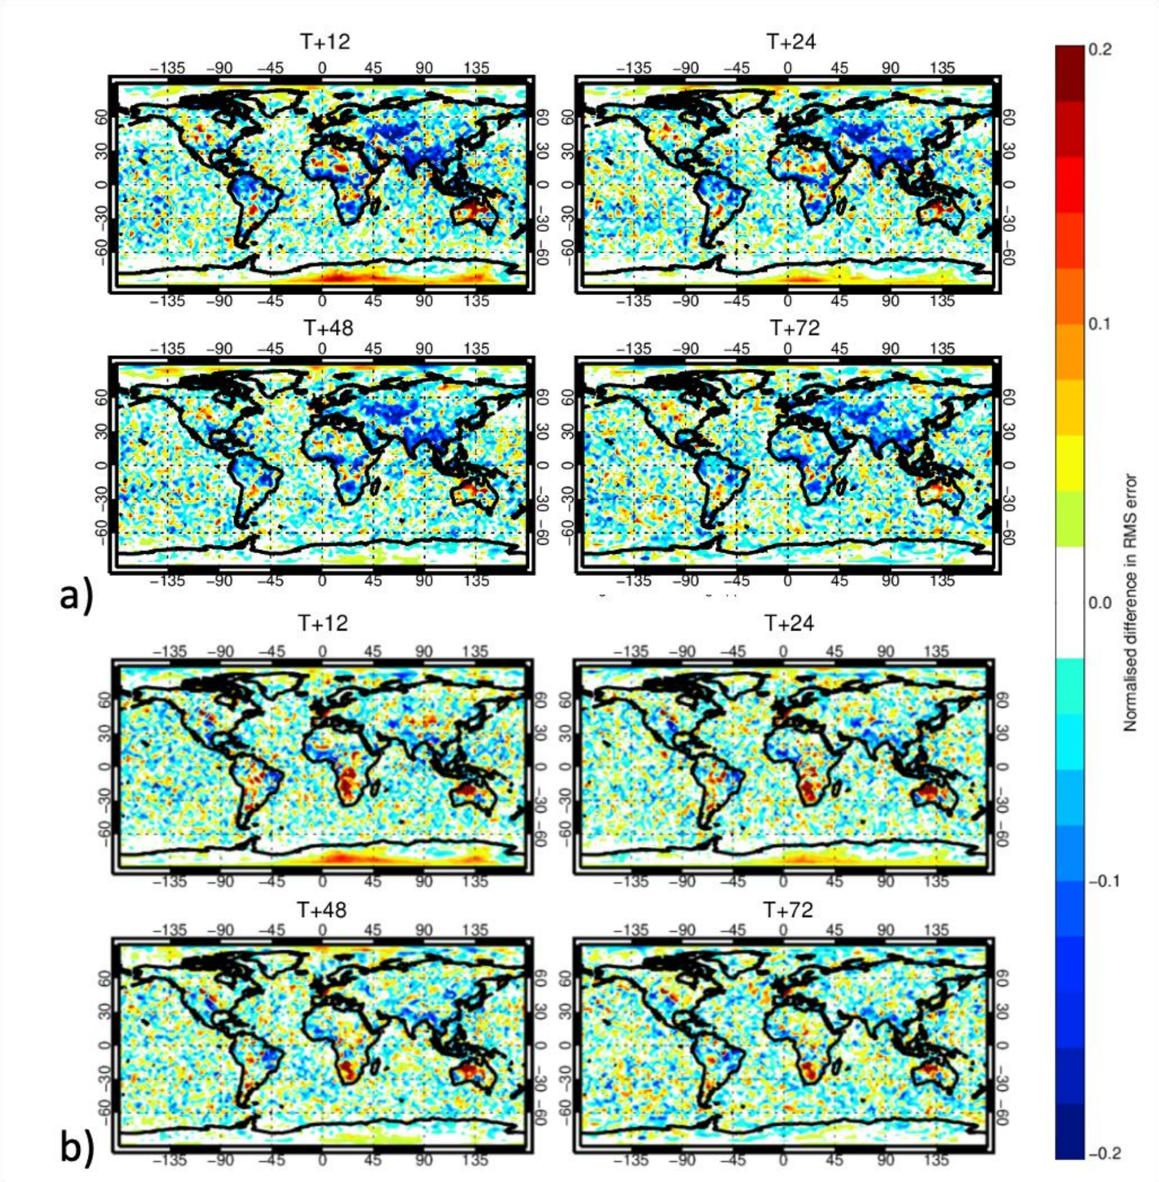

Figure 5: Relative differences in (a) T2m and (b) RH2m RMS forecast error for different lead times (12 hour to 3 days) verified against operational analysis, averaged over the NH, the tropics and the SH along the boreal summer 2022, comparing configuration with the LDAS activated in uptraj_0-2 with STL1 enabled and assimilated into the SEKF in each outer loop with the weakly coupled system.

Figure 6 shows the global averages of FG departures RMS against surface-sensitive observations comparing all three configurations with that of the weakly coupled control taken as a reference (vertical line). Similar to Figure 4, an increase (decrease) in RMS indicates a degradation (improvement) of the background forecast's fit to observations. The verification regarding T2m SYNOP and METAR stations reveals that enabling and assimilating XCV-STL1 without outer loop coupling is degrading the forecast scores – corresponding to increased first-guess departures (green line). In contrast, assimilating XCV-STL1 (red and black lines) tends to reduce the T2m RMS, with somewhat larger benefits observed when combined with outer loop coupling (black line). As previously seen for the verification against operational analysis, scores on a global average are degraded for RH2m (Q2M on the figure), whereas both outer loop coupling with or without assimilating the XCV shows only little differences.

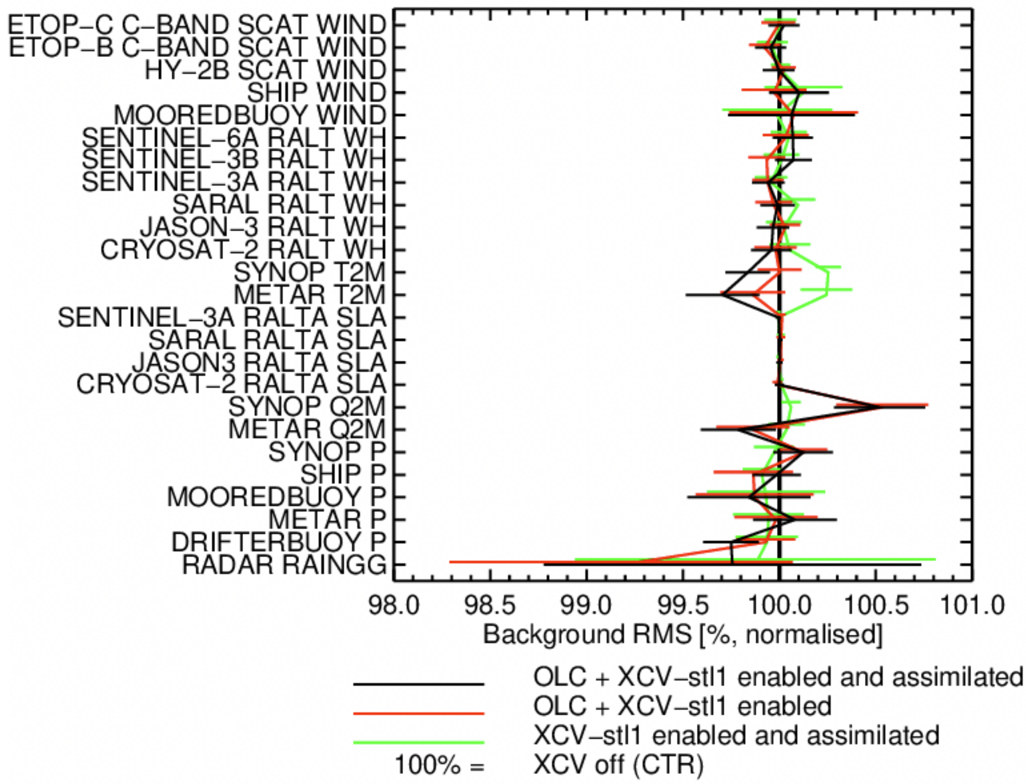

Figure 6: Globally averaged relative differences in RMS first-guess departures regarding near surface-sensitive observations along the boreal summer comparing configurations regarding STL1 enabled and assimilated in uptraj_2 without exchanging land and atmospheric initial conditions (green line), the LDAS activated in uptraj_0-2 with STL1 enabled (red line), and LDAS activated in uptraj_0-2 with STL1 enabled and assimilated into the SEKF in each outer loop (black line), verified against the weakly coupled system (vertical line at 100%).

Figure 7 illustrates the Scorecard for the experiments based on configuration (4) compared to (1) for observations verified against observations – showing significant impact for a 95% confidence interval – indicated by framed boxes – for different atmospheric variables and lead times, averaged over different areas. Significant impact is shown for T2m and 2 metre dewpoint temperature (D2m) forecasts in the medium range in the NH and tropics. Forecast scores for other observations are generally neutral, showing no systematic impact. Upper air atmospheric scores are also mostly neutral. Similar as seen in the spatial distribution of the forecast error against operational analysis in terms of RH2m (Figure 5b), for Australia and New Zealand, the scores against observations for D2m are also significantly degraded.
Figure 8 presents in more detail the actual magnitude of the relative differences in first-guess departures for T2m and D2m, corresponding to the boxes in the Scorecard for the NH, tropics and Australia/New Zealand, with positive values indicating improvements relative to the weakly coupled control experiment. As shown in Figure 8a-d, the tropics and NH exhibit

improvements of comparable magnitude and lead time (slightly less for D2m), whereas Australia and New Zealand (Figure 8e–f) display opposite behavior.

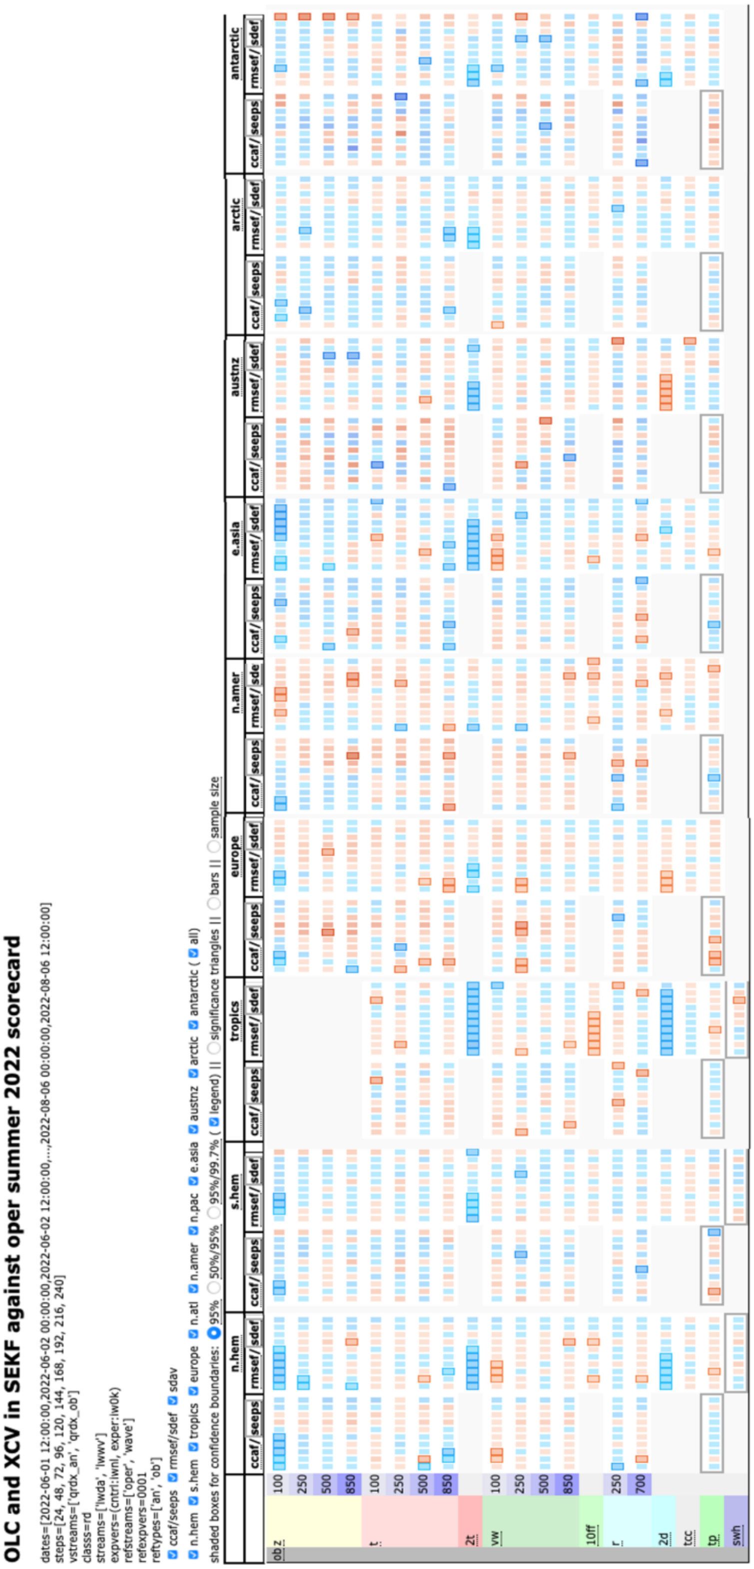

Figure 7: Scorecard showing relative differences in first-guess departures against observations for different areas (columns), variables (rows) and lead times (boxes), with framed boxes referring to a 95% confidence interval; blue ≙ improvements; red ≙

degradations, comparing the configuration with the LDAS activated in uptraj_0-2 with STL1 enabled and assimilated into the SEKF in each outer loop with the weakly coupled system.

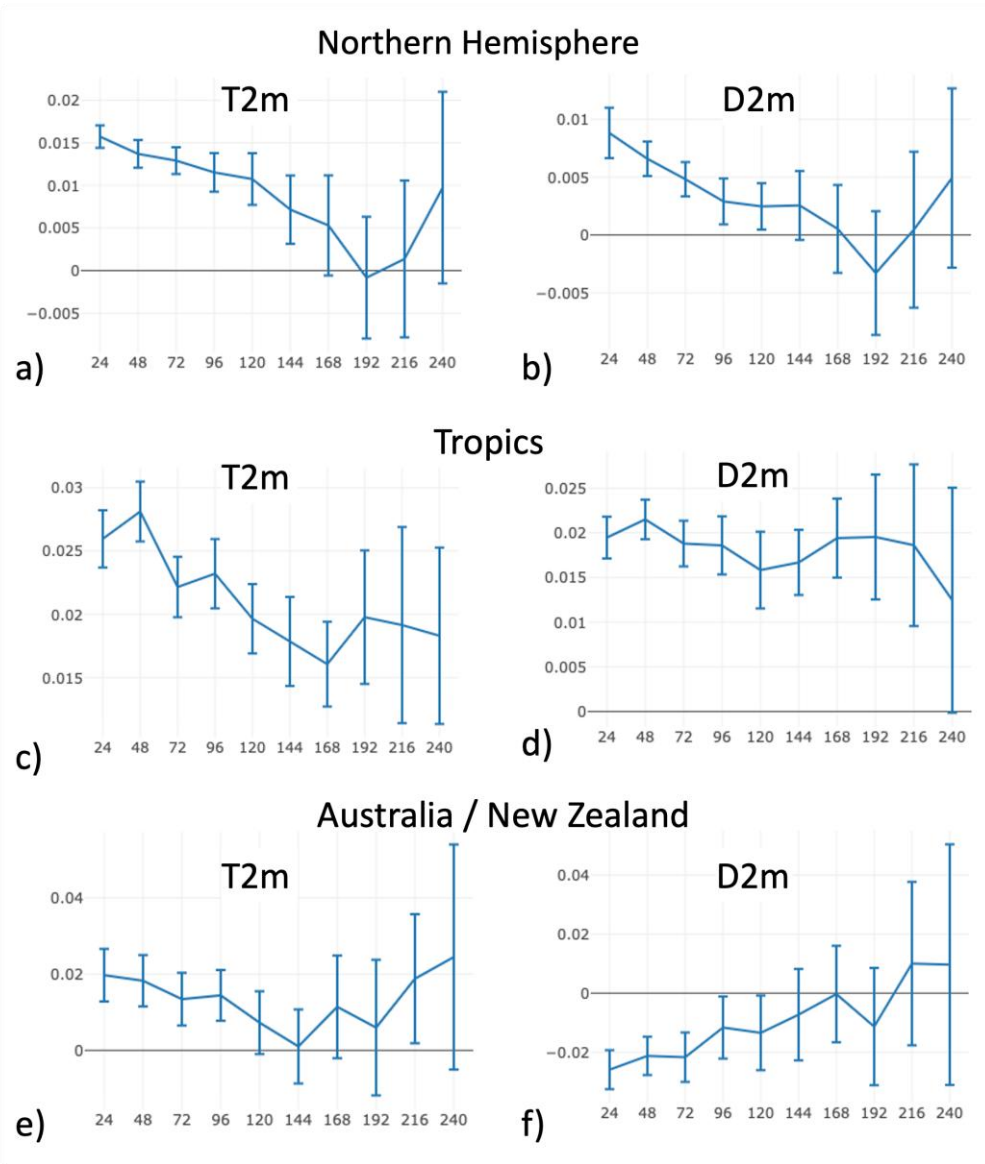

Figure 8: Relative differences in first-guess T2m and D2m departures against observations for the Northern Hemisphere (a and b), the Tropics (c and d), and Australia/New Zealand (e and f), verified against observations, comparing the weakly coupled control to the configuration with the LDAS activated in uptraj_0-2 with STL1 enabled and assimilated into the SEKF in each outer loop. Positive values indicate improvements with the coupled data assimilation system.

4.5 Discussion

The XCV framework allows adjusting the background of land interface variables (skin temperature, first-layer soil temperature, etc.) by including the respective tangent-linear and adjoint model in the 4D-Var optimization. Hereby, the increments added to the XCV variables are based on the entire observational network assimilated into 4D-Var, allowing to investigate the exploitation of atmospheric observations to better constrain the surface. Yet, unlike for the atmospheric state variables that are well-constrained having the Jacobians of satellite channels spanning several variables with mixed vertical sensitivity, there are no explicit land observations assimilated into 4D-Var and land surface-sensitive channels are deliberately filtered. Thus, applying the XCV framework in an outer loop coupled DA context allows the respective land-surface variables that are included in the XCV to be indirectly constrained by the observations assimilated in the LDAS through balancing between land and atmosphere.

As pointed out in Section 4.3, enabling the XCV-STL1 in the 4D-Var control vector without assimilating the analysis into the LDAS does not improve the forecast scores. This can be due to the model's short forecast of STL1, which is used to generate the land initial conditions for the next DA window, being initialized from the SEKF analysis rather than the XCV analysis obtained from 4D-Var.

This may reflect a physical system memory effect, with the land model having a long memory and therefore retaining the information from the STL1 updates in the forecast, while the atmosphere has a short memory and loses that information quickly. Also, enabling XCV solely as a sink term might lead to better balancing of the information content of the observations that are assimilated into the atmosphere among the atmospheric state variables, but the corresponding increments are only added to STL1 and used in 4D-Var, and not propagated to the next DA window, which might cause inconsistencies between land and atmosphere. In case the XCV analysis fields are assimilated into the outer loop coupled LDAS as pseudo-observations, the information content of the atmospheric observations is propagating to the land and hence influences the analysis the short forecast is initialised from with notable impact on the initial conditions of the next DA window.

Similar as in the exploration of different outer loop coupling configurations (results presented in D2.1), enhanced coupling through activation and assimilation of the XCV-STL1 degrades the skin temperature forecast in some areas (not shown here – see D2.1). Further experiments were conducted and are briefly presented below.
Two methods have been investigated to alleviate the negative impact that goes hand in hand with enhanced coupling. First, scaling the observation error covariance matrix used in the SEKF across outer loops – specifically by inflating R in the later outer loops – rather than keeping it constant leads to comparable or improved near-surface NWP scores (see reduced first-guess departures in Figure 9). Hereby, the land realism is better maintained because less increments are added to the land variables while the positive land-atmosphere coupling effect of balancing the initial conditions is still given. Second, optimizing the skin thermal conductivity of land types – specifically by reducing the conductivity over bare soil – resulted in a more stable diurnal cycle. The latter one serves as an example of how data assimilation can be used for model parameter estimation, a topic that remains under active investigation.

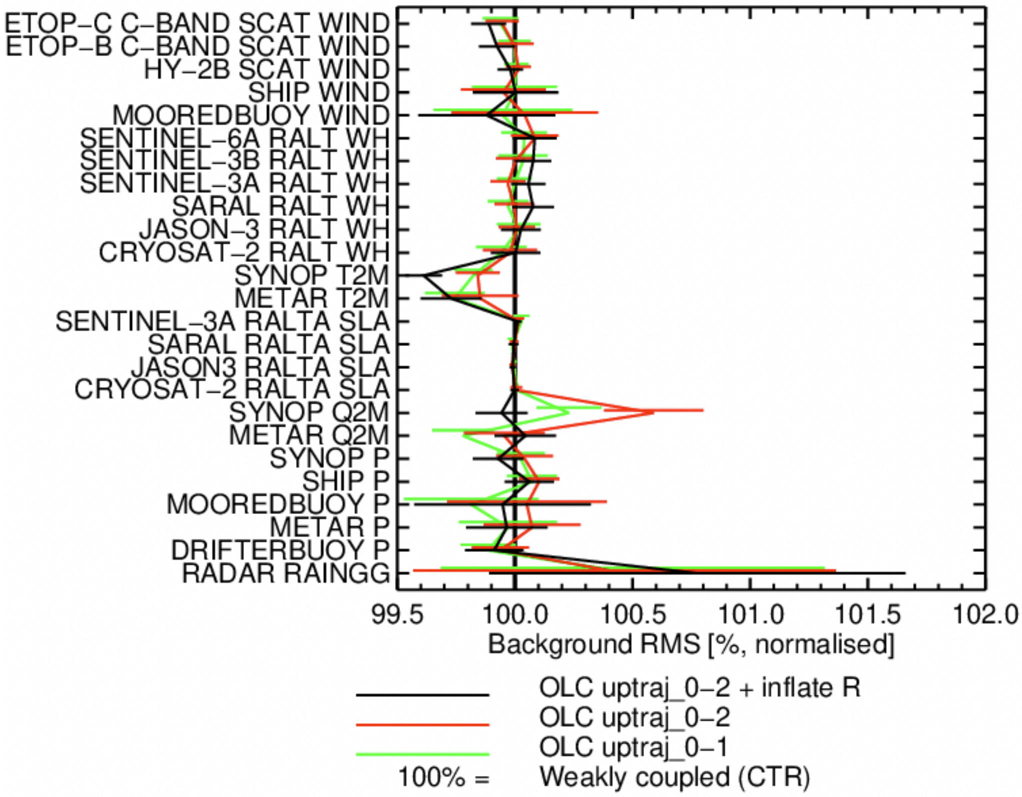

Figure 9: Globally averaged relative differences in RMS first-guess departures regarding near surface-sensitive observations along the boreal summer comparing configurations regarding the LDAS activated in uptraj_0-1 with constant R (green line), the LDAS activated in uptraj_0-2 with constant R (red line), and LDAS activated in uptraj_0-2 with inflated R for later outer loops (black line), verified against the weakly coupled system (vertical line at 100%).

5 Summary and Conclusions

This report finalises the major infrastructure developments regarding outer loop land-atmosphere coupling that were conducted in the ECMWF IFS system in preparation of the next generations of the C3S for global reanalysis and seasonal prediction. The new system provides an efficient and flexible setup for exchanging updated land and atmospheric initial conditions between components in different outer loops under the OOPS environment, making it well-suited for operational use at ECMWF. This report summarizes the implementation carried out in WP2 T2.3, including the conducted NWP experiments and initial assessment of an enhanced coupled system designed to better exploit currently underused land surface-sensitive temperature observations within an outer loop coupled DA framework. In particular, the first-layer soil temperature – provided within the 4D-Var optimization through the XCV framework – is assimilated as pseudo-observations to update soil temperature in the LDAS.

Enabling and assimilating the XCV first-layer soil temperature in combination with the strongest coupling configuration, where the LDAS is activated in the first three 4D-Var outer loops, revealed conditionally improved forecast skill for the near-surface variables. The results suggest a potential benefit of enhanced coupling when combining the update of initial conditions with the exchange of analysis fields through pseudo-observations. This work indicates that interface observations may help constrain the land and atmosphere through balancing, with potential benefits for surface temperature that can vary considerably within the 12-hour data assimilation window; however, these benefits need to be further investigated and evaluated.

In addition, enhanced coupling can involve large updates to the land surface or a suboptimal distribution of increments within the model, which manifests as local degradations in skin temperature (as discussed in D2.1) and relative humidity in regions that may be affected by model bias such as Australia and New Zealand. Moderating the increments applied to the land – for example, by scaling the SEKF observation error across different outer loops – offers a way forward to preserve land realism while still leveraging the benefits of balanced initial conditions in the outer loops. Exploring the use of additional interface observations or integrating further surface-sensitive channels within the XCV framework promises to further constrain the land surface and will be a focus of future work, with the aim of addressing challenges arising from heterogeneous land surfaces.

6 References

Desroziers, Gérald, Loık Berre, and Bernard Chapnik: Objective validation of data assimilation systems: Diagnosing sub-optimality, In *Proceedings of ECMWF Workshop on diagnostics of data assimilation system performance*, pp. 15-17. 2009.

de Rosnay, P., Browne, P., de Boisséson, E., Fairbairn, D., Hirahara, Y., Ochi, K., et al. (2022) Coupled data assimilation at ECMWF: current status, challenges and future developments. Quarterly Journal of the Royal Meteorological Society, 148(747), 2672–2702. Available from: https://doi.org/10.1002/qj.4330

English, S., D. Salmond, M. Chrust, O. Marsden, A. Geer, E. Holm et al., 2017: Progress with running IFS 4D-Var under OOPS, ECMWF Newsletter No. 153, pp. 13–14. https://www.ecmwf.int/en/newsletter/153/news/progress-running-ifs-4d-var-under-oops

Herbert C., P. de Rosnay, P. Weston, and D. Fairbairn: Towards unified land data assimilation at ECMWF: Soil and snow temperature analysis in the SEKF, QJRMS 2024, DOI: https://rmets.onlinelibrary.wiley.com/doi/10.1002/qj.4808

Ingleby B., G. Arduini, G. Balsamo, S. Boussetta, K. Ochi, E. Pinnington, and P. de Rosnay: "Improved two-metre temperature forecasts in the 2024 upgrade", ECMWF Winter 2024 Newsletter, doi: https://www.ecmwf.int/en/newsletter/178/earth-system-science/improved-two-metre-temperature-forecasts-2024-upgrade

Mahfouf, J.-F., K. Bergaoui, C. Draper, F. Bouyssel, F. Taillefer, and L. Taseva (2009), A comparison of two off-line soil analysis schemes for assimilation of screen level observations, J. Geophys. Res., 114, D08105, doi:10.1029/2008JD011077.

Massart, S. (2023), Extended control variable (XCV): skin temperature background correction for the assimilation of clear-sky microwave radiances, ECMWF Technical Memorandum, DOI: 10.21957/4bbf1919d5.

Penny S.G., Akella S., Alves O., Bishop C., Buehner M., Chevallier M., Counillon F., Draper C., Frolov S., Fujii Y., Karspeck A., Kumar A., Laloyaux P., Mahfouf J.-F., Martin M., Peña M., de Rosnay P., Subramanian A., Tardif R., Wang Y., Wu X.: Coupled Data Assimilation for Integrated Earth System Analysis and Prediction: Goals, Challenges and Recommendations. World Meteorol. Org. (WMO), WWRP 2017-3, 2017. https://library.wmo.int/doc_num.php?explnum_id=10830

Rabier, F., Järvinen, H., Klinker, E., Mahfouf, J.-F. and Simons, A. (2000) The ECMWF operational implementation of four-dimensional variational assimilation. I: Experimental results with simplified physics. Quarterly Journal of the Royal Meteorological Society, 126(564), 1143–1170. https://doi.org/10.1002/qj.49712656415

Document History

Version	Author(s)	Date	Changes
1.0	Christoph Herbert, Peter Weston, Patricia de Rosnay, Sébastien Massart and Massimo Bonavita	24/11/2025	Initial version
1.1	Christoph Herbert, Peter Weston, Patricia de Rosnay, Sébastien Massart and Massimo Bonavita	26/11/2025	Revised version after internal WP review
1.2	Christoph Herbert, Peter Weston, Patricia de Rosnay, Sébastien Massart and Massimo Bonavita	Dec 2025	Revised after internal project review

Internal Review History

Internal Reviewers	Date	Comments
Giovanni Conti (CMCC) and Xiaohua Yang (DMI)	Dec 2025	Comments on v1.1